Detecting Bidders Groups in Collusive Auctions

Timothy G. Conley and Francesco Decarolis*

January 4, 2012

Abstract

This paper studies entry and bidding in procurement auctions were contracts are awarded to the bid closest to a trimmed average bid. We characterize equilibrium under competition and show that it is weak due to strong incentives for cooperation. We present statistical tests motivated by a model of coalition entry and bidding. We show that our tests perform well in a validation dataset with known cartels. We also use them to investigate cooperation in a larger, more representative dataset where cartels are suspected but not known. We detect several suspiciously cooperative groups with potentially substantial, *positive* effects upon auctioneers' revenues.

JEL: DL22, L74, D44, D82, H57.

*Conley, University of Western Ontario. *E-mail address:* tconley3@uwo.ca. Decarolis, University of Wisconsin - Madison. *E-mail address:* fdc@ssc.wisc.edu (corresponding author). The authors would like to thank their colleagues at their institutions as well as Susan Athey, John Asker, Lanier Benkard, Marianne Bertrand, Peter Cramton, Jeremy Fox, Jakub Kastl, Jonathan Levin, Greg Lewis, Gustavo Piga, Rob Porter, David Rivers, Pierluigi Sabbatini, Jesse Shapiro, Giancarlo Spagnolo and Charles Zheng for the useful suggestions. We thankfully acknowledge the aid with the data collection of the Bank of Italy, the Italian Authority for Public Contracts and the Legal Office of the municipality of Turin.

"*....At the first meeting they said: "Why should we kill ourselves and make laugh those coming from outside?" Here (*i.e., in Turin*) firms from the South were coming and getting the jobs, getting the averages, they used to came with 20, 30 or 40 bids, they used to get the jobs and then what was left for us?...*" (Confession of Bruno Bresciani, found guilty of having rigged 94 average bid auctions and other related crimes; convicted to 7 years of jail in April 2008)

1 Introduction

In recent years, economists have contributed to designing new auction markets for activities ranging from the sale of spectrum licenses for mobile operators to that of electricity supply contracts. However, the extent to which these auctions can deliver the intended results crucially depends on how bidders respond to strategic incentives. In this paper, we present the case of a large auction market for the procurement of public works in Italy and show the sophisticated response of bidders to the incentive to coordinate entry and bidding to rig the mechanism. We also introduce two statistical tests that work well at detecting groups of cooperating firms and that could be applied to other markets presenting similar incentives.

The auctions that we study are called 'average bid' auctions (ABAs) and have been used in Italy since 1999 for the procurement of the vast majority of public works. Similarly to the procurement auctions used in various other countries listed in Table 1, their main feature is that the winner is decided through an algorithm that eliminates all bids that are deemed 'too good to be true.' In the procurement of works, there is often concern that in first price auctions adverse selection can lead to a trade-off between a high winning discount and a low ex post performance. To better understand how the awarding of contracts works under ABA, suppose that a public administration (PA) announces that it is willing to pay up to a certain reserve price to have some public work executed. Firms submit bids in the form of discounts on this reserve price. In a standard first price auction (FPA), the highest discount wins. Instead, the rule in place in Italy uses an algorithm to exclude all the discounts that are above a certain threshold related to the average of the discounts. The firm with the highest non eliminated discount wins and is paid its own bid to perform the work. Although ruling out bids that may be too good to be true reduces the risk of ex post default by firms, bidders' incentives are deeply distorted relative to a standard FPA.

Auctions like these are 'collusive auctions' in at least two senses. First, since the highest discount is always eliminated, there will be one equilibrium in which all firms submit the lowest possible discount, zero, i.e. they bid the reserve price. In this case, the auctioneer procures the contract at the worst possible terms from her perspective, exactly the same result as when all bidders form a single coalition. Second, the fact that the awarding rule is based on a function of the average bid implies that a coalition of firms can manipulate the awarding process by using multiple bids to pilot the relevant threshold. Our analysis reveals that in the environment that we study, firms do indeed form coalitions that coordinate the entry and bids of their members. This incentive to cooperate is at the heart of our study, which addresses it from both a theoretical and an empirical perspective.

We begin by formulating a model of firms' entry and bidding decisions. In our model, if firms act independently, the unique equilibrium is for all firms to offer a discount of zero over the reserve price. Since the data reveals a very different behavior, we enrich the model allowing for the presence of coalitions of firms. Relative to independent firms, firms within the same group should be more likely to enter together and to bid on the same side of the bid distribution. The latter action serves to tilt the awarding threshold toward the portion of the bid distribution where the group's bids lie.

We then develop two statistical tests, one for entry and one for bids, to capture these behavioral features. Our entry test is motivated by the simple idea that, in order to bid in a cooperating manner, a group of firms must be present together in sufficient numbers in an auction. Therefore, a group of cooperating firms may be detected by comparing its entry and bidding behavior to a comparable set of control groups' behavior. We compare the frequency of joint entry between a suspect group and a collection of randomly chosen groups whose members are comparable to firms in the suspect group. Our bid test is motivated by a search for groups employing a cooperative strategy to pilot the trimmed average bid toward one of their members. We exploit the exact rules of these AB auctions to construct a test statistic tailored to measuring the extent to which a given groups' bids move the threshold that determines a winner. We then compare this measure of 'trimmed mean piloting' for the suspect group versus a control group comprised of a randomly selected groups of comparable bidders. The repeated nature of the auctions may enhance the power of the test even when there are multiple groups that are simultaneously trying to rig the same auction. However, testing with multiple auctions has an important set of caveats that we discuss in the text.

When we apply our tests to known groups of cooperating firms, our tests perform well in detecting these groups. In particular, we use 276 ABAs for roadworks held by the city of Turin between 2000 and 2003. We refer to these auctions as the Validation data. In 2008, the Turin's Court of Justice ruled that these auctions had been rigged by 8 groups made up of 95 firms.[1] Each group strategically submitted bids to affect the awarding of the contract. According to Italian law this activity is a crime. These groups were labelled cartels[2] and their members fined, some of them even sentenced to jail. For our purposes, this is an ideal sample to validate our tests because we can check whether the tests are able to identify the 8 cartels sanctioned by the court. The results that we obtain are very much supportive about the capacity of our tests to correctly detect cartels. Of the 8 cartels, the only one for which we do not find systematic evidence of cooperation is the one that the court sanctioned less because its members rarely coordinated bids.

We then turn to the problem of detecting groups in auctions where we have no prior knowledge of their presence. We look at a dataset of 802 ABAs held in the North of Italy between 2005 and 2010. We refer to these auctions as the Main data. Many of the observed features of these ABAs resemble those of the ABAs in the Validation data. Given the large number of firms in these auctions, we suggest various ways to reduce the set of firms

[1]The decision took into account some of the empirical evidence that we will discuss as well the confessions of some ring members and the phone calls and emails intercepted by the police.

[2]We will follow this sentence in referring to these 8 groups as cartels although our discussion of Harrington (2011) clarifies that they are cartels in the legal but not in the economic sense of this word.

to analyze. Our favorite method constructs candidate groups starting from the network of relationships connecting firms along various observable dimensions: overlaps in the identities of owners and managers, the exchange of subcontracts, the formation of temporary bidding consortia and geographical proximity. Using groups constructed in this way, we can then apply our tests. Based on these tests, we detect numerous groups of firms that appear to be engaging in the coordination of their bids and entry. In particular, our favorite conservative estimates suggest that these groups affect no less than 30% of the auctions. We then argue that this cooperative behavior likely produced large savings for the auctioneer. This unusual claim is made relative to the benchmark competitive case in which all firms offer a discount of zero and implies savings for the auctioneer of about 13% of the reserve price on average. The only harm occurs to bidders outside the groups that are less likely to win and, when they win, they do so at a worst price than under competition. Finally, we present a clear illustration of how quantitatively important is bidders' reaction to the change of the auction incentives. We analyze a change in the regulation that replaced the ABA with the FPA for certain types of contracts producing the exit from the market of hundreds of firms. We use our tests to classify exiting firms between those belonging to groups and those not. Our findings suggest that, among the 774 firms whose exit from the market appears associated with the introduction of the FPA, 159 of them (or 21%) belong to groups, while the others are independent. For the latter firms, exit is most likely driven by their inefficiency which does not allow them to effectively compete in FPA. For firms in the groups, instead, it is ambiguous whether a firm exits because it is inefficient or because it is a shill of some other group member. Shills are those firms set up by another firm for the only reason to have more bids to manipulate the ABA and, clearly, are useless in the FPA. In the court case, it is argued that some of the convicted firms possibly played the role of shills. Often, they were formally owned and managed by wife, sisters and mother of the male owner of the "real" construction firm. Although we can rely only on incomplete data to assess the gender of the firms owners and managers, we find some weak evidence of more female owners and managers across the 159 exiting firms belonging to groups relative to both other exiting and non exiting firms.

Our paper has four main contributions. Its most direct contribution is to develop statistical tests that can help courts evaluating cooperation in ABA or, with slight modifications, coordination mechanisms with similarly manipulable awarding rules. In public procurement, examples of such mechanisms abound. For instance, in Medicare both the system of low income subsidies in Part D and the DEMPOS auctions for durable medical equipment have manipulable thresholds. The second direct contribution is to show that ABA is conducive to the formation of groups. These auctions receive very little attention in the literature but they exist in numerous countries listed in Table 1 and, among them, they are of major economic relevance at least in Colombia, China, Italy and Japan. The third and more general contribution is that by clearly showing that firms' response to the auction incentives is both highly sophisticated and quantitatively very large, we contribute to a growing literature that advocates the use of accurately designed auctions for public procurement. Finally, our results are useful for the design of antitrust regulations. In particular, as discussed in Harrington (2011), there are differences between the legal definition of collusion, which typically refers to every action that firms take to coordinate prices, and the economic one, which typically means that the price that results from the coordination of the firms' actions is above the one achievable under competition. In this respect, we present a striking case in which the legal and economic definitions of collusion lead to totally different evaluations of the damages caused by bidders' cooperation to the auctioneer's revenues. Therefore our paper argues against the usage of automatic sanctions punishing all types of cooperation and in favor of a careful economic analysis of the markets.

Literature: By studying bidders' cooperation, this paper is most closely related to the literature on collusion in auctions. In auction design, collusion is generally regarded as a first order concern (Klemperer, 2004) and has received substantial attention from the theoretical literature.[3] Our study, instead, contributes to the much smaller empirical literature on collusion in auctions. This literature can be roughly divided into two groups: the studies of collusion practices in markets where the presence of cartels' existence has been proved by court (Asker, 2010, Pesendorfer, 2000, Porter and Zona, 1993 and 1999) and the studies that try to devise methods to distinguish competition from collusion in environments where the presence of collusion is only a possibility (Bajari and Ye, 2003). Both approaches have led to the flourishing of a literature within industrial organization concerned with "screens for collusion" (i.e., statistical tests to detect collusion, see the review by Abrantes-Metz and Bajari, 2010). In this paper, we take an intermediate approach: we use information from auctions where collusion was proved, but we do so in order to devise an empirical methodology that allows assessing the likelihood of groups in markets where their presence has not yet been proved. The motivation of our approach is based on the idea of Hendricks and Porter (1989) who explain that collusion is tailored to the specific rules of the auction and the environment. Therefore, we use data from auctions with collusion to learn about the behavior of groups and then search for evidence of this behavior in other similar auctions.

Finally, our analysis also contributes in two ways to the large literature on public procurement auctions. First, it contributes to the study of collusion in public procurement auctions (a review of cases for the US is contained in Haberbush, 2000) and, more specifically, in auctions for roadwork jobs (similarly to Porter and Zona, 1993 and Ishii, 2006). Secondly, it contributes to the study of mechanisms similar to the ABA, which are little studied in economics but widely used in practice. These auctions were introduced by civil engineering literature.[4] However, there are few theoretical results on the effect of rules similar to ABA: Albano et al. (2006), Engel et al. (2006), Katzman and McGeary (2008), Decarolis (2011) and Cramton et al. (2011). The latter study also presents the results of a lab experiment involving the auction by Medicare to procure medical equipment.[5] Instead, Decarolis (2012) analyzes a different part of Medicare, Part D. His analysis extends the statistical tests de-

[3]The seminal studies in the theoretical literature include Robinson (1985) addressing the strength to collusion of first price relative to second price auctions and the studies on cartels' behavior in second price or English auctions (Graham and Marshall, 1987, and Mailath and Zemski, 1991) and in first price auctions (McAfee and McMillan, 1992). Recent work on collusion and auction design is Marshall and Marx (2006).

[4]Ioannou and Leu (1993) and Liu and Lai (2000) are the original studies proposing such mechanisms to address the trade-off between awarded price and ex post performance caused by FPA. Instead, for the different solution to this problem proposed by the economics literature see Spulber (1990).

[5]This is the multiunit-median bid auction that Medicare uses for the procurement of durable medical equipment (known as DEMPOS auction).

veloped in our paper to study how firms price their Part D insurance plans to manipulate the public subsidy that they receive.

The outline of the paper is as follows: the next Section provides a description of the market and our data sources, Section 3 presents a model of firms' entry and bidding, Section 4 presents our econometric tests and investigates their performance on the Validation data, Section 5 discusses the case of testing with no prior knowledge about groups, Section 6 illustrates the results obtained by applying the tests to the auctions in the Main data and, finally, Section 7 concludes.

2 Description of the Market

In this Section, we describe both the institutions and the datasets that we will analyze. We study auctions held by Italian public administrations (PAs) to procure contracts for simple roadworks in Northern Italy. We are motivated to study these auctions in particular because for the PA of Turin we have access to what we call a Validation dataset as a result of legal cases where sets of firms were convicted for collusion in these auctions. These Validation data auctions are comparable in other aspects to the remainder of our data, which we refer to as our Main dataset.

For these roadwork contracts, PAs are typically required to select the contractor through sealed bid price-based auctions.[6] A small fraction of these auctions are of the well known first price auction (FPA) type, but the vast majority are 'averaged-average with tail trimming' auction, to which we will refer as average bid auction (ABA). The regulations of ABA and FPA are identical in everything except for how the winner is identified. In both cases, the PA announces a job description and a reserve price that is the maximum it is willing to pay. Then firms submit sealed bids consisting of discounts on this reserve price. However, while in FPA the highest discount wins, in ABA the winner is found as follows: a) bids are ranked from the lowest to the highest discount; b) a trim mean (A1) is calculated disregarding the 10 percent of the highest and lowest discounts; c) a new mean (A2) is calculated as the average of those discounts strictly above A1, disregarding those discounts excluded for the calculation of A1; d) the winning discount is the highest discount strictly lower than A2. Ties of winning discounts are broken with a fair lottery. The winner is paid his bid to perform the work. Figure 1 offers an example with 17 bids: the winner is denoted D^{Win} and, in this case, it is the 7th highest discount.

The ABA described above was introduced in 1999 and, until June 2006, it was the compulsory mechanism for the procurement of almost all contracts with a reserve price below €5 million. In this period, approximately 80% of all the contracts for public works were procured by PAs using ABA, resulting in a total reserve price of the auctioned contracts of approximately €10 billion. Starting from July 2006, a series of reforms required by the European Union removed the mandatory use of ABA and extended the use of the FPA. However, even after these reforms ABA remained the most frequently used procurement

[6] A detailed discussion of the regulation is contained in Decarolis (2011) and Decarolis et al. (2010).

format.[7] In this paper, we do not consider the auctioneer's problem of choosing among different auction formats, but we focus on ABA to study bidders' behavior in this format.

A) Main Data

Our Main data contain 1,034 auctions held by counties and municipalities between November 2005 and May 2010. All auctions involved the procurement of simple roadwork contracts (mostly paving jobs, worth below €1 Million) and were held in five regions of the North of Italy (Piedmont, Liguria, Lombardy, Veneto and Emilia-Romagna). The choice of the sample is motivated both by the relevance of these contracts, which are the most frequently procured public works, and by the need to assure the comparability of the auctions, despite the fact that they were held by different PAs and at different points in time. This comparability seems confirmed by the fact that we observe a substantial fraction of firms bidding repeatedly both over time and across auctions of different PAs.

As discussed above, a switch from ABA to FPA was gradually ongoing during the period we study. Our Main data consists of 802 ABAs and 232 FPAs. Figure 2 shows the geographical distribution of the 802 ABAs. Table 2 presents some summary statistics separately for the two types of auctions. Comparing the statistics for the two sets of auctions reveals several differences in terms of bidders entry and bidding. As regards entry, the number of bidders is several times larger in ABAs than in FPAs: on average there are 7 bidders in an FPA and 51 in an ABA. As regards bidding, the winning discount is on average 13 percent in an ABA, while it is 30 percent in an FPA. Moreover, in ABAs there is substantially less within-auction variation in the bids than in the FPAs: this is shown by both the lower within-auction standard deviation of bids and the lower difference between the winning discount and the next highest discount in the ABAs relative to the FPAs. This latter variable, sometimes defined as 'money left on the table' is on average 4.5 percent of the reserve price in an FPA but only .2 percent in an ABA. Finally, in the right panel of Table 2, we report summary statistics for the bidders. There are approximately 4,000 firms that bid at least once. They exhibit strong asymmetries both in their characteristics (like capital) and in their performance in the auctions (like the number of victories). Although we do not report the data broken down by the format in which the firms participate, on average the firms bidding in FPAs have a larger size (in terms of capital and number of employees) and are located closer to the work area.

B) Validation Data

The ABAs in the Validation data were collected by the legal office of the municipality of Turin as part of a legal case against several firms accused of having committed auction rigging. This dataset consists of 276 ABAs held by the municipality of Turin between 2000 and 2003 to procure roadwork jobs. There is a substantial overlap of bidders among the Main and Validation data which underscores the comparability of the ABAs in the two datasets. On April 2008 the Court of Justice of Turin convicted the owners and managers of numerous

[7] Decarolis (2011) uses these changes in regulation to study the effects of a switch from ABA to FPA on various auction outcomes, like the winning bid, measures of performance (delays and cost overruns), entry and subcontracts. Our study is complementary to Decarolis (2011) in the sense that we provide a foundation for many of his results based on firms' cooperation in ABA.

construction firms. The court documents identify a network of 95 firms that operated in 8 cartels.[8] We use the term cartels to follow the court terminology and to better distinguish these 8 groups from the candidate groups of cooperating firms in the Main dataset. These cartels were very successful in their activity. Despite representing no more than 10 percent of the firms in the market, they won about 80 percent of all the auctions held in the Piedmont region between 2000 and 2003. Cartels were formed mostly by firms geographically close to each other and to Turin. This is unsurprising as proximity to other group members is plausibly related to lower costs of coordinating actions and of exchanging favors.[9] Proximity to Turin surely provides cost advantages for execution of road construction contracts. Figure 3 shows the geographical location of the cartels and indicates each with a capital letter, from A to H, that we also use in Table 3 and throughout the remainder of the paper. Two cartels, G and H, despite having all members close to each other, are located far from Turin. According to the court decision, these cartels did not want to win the auctions to perform the jobs, but only to resell them through subcontracts. Finally, as Table 3 shows, these 8 cartels are quite heterogenous in their size, entry and victories.

In addition to the asymmetries across cartels, there are also significant asymmetries within cartels. The bottom panel of Table 4 reports summary statistics for both the firms inside and outside the cartels. Given that this sample was assembled to compare alleged colluders with independent firms, it is not surprising to see that all variables measuring outcomes of the auctions (entry, victories, subcontracts, etc.) take larger values for the members of the cartels. As regards the auctions themselves, the top panel of Table 4 suggests that these auctions are similar to those in the Main data described in Table 2 on the basis of entry and of dispersion of the bids. Interestingly, the average winning discount is higher in these 'colluded' auctions than in those of Table 2, 17.4% compared to 13.7%.

C) Descriptive Evidence on Firms' Behavior in the Two Datasets

The importance of the Validation data is that for its auctions we have a rather clear idea of what firms were doing and why. Indeed, several of the persons involved in the agreements made confessions to the court in an attempt to reduce their sentence. Moreover, phone calls and emails where recorded by the police for almost three years and portions of these conversations became publicly available with the sentence. The picture that emerges describes a complex environment in which cartels compete against each other (although in some occasions some of them form short term agreements) and against numerous independent firms. Three specific features of both bidding and entry emerge clearly.

C.1) Predictable Winning Bid Range

The first feature of the bid distributions is that a basic range for winning discounts is

[8]Turin Court of Justice, 1st criminal section, April 28th, 2008, sentence N. 2549/06 R.G.. Of the 95 suspect firms, the sentence convicts 29. Proscription lead to the acquittal for 2 firms. The judgment of the other firms was decided in different court cases. In our study we consider the full network of 95 firms.

[9]Porter and Zona (1993) suggest various reasons for why cartels emerge in the type of market studied in this paper: (1) bids are evaluated only along the price dimension and so product differentiation is absent; (2) firms are relatively homogeneous because of the similar technology and inputs; (3) every year there are many auctions and they take place quite regularly; (4) there are legal forms of joint bidding; (5) the same firms repeatedly interact, (6) ex post the auctioneer discloses the identities and bids of all bidders. These six reason likely played an important role for both the formation and the stability of the groups that we study.

predictable across auctions within a PA. The winning bids are almost always near the approximate mode of the bid distribution, which in the Validation data is around 18%. In our Validation data, court documents confirm this as some convicted firms revealed that it was known to all players in this market that most bids were always placed around a known likely range for the winning bid. Figure 4 illustrates this for one Validation data auction. Individual bids are plotted in increasing order with discounts on the vertical axis. There is a clear approximate mode in the distribution around a discount of 18% with the winning bid highlighted by the thick line on the edge of this mode. Auctions for this PA within a year of this auction have very similar approximate modes and winning bids are consistently near these approximate modes. This basic pattern occurs in the Main dataset as well. This evidence about predictability of modes and range containing winning bids is confirmed by accounts given by market participants about firms' bidding policies and is consistent with the large amount of information about past auctions available to bidders.[10] In fact, Decarolis (2011) finds that, across PAs in our main dataset, in their ABAs there is a strong tendency for the winning bids to remain nearly identical across the auctions of the same PA regardless of the type of job, reserve price, duration of the work, etc. There is variation across PAs in these approximate modes and winning bid ranges.[11]

As we will show in the next Section, these empirical regularities are relevant for us in four main ways. The fact that the mode and winning discounts are substantially greater than zero is evidence against firms acting independently and in favor of there being cooperative groups. The predictability of the range for winning bids motivates our treatment of independent firms as being able to predict the range where winning bids will lie and choosing to confine their bids to this range. This predictability of the winning bid range is also consistent with a cooperative strategy where a subset of a group of collaborating firms pilots the trimmed mean towards another member's bid. Finally, the similarity of the bid distribution modes and ranges for winning bids across auctions provides some reassurance that a common equilibrium is being played in the auctions we pool in our datasets.

C.2) Average-piloting Bids

The second feature about bidding is that, despite the fact that most bids are typically in a range near the winning discount, there are often some extremely high and/or low discounts. The explanation offered in the court documents is that sometimes bids are not placed to win but to pilot the average. The primary mechanism identified by the courts used for rigging ABAs is in fact sets of firms bidding to pilot the trimmed means determining winners towards their cartel members' bids. The bidders themselves refer to these very high/low bids as 'supporting bids' because they are too extreme to have any chance of winning the auction, but can help a connected firm to win. In Figure 4, the nine highest discounts illustrate well the idea of supporting bids. Recall that the vertical axis is the discount offered while the horizontal axis lists the bidders in an increasing order of their discounts. Different symbols

[10] The sources of information are both public, as the regulation mandates the publication of the auction outcomes on the notice board of the PA, and private, as a very competitive market exists for firms collecting and reselling information about public auctions. Coviello and Mariniello (2011) provides a detailed account of the disclosure of information in these auctions.

[11] The different promptness and accuracy with which different PAs release information about past auctions (documented in Coviello and Mariniello, 2011) is likely one of the reasons for these differences across PAs.

indicate different cartels with the cross representing firms not in groups. The majority of discounts are near the 18% approximate mode. However, several members of the cartel, represented by a circle, submitted discounts that are 'discontinuously' greater than those of all other bidders. In this case, their strategy was successful in making a member of their coalition win the auction (the thick blue line). Numerous similar cases are present in the Validation data. Moreover, numerous extreme discounts suggesting a clear piloting of the awarding threshold are present also in the Main dataset. It is routine for there to be clusters of bids in the tails of the distribution separated by a substantial distance from the bulk of the bids.

C.3) Entry of Connected Firms

The last relevant behavioral feature regards joint entry of firms. It is illegal for two firms sharing the same majority shareholder to submit bids in the same auction. However, the case in Turin reveals that entry by closely connected firms was common. Several of the firms composing the 8 sanctioned cartels shared some shareholders but always entered auctions together. Moreover, some of them also shared managers, ownership by members of the same family, registration at the same street address, or they systematically exchanged subcontracts. Since we observe all these characteristics for the firms in the Main dataset, we know that in both datasets it is extremely common to find several closely connected firms entering the same auction. In Section 5 we show how to exploit these observable links between firms to construct candidate groups of cooperating firms.[12] Sometimes the connections between firms in the Validation data were so strong that the court argued that some firms could have been considered shills of some other firm in the same cartel: firms existing for the sole purpose of allowing the original firm to place multiple bids. However, not even the court could convincingly identify which firms were shills because that requires observing a counterfactual environment where firms do not gain from having multiple bids. In Section 6 we explore this issue in greater detail, but for most of our analysis it will be convenient to think of a group as a collection of different firms that delegate their actions to a common mediator.

Overall, common features between the Main and Validation data discussed in this Section strongly suggest they are comparable and lessons learned from our Validation data will be valuable in analysing the Main data.[13] We begin our analysis introducing a model of group behavior.

[12]The distribution of these firm linkage variables is quite similar in the Validation and Main data sets. For both subcontracting and the three variables measuring ownership, management and white collar workers, both rank sum and t-tests comparing means fail to reject at the 5% a null of equal distributions.

[13]One additional reason for looking for groups in the Main data is that the criminal courts of two Northern cities (Treviso and Vicenza) recently started cases similar to that of Turin. Other trials are occurring also in the South. The Italian Antitrust Authority also expressed concerns about the risk of collusion posed by mechanisms similar to the ABA when they were still a new mechanism used by few PAs, see AGCM, 1992.

3 Model of Participation and Bidding

This Section presents a model of firms' entry and bidding decisions. Its implications are used to develop our tests for coordinated entry and bidding. We assume that a single contract is auctioned off. Firms are either independent or part of groups. We model the decision problem of independent firms in two stages: in the first stage firms observe their cost for preparing the bid and then, in the second stage, those firms that decided to pay the preparation cost learn their cost of completing the job and then bid. Entrants know both that N bidders will enter and whether some will act in groups. We abstain from modeling the inner working of these groups. Instead, we assume that a group is a collection of firms that delegate to a common mediator their entry and bidding decisions in exchange for a share of the group's joint profits. The mediator observes their costs and acts to maximize their joint profits.

Characterizing the behavior in the bidding stage

In the auction, a subset of cardinality N of the potential bidders places a bid. There are N^I independents and N^g firms in group g, with $g = 1, ..., G$. We assume throughout the paper that $N > 4$. Each firm j has cost c_j of completing the job. Assume that $c_j \in [c^l, c^h]$ for all $j = 1, .., N$ and that each entering firm draws its cost from a distribution $F_C(.)$ that is absolutely continuous. Before bidding, firms also observe the maximum price, R, that the auctioneer is willing to pay (the reserve price). This price is non-binding for all firms. A firm submits a sealed bid, $b \in [0, 1]$, consisting of a discount over R. Therefore, the expected profit for an independent firm j that enters and bids b_j is: $E_I(\pi(b_j)) = [(1 - b_j)R - c_j] \Pr(b_j$ *wins*). The winner is determined according to the Italian ABA described in the previous Section, thus b_j wins if it is the highest discount strictly below $A2$.[14] We begin the analysis by looking at the case in which there is no group.

Proposition 1: Without groups, there is a unique Bayesian Nash equilibrium (BNE) in which all firms bid a discount of zero percent (zero-discount equilibrium).

All other proofs are in the Appendix. However, for this proposition it is straightforward to see that zero-discount strategy profile is an equilibrium: any unilateral deviation leads to a zero probability of winning, while bidding zero gives a $1/N$ probability of winning and making a profit (since the reserve price is non binding). Uniqueness, instead, is due to two details of the ABA rule: (i) the winner has the highest discount strictly lower than $A2$ and (ii) the discounts disregarded to calculate $A1$ are not eliminated, unless they are also greater than $A2$. Thus, no bidder wants to submit the single highest discount. However, even if all discounts are identical, a single bidder will have a profitable deviation unless they are all zero. If N-1 discounts are equal to some $b > 0$ and one discount is equal to zero, this latter discount wins because it is the closest from below to $A1 = A2 = b$. Uniqueness of the equilibrium is a feature of the detail of the Italian rule. Indeed, an alternative rule

[14] For readability we report again the exact rule: the discounts' trim mean, A1, is computed as the average discount disregarding the highest and lowest 10 percent (rounded to the highest integer) of discounts; then A2 is calculated as the average of the discounts greater than A1 disregarding the discounts excluded for the calculation of A1; the discount closest from (strictly) below to A2 wins. The winner is paid his own price and ties of winning discounts are broken with a fair lottery. If all bids are equal, the winner is selected with a fair lottery. A tie of winning bids is broken with a fair lottery.

stating that the winner closest to the average wins can have a multiplicity of equilibria. However, these equilibria would also have the feature that all discounts would be identical. This pooling of discounts implies that the allocation produced is inefficient. Moreover, in the unique equilibrium of the ABA the auctioneer's revenues are minimized.[15] These results are an artifact neither of bidders symmetry nor of the presence of a known number of bidders since relaxing these assumptions leaves Proposition 1 unaltered.[16] However, a coalition of firms can break this equilibrium.

Proposition 2: The zero-discount strategy profile is not an equilibrium unless all bidders are independent, or they all belong to the same group, or all groups are smaller than the minimum breaking coalition (which is 2 plus 10% of N rounded to the next highest integer).

The minimum breaking coalition denoted N^* and defined in Proposition 2 is the smallest group of bidders that can break the zero-discounts equilibrium. If all firms are bidding zero, then a coalition of N^* firms can submit all discounts strictly greater than zero (for instance $N^* - 1$ discounts equal to $\varepsilon > 0$ and one equal to $\varepsilon/2$) and win for sure. Although the winning firm receives a lower payment form the auctioneer, there is always an ε small enough to make this strategy strictly more profitable for the group than bidding zero and winning with probability $1/N$.[17]

The insight from Proposition 2 is that, built into the ABA, there is a powerful incentive to induce firms to form groups.[18] In FPA, a coalition of bidders will stand a good chance of profiting by deviating from the competitive equilibrium only if it contains the bidders with the lowest cost draws. Instead, the incentive to cooperate is stronger in ABA since (by Proposition 2) any coalition, regardless of its members' costs, can profitably break the zero-discounts equilibrium. However, once we allow for the possibility of group bidding, the characterization of the full set of equilibria becomes challenging because of the complexity of the strategy space. Nevertheless, from the previous Section we know that the whole market exhibits some clear characteristics: auctions have a large number of bids and most of these bids are concentrated around an easily predictable range. Therefore, instead of trying to characterize the whole set of equilibria we assume that it is common knowledge that a large number of bidders will bid close to a known range and ask whether this behavior can be rationalized as an equilibrium of a game where groups are active. The following proposition gives a positive answer to this question. To begin, let us indicate by b_f as a common forecast bid around which most bids are placed. Formally, define b_f^I as a profile of discounts for the independent bidders such that: all their discounts are in $[b_f - \eta, b_f]$ with $b_f \in (0, 1)$ with

[15]Decarolis (2011) shows that if cost uncertainty persists after bidding and defaults are possible (at a privately observed cost), the ABA can dominate the FPA in terms of revenues. Since in this study we are not concerned with the auctioneer's choice of mechanisms and defaults never happen in our data we ignore these considerations.

[16]The key assumption driving equilibrium uniqueness is the absolute continuity of the cost distribution. Examples of multiple equilibria can be found with discontinuous distributions. Asymmetric bidders examples that replicate such discontinuity can also produce multiple equilibria.

[17]Instead, if all groups are smaller than N^*, then no individual firm or group has an incentive to bid more than zero. Additionally, if all firms belong to a single large group, the winning bid must be equal to zero.

[18]The process through which a group of distinct firms forms differs from that of a firm creating a group with its own shills. We take groups as given and do not model the formation stage.

$b_f > \eta$ and $\eta > 0$ but small. To simplify the analysis, we assume that there is a single group.

Proposition 3: For any b_f^I and N^I and any small $\varepsilon > 0$, there are two values, N^{g*} and N^{g**}, $N^{g**} \geq N^{g*}$, such that if the group size is at least N^{g*} but less than N^{g**}, then there is an ε-equilibrium[19] in which all group's discounts are clustered above b_f. Instead, if its size is at least N^{g**}, there is an ε-equilibrium in which all group discounts are clustered below $b_f - \eta$.

This proposition says that if many independents bid within a narrow interval $[b_f - \eta, b_f]$, then a group that is not too large, in between $N^{g*}$and N^{g**}, will cluster its discounts above this range. By clustering discounts on the higher side of the discount distribution, the group increases its chance to win by moving the interval containing the winning bid, $[A1, A2)$, toward the side of the distribution where its bids are located. In the proof presented in the appendix, it is shown that the group strategy entails bids randomization. This is essential to avoid that the independents outguess where the group will push $[A1, A2)$. By clustering and randomizing together its bids, the group makes the probability of an independent winning by deviating from b_f^I arbitrarily small. Given b_f^I, it might seem intuitive that a group would prefer pushing downward the interval $[A1, A2)$ and winning with a discount lower than b_f. However, for such a strategy to be an equilibrium the group needs to have a rather large size, i.e. greater than N^{g**}. This is because moving A2 upward typically requires less bids than moving it downward (see the appendix for the exact details). Finally, our last proposition further differentiates the bidding behavior of groups and independents.

Proposition 4: Assume that $[b_l, b_h]$ with $b_h > b_l \geq 0$ is the bids support resulting in an equilibrium with at least one group, then there cannot be any Bayesian-Nash equilibrium in which an independent firm bids b_h. Instead, only a group member can place b_h.

The idea is that b_h is a strictly dominated discount for an individual firm but for a group it can serve to tilt the average upward. In this case, clustering of all group discounts on the high end of the distribution should occur as otherwise b_h is weakly dominated by any lower discount. The argument does not necessarily extend to b_l as b_l might be the only individually rational discount to cover high production costs. However, the argument also applies to b_l if the winner is allowed to resell (part of) the contract, for instance via subcontracts.

When multiple groups are active and many independents bid close to b_f, there is still an incentive for each of them to manipulate the average by moving the interval $[A1, A2)$ toward the group's own bids. Thus, clustering strategies maintain their intuitive appeal. Analogously, the incentive to randomize in order to make the location of $[A1, A2)$ unpredictable to rivals is even stronger since now a group that always clusters on a single side of the distribution will be easily outguessed by the other group(s). This is particularly relevant given the repeated nature of the auctions that we observe. In this scenario, the incentive

[19] In an ε-Nash equilibrium (Radner, 1980) no player has a deviation leading to a gain greater than ε. This is a full information concept and we can apply it because, by focusing on a situation in which the social norm is to center discounts around a known range, we are implicitly assuming that doing so is profitable. This implicit assumption is justified by two findings in Decarolis (2011). He shows that when FPAs replaced ABAs: (i) the winning discount more than doubled and (ii) the value of subcontracts declined by a third. Thus, in ABA the winning discount is much less than what the most efficient firm would be willing to offer and, even if an inefficient firm wins, it will resell part of the work via subcontracts.

for independents to always bid close to b_f resembles buying a lottery ticket that can pay off because the different groups do not coordinate in how they push the location of $[A1, A2)$. Ex ante, to each independent all values within a small ε of b_f are approximately worth the same and independents can be thought of as mixing within this interval. More intuitively, by the purification theorem this behavior can result from playing a pure strategy based on their privately observed cost. Therefore, since independents draw their cost independently from the same distribution, their bids will also be independent. Hence, knowing that the bid of the independent firm i, b_i, is on the high (low) part of the distribution does not provide any information on the location of the bid of another independent firm j, b_j, is. If instead, i and j are in the same group, then knowing that b_i is on the high (low) end of the distribution makes it more likely that b_j is on that side too. In the next Section, our bid test is based on discovering groups of firms clustering their bids to pilot $[A1, A2)$.

Characterizing the behavior in the participation stage

The above discussion clarifies how a group can improve its expected payoff by coordinating a sufficient number of bids in an ABA. From the court case in Turin, we know that groups play complex participation strategies sometimes involving bribing firms to bid with them for a single auction. However, since our data is not rich enough to measure precisely this phenomenon, we abstain from modeling the inner working of the groups' entry choices. Instead, we simply assume that a group's mediator trades off the benefit of one additional bid to manipulate the average against the cost of the additional bid preparation. The meaning of this latter cost is likely different for independents and group firms since for group firms it consists of a higher probability of being sanctioned by a court. Nevertheless, we expect that a firm in a group is more likely to enter if the other $N^* - 1$ firms from the same group enter because of the greater expected payoff. For the independent firms, instead, we assume that every firm j independently draws a participation cost $q_j \sim F_Q(.)$. Thus, if we define the expected profit before independent firm j observes its cost as $E_I(\pi_I)$, then j follows the cutoff rule: enter if $q_j \leq E_I(\pi_I)$ and stay out otherwise.

Since independent firms' expected profit from participating is constant, as long as they independently draw their entry cost, their entry choices are also independent. On the contrary, the entry of a group member is more likely when at least $N^* - 1$ other members enter too. Our participation test is based on this idea.

4 Econometric Tests

4.1 Participation Test

Participation patterns among groups of firms within a suspected set of cooperating firms have considerable potential to identify collaborators. Firms participating in our conjectured average-piloting strategy have to coordinate their entry in sufficient numbers. Coordinated entry of at least the minimum breaking coalition size from Proposition 2 allows a group to be certain of being able to influence the winning bid thresholds. Non-cooperating firms' entry decisions will be independent provided that their cost draws are independent.

The logic behind our participation test is to compare the participation patterns of a group of firms g comprised of firms suspected of cooperation with participation patterns in a control set of groups. Groups with randomly selected members from a comparable set of firms comprise a natural control set. If, for instance, group g has 5 members we can compare frequency of its members' participation in the same auction with the frequency of coincident participation for a randomly selected set of 5 firms. A key consideration in practice will of course be the choice of the set of comparable firms from which to choose the random comparison groups. Firm characteristics like location, size, etc. will certainly be important criteria for selecting comparable firms to those in g. For ease of exposition, we present our participation test without explicit conditioning on firm characteristics and simply denote the set of comparable firms, potential participants in each auction, as M. A discussion about M follows the introduction of our participation test.

Formally, our participation test is a test of the null hypothesis that a suspect group g with N^g members has the same distribution as a group comprised of N^g randomly selected firms from the set of potential participants M. Drawing N^g firms from M without replacement, we obtain $\binom{M}{N^g}$ combinations. Define H to be the set of all these combinations. Defining T as the total number of auctions and using the indicator that $d_{it} = 1$ for firm i attending auction t, we can define the frequency of auctions participated by all N^g members as:

$$f^g = \frac{1}{T} \sum_{t=1}^{T} \Pi_{i \in g} d_{it}$$

In the same way, we can define the analogous frequency for firms in the set $h \in H$:

$$f^h = \frac{1}{T} \sum_{t=1}^{T} \Pi_{i \in h} d_{it}$$

Our test decides whether firms in g have unusually coordinated entry by determining whether f^g is a tail event relative to the distribution of f^h induced by the random selection of group h, i.e. multinomial with equal probability on each element of H. This is commonly referred to as randomization or permutation inference (see Rosenbaum 2002). A one sided test of our null at the 5 percent significance level corresponds to the following decision: reject if f^g is greater than the 95th percentile of the f^h distribution. The f^h distribution can be exactly calculated or approximated via simulation.

The choice of the set of comparable firms M will be a key decision for implementation of our participation test. For the Italian roadwork procurement auctions that we study, participation is undoubtedly a function of firms' characteristics. Formal legal restrictions impose that a firm can bid in an auction only if it has a certification for both the job's type of work and for at least the contract reserve price.[20] Moreover, given the nature of road construction, transport costs will surely be important with proximity to the job site conferring cost advantages. Therefore, it is essential for the validity of our test that when we

[20]A firm is a potential bidder if at the same time: (i) the firm has the legal qualification to bid, (ii) has bid at least once in the county where the auction is held and (iii) has bid at least once in the region where the auction is held in the same year of the auction.

construct the control groups they match the suspect group along those firms' characteristics that determine entry. Otherwise, we might observe a difference in participation between the suspect group and a control group exclusively because their cost conditions induce a very different entry pattern.

The choice of the number of firms in g is also an important decision. Relatively large and small choices of g may be the most informative. When the group is large, for a fixed set M, power should be good as coincidental attendance of a large group of independent firms will be unlikely. Moreover, by Proposition 2 we know that a small two-firm group should also have good power as the minimum breaking coalition must have at least three firms. Although the concept of minimum breaking coalition has been defined relative to a scenario in which all firms outside this coalition behave independently, it is nevertheless suggestive that size-two groups formed from a set of cooperating firms might be less likely than size-two groups composed of independent firms to have both members coincidentally attend an auction.

It is important to note regardless of how well we use firms' characteristics to determine M, this set is very likely to contain both independent firms and undetected cooperating firms. Thus our null distribution under no cooperation is likely not an approximation of the conduct of independent firms. The data in any real application will be inherently a mixture of independent firms and undetected cooperating firms. In a typical non-validation style dataset of course we will not know which firms are independent and thus cannot construct a reference distribution for any null involving only independent firms for comparison to f^g. We anticipate a loss in power when forced to use such mixture datasets compared to the infeasible case with identifiable independent firms. To understand the magnitude of this power loss, below we investigate the performance of an analogous test exploiting identities of independent firms in our Validation data.

Validation Data Results: We report the results obtained for the Turin cartels in Figure 5. For each of the 8 cartels, the figure shows the frequency of participation of subgroups of all sizes. The red dotted lines are the 5^{th} and 95^{th} percentiles of the reference distribution formed from randomly chosen firms. For example, focus on panel (a), we observe the largest subset of cartel B that jointly enters has size 16. However, the 95^{th} percentile of the reference distribution for such a large group is approximately zero. Indeed, the 95^{th} percentile of the reference distribution is estimated to be positive only for subgroups no larger than 10. Across cartels, the frequency of joint entry for larger-sized suspect groups is much higher than that of the 95^{th} percentile of the reference distribution. Larger-sized groups provide clear rejections of the null of non-coordination in entry. A second relevant aspect for cartels B and C is that small subsets, of size 2, 3 and 2 respectively, have joint participation frequencies that are *lower* than the 5^{th} percentile of the reference distribution. Therefore, firms in B and C exhibit behavior consistent with a cartel considering minimum breaking coalition size when coordinating entry.

4.2 Bid Test

Our bid test is based on the type of coordinated bidding described by our theoretical analysis and motivated by anecdotal evidence provided by court documents associated with our Validation data. Despite the complex nature of the game, we showed in Proposition 3 a basic strategy for a cooperating set of firms: members clusters their bids in an attempt to pilot the relevant averages.[21] We exploit the details of our ABA mechanism to construct a test statistic that should be sensitive to the average-piloting behavior that such cooperating firms would use. For ease of exposition, we present our test without conditioning on any firm or auction characteristics. Extensions conditioning on observable characteristics are straightforward and discussed in following Sections.

We base our test on a measure of how much influence a given set of suspected firms has upon a trimmed mean discount ($A1$) for an auction. First, define a group g suspected of piloting averages. We consider an auction with N total firms with N^g firms in group g and N^{-g} firms not in this group. We define $B^g = \{b_1^g, ..., b_{N^g}^g\}$ as the ordered (from small to large) set of discounts from group g and $B^{-g} = \{b_1^{-g}, ..., b_{N-N^g}^{-g}\}$ as the ordered set of remaining discounts. The trimmed mean throwing out N' discounts[22] on either end is:

$$A1^g = \frac{1}{N^{-g} - 2N'} \sum_{i=N'+1}^{N^{-g}-N'-1} b_i^{-g}.$$

This statistic $A1^g$ will be systematically lower/higher than the trimmed mean of all the discounts if the group is trying to pilot the overall trimmed mean up/down. Formally, we test the null hypothesis that the firms in group g are not cooperating to pilot the overall trimmed mean. Our operational definition of 'not cooperating' is that firms are bidding independently.

A natural approximation of the distribution of $A1^g$ under the null hypothesis of no cooperation is that generated by randomly selecting a group of the same size as g, N^g, from the full set of discounts $B^g \cup B^{-g}$. Randomly drawing without replacement N^g discounts out of $B^g \cup B^{-g}$ results in N choose N^g combinations. Define S to be the set of all these combinations of ordered (from small to large) discounts so clearly $B^g \in S$. The trimmed mean without a combination $s \in S$ is:

$$A1^s = \frac{1}{N^{-g} - 2N'} \sum_{i=N'+1}^{N^{-g}-N'-1} b_i^{-s}.$$

and the distribution of $A1^s$ is multinomial with equal probability on each combination $s \in S$. When S is too large to compute this distribution exactly, it can be approximated via simulation.

[21] Notice that this prediction was derived under the assumption that firms know the likely range of winning bids. This assumption is justified by the detailed data on past auctions that bidders can easily collect thanks to the transparency requirements that regulation imposes on the PAs.

[22] N' is the 10 percent of the number of N rounded up to the next highest integer.

Our test decides whether a group of firms has unusually coordinated discounts by checking whether the realization of $A1^g$ is a tail event relative to the distribution of $A1^s$. Both Proposition 3 and 4 clearly indicate that a clustering in the upper tail of the discount distribution is characteristic of a cooperating group. Clustering on the lower tail could be either an intentional choice of a cooperating group or the result of high cost conditions for some independent firms. To capture these different possibilities, we consider both the one and the two-sided version of the test. [23]We approximate the $A1^s$ distribution via simulation, drawing a large number of times from S and calculating for each draw s the corresponding $A1^s$.

It is again important to note that our approximation of the null distribution under no cooperation is not likely to be an approximation of the conduct of independent firms. In any real application in which we are motivated to test for cooperative behavior, we anticipate that our bid data will be inherently a mixture of bids from independent firms and undetected cooperating firms. In a typical non-validation style dataset of course we will not know which firms are independent and cannot construct a reference distribution for the null of independent firms for comparison to $A1^g$. In the web appendix we use our Validation data to investigate performance of our bid test with a reference distribution of independent firms to better understand our procedure.

Multiple Auction Testing

Our bid test as stated above applies to a single auction. With data on multiple auctions it may be feasible to base tests on a group's behavior across auctions. This has the potential to increase our power to detect the type of cooperative bidding behavior our analysis is focused upon. However, testing across multiple auctions presents a formidable challenge if we allow for firm-level persistent idiosyncrasies in behavior. Our operational definition of non-cooperation needs to be augmented with respect to a firm's actions in multiple auctions. In particular, even if it is reasonable to use a benchmark that non-cooperating firms act independently within an auction, we should allow for a given firm's actions to be correlated across auctions in which it participates. The bottom line is that when taking a set of firms s, the set of $A1^s$ outcomes across multiple auctions should not be considered independent.

First consider a bid test across two auctions. We form a joint test statistic for a suspect group g of firms who participated in both auctions with bid test statistics $A1_1^g$, corresponding to the percentile p_1^g of the reference distribution, and $A1_2^g$, corresponding to the percentile p_2^g of the reference distribution. Our joint test statistic J^g describes the extent to which these percentiles are extreme, either small or large, across the two tests. For below-median percentiles we use the percentile itself and for percentiles above the median we use one hundred minus the percentile as a measure of how far it is in the tail. To aggregate across

[23]A two-sided version of this test at say the 5 percent significance level corresponds to the following decision: reject the null if $A1^g$ is not between the 2.5 and 97.5 percentiles of the distribution of the $A1^s$. One-sided tests likewise will reject if $A1^g$ is higher or lower than the corresponding critical values given by the appropriate tail percentile of the $A1^s$ distribution.

auctions we add the individual 'tail percentile' measures forming our statistic as:

$$J^g = \sum_{i=1}^{2} p_i^g 1\{p_i^g < 50\} + (100 - p_i^g) 1\{p_i^g \geq 50\}$$

where $1\{\cdot\}$ is the indicator function. This test statistic will take on small values if both test statistics are tail events and larger values otherwise. J^g clearly involves the same set of firms g in both auction one and two, whose p_1^g and p_2^g statistics could be arbitrarily dependent. In order to capture dependence across auctions in $A1_1^g$ and $A1_2^g$, we need to use the corresponding distribution for a randomly selected group s that participates in auctions one and two. Our reference distribution for J^g under the null hypothesis of no cooperation is the implied distribution of:

$$J^s = \sum_{i=1}^{2} p_i^s 1\{p_i^s < 50\} + (100 - p_i^s) 1\{p_i^s \geq 50\}$$

We approximate the distribution of J^s under the null via simulation, randomly selecting groups s and constructing J^s for a large number of draws s. This joint test is trivially extended in principle to any number of auctions by redefining J^s to depend on bid test outcomes from all the auctions.

In a scenario where all firms attend all auctions there is by construction no difference across firms in attendance patterns. However, in applications like ours an important issue with a joint test arises because not all firms attend all auctions and independent firms are less likely to jointly attend auctions together compared to firms acting cooperatively. In a setting where not all firms attend all auctions, we will still approximate the distribution of J^s under the null of no cooperation via simulation that randomly selects a group s. However, it will only be feasible to construct the statistic J^s when the group of firms s attends both auctions 1 and 2. Thus our reference distribution under the null implicitly conditions upon attendance at these two auctions. This is unavoidable unless we have an explicit model for how our bid tests are correlated across auctions.

Conditioning on auction participation has an important effect upon the composition of the reference or control distribution for our bid test. As noted in the previous Section, even in single auction case our approximation for the null distribution of no cooperation is likely to be a mixture of bids from independent firms and cooperating firms. When we condition upon attendance at two auctions, there will be a change in the composition of the approximate null distribution. The proportion of independent firms relative to undetected cooperating firms will decrease as attendance is required at an increasing number of auctions. For typical participation patterns of independent firms we expect that if we conditioned upon all the members of a group attending dozens of auctions, the large majority of firms left would be those in groups. Thus, as the number of auctions jointly considered increases, there is a cost in terms of power eventually decreasing due to this composition effect. Of course, this may be offset by the usual power benefits of multiple testing. We use our Validation data to investigate these costs and benefits and to calculate what is in a sense an optimal number of auctions to jointly test.

Validation Data Results for single-auction and multi-auction bid test: All of our tests are implicitly conditioning on several factors that disciplined our dataset construction. As discussed before, all auctions involve roadwork jobs, which are among the simplest and more standardized types of public works, and were procured by the same PA within a period of three years. As regards bidders, these must be firms that at the time of the auctions were in possession of the right legal qualification for this type of jobs. As shown in Table A.2 in the Web Appendix, after conditioning via sample construction, we could not find any firm attribute that (alone or jointly with others) was robustly associated with firms' bids in ABAs.[24] Therefore, in the rest of this paper we will use the bid test without explicitly conditioning on bidders observables. Nevertheless, as noted above our multiple auction bid test implicitly conditions on participation patterns.

Single-auction bid test results The results of the single-auction bid test for each of the 8 cartels across all the auctions are summarized by the histograms in Figure 6. Each histogram describes for all the auctions in the Validation data the percentile of the distribution of $A1^s$ to which $A1^g$ corresponds. A small percentile for $A1^g$ is consistent with a group trying to pilot the winning discount up and a large percentile is consistent with the group trying to pilot the winning discount down. Almost all histograms in Figure 6 have peaks close either to zero or to 100 indicating that the hypothesis of no cooperation is rejected in most of the cases.

It is important to note that we do not consider these single-auction tests to be independent of each other. Therefore, we do not have a precise prediction for the shape of the histograms in Figure 6 under the null of no cooperation. Nevertheless, these histograms are still very useful. Histograms that are very skewed will generate the same conclusion of evidence against the no cooperation null for many different sets of prior beliefs about the strength of dependence across auctions. Our strong prior belief is that dependence across auctions is weak enough here for large departures from uniform to be taken as strong evidence against the no cooperation null hypothesis.

Despite reasons to believe that piloting the winning discount up may be easier than piloting it down, our results suggest that piloting in either direction occurs. Cartels A, E, and H appear to have often piloted the winning discount down. Cartels B, C, G and F appear to have a tendency to push it up. Both directions appear frequently enough to warrant checking both rather than relying on a one-tailed test. The auctioneer of course is not indifferent to the direction of piloting as a upward move in winning discount increases revenue. Finally, our results for cartel D do not appear consistent with systematic piloting behavior. Cartel D is unique among the court-identified cartels as although its members were identified as cooperating, they were not charged for 'criminal association' because their cooperation was sporadic. Thus, in fact our finding of not detecting unusual cooperation in bidding for cartel D's members is not surprising.

[24]This table presents OLS regressions comparing the determinants of bids separately in two samples, one of FPAs and one of ABAs. Cost proxies appear predictive in the FPAs sample. However, in the ABAs sample they are not significant predictors of bid patterns.

Multi-auction bid test results We illustrate our multi auction bid test for cartel C in Table 5. The first choice needed to implement this test is the size of group to use. While using a group larger than the minimum breaking coalition is reasonable, too large a group will make the test useless as not enough independent firms will jointly participate in auctions. We use the smallest group size that is a local maximum of the participation distribution in Figure 5, provided it is at least four. For cartel C this is a group size of 5. We apply the joint participation test to the 5 members of cartel C with the highest joint participation.

Results are reported in Table 5. These 5 firms jointly entered in 51 auctions. Thus, in principle we could form the J statistic for each K-touple of auctions for K=2,..., 51. In the table we report only the results of the test for K=2, 4, 6, 8, 10 since with higher values of T it becomes harder to find control groups of firms that jointly entered in the same auctions. When, for instance, we look at the case of a pair of auctions, K=2, various combinations of two auctions exist out of the 51 auctions jointly entered. Instead of arbitrarily picking one, we pick at random many of these pairs and perform our participation test on each pair. In Table 5, for K=2 we perform the test using 531 different pairs that are randomly drawn from all pairs of these 51 auctions with at least 30 participants in common. For each of the 531 pairs, we construct the J^s control distribution using the set of all other pairs of firms that enter the same two auctions. For each of the associated J^g we compute its p-value in a two-sided test of no cooperation. Table 5 summarizes the distribution of these p-values across the 531 pairs by reporting its 10th 50th and 90th percentiles. Again, we note that these p-values are not from independent sets of auctions, their distribution needs to be considered along with prior information/assumptions about the strength of dependence. As for the single-auction tests, our strong prior beliefs are that this dependence is weak enough for substantial fractions of p-values less than 5% to be taken as evidence against the no-cooperation null.

Detailed results for each of the cartels, analogous to those in Table 5, are reported in Table A.4 in the Web Appendix. Additionally, these results are summarized in Table 6 which reports for each cartel: the size of the subgroup used, how many auctions the firms in this group won, whether the median p-value over all tests was .05 or less for at least one K-touple of auctions, the smallest K-touple of auctions where this median p-value was .05 or less. There is considerable variation across cartels in the number of auctions that need to be considered for the median p-value of these tests to be less than .05 in the six cartels where this occurs. The results for cartel D are unsurprisingly similar to the individual bid test results, there is little evidence against the null of no cooperation for this weak cooperation cartel.

Not detecting cooperation in cartel G is a bit surprising as the individual auction test provided evidence of cooperation in this cartel. The explanation lies in the structure and behavior of cartel G: this is a relatively large group of 16 firms, but only 5 of them win auctions. The non-winning partners always place supporting bids, generally consisting of very high discounts and the few designated winners always bid closer to the center of the distribution. This implies that when we look at the single-auction bid test using all the firms, we often detect these firms as a group (see panel c in Figure 6). What allows this cartel to evade detection in the multiple auctions testing is that the designated winners are the only groups that frequently participate together, individual supporting bidders participate sporadically. Therefore, our group selection method, selecting a subset of 4 firms within cartel G that jointly participate the most, results in a subgroup of 4 firms who are frequent winners and do not bid in an unusual manner. This highlights an important caveat that because it conditions on particular participation patterns of the chosen groups, the power of our joint test to detect non-cooperation can be less than our single-auction bid test.[25]

5 Testing Cooperation with Unknown Groups

Our testing methods can in principle be applied to any candidate group. In applications with a small number of firms, all possible groups could be examined. However, this is computationally infeasible in the situations like that in our Main dataset with hundreds of bidders. Feasible strategies for selecting groups of firms will of course depend on the available information. We investigate approaches that are feasible in two extreme scenarios. First, we present a method feasible with our data based on using firm characteristics. Our Validation data allow estimation of predictions of cooperative links between a pair of firms based on their characteristics. The fact that our Main dataset is comparable to the Validation data allow us to use this estimated model to predict links and groups in the Main dataset. We examine both the 'in-sample' performance of this method using the Validation dataset itself as the target as well as its performance using our Main dataset. Second, we examine strategies for choosing groups in a poor information case where we have no firm characteristics and hence resort to using participation patterns to define prospective groups and then, given candidate groups, employ our bid test. We make no claim that either group selection method is optimal, leaving the question of optimal group selection for future research.

Good Data Scenario. Our group selection method in this scenario has three steps.

Step 1: In both our Validation and Main datasets, we observe measures of firms' association along three dimensions: common ownership and management, formation of temporary bidding consortia and exchange of subcontracts. Using the Validation dataset, we construct all pairs of firms that can be formed by linking each one of the 95 known cooperating firms to any of the other bidders, through any of these three association measures. This results in 775 pairs. Since in this dataset we know the composition of the 8 cartels, we can estimate a model predicting which of these pairs are in the same cartel given their characteristics. We estimate a probit model where the dependent variable is one if the pair is in the same cartel and zero otherwise. Table 7 shows that the characteristics that we are analyzing help in predicting group membership. We also include measures of the geographical proximity between firms. Specification (1) in Table 7 indicates a positive association between the probability of being in the same cartel and exchanging subcontracts, sharing personnel, being located in the same county and having bid together in a legal consortium. In our favorite specifica-

[25]All these results were obtained with control groups that are known to be a mixture of independent firms and cartel members. The cost of this mixture versus a usually infeasible control distribution of only independent firms is explored in Table A.5 in the web appendix. Relative to those in Table 6, we observe the expected increase of power and we achieve evidence of cooperation for cartel G.

tion, model (2), we also use certain interactions of these links between firms to improve the predictive capacity of the model.

Step 2: We use our estimates of cartel membership probit model from the Validation dataset to generate predicted cooperative group membership probabilities for pairs of firms from a target dataset. We focus on the firms participating in auctions most frequently. The top 10% of firms in the target data (their number denoted N) in terms of participation are paired with each other firm in the sample with which they have at least one linkage due to common ownership and management, formation of temporary bidding consortia, or exchange of subcontracts. We do not use physical proximity to determine potential pairs as it is strongly related to contract execution costs as well as intercartel coordination costs. For each of these pairs, we construct a predicted probability of cooperative group membership. The complements of these predicted probabilities are interpreted as a dissimilarity array.

Step 3: We use the constructed dissimilarity array from step 2 with a standard hierarchical clustering algorithm (Gordon, 1999) to partition the firms into clusters. In the first round of the algorithm, all firms are singleton clusters. In the next rounds, firms (or groups of firms) are associated together on the basis of their average dissimilarity. The process stops when a maximum tolerance for dissimilarity is reached. The clustering algorithm has a tendency to yield some very large and small clusters that we trim away to arrive at a set of candidate groups. Since the procedure entails arbitrarily chosen tolerance parameters, we provide its exact details in the Web Appendix in the note to Table A.3. We experimented with different parameters and settled with those reported in the note to Table A.3.

The 'in-sample' performance of this group selection method with our Validation data is illustrated in Table 8. Note that our method should work well in this case as it was in a sense tailored to this dataset. The first column is an integer enumerating each of the 14 clusters created by our 3-step procedure. The second column is a letter identifier of the known cartels most often represented in each of the 14, labeled from A to H. The following column reports the size of this this subgroup. The following two columns report the number of members from different known cartels and the number of independent firms. The last two columns report, respectively, the total number of victories of the members of the cluster and which, if any, of our two tests leads to a rejection. One sided participation tests rejecting if the frequency of participation is above the 95th percentile of the reference distribution were conducted for the groups' largest joint participation size. Two-sided single-auction bid tests were conducted just as for the Validation data. As a simple summary of these tests, we label a group as being unusually cooperative if at least 30% of the p-values of these single-auction bid tests rejected no cooperation we label that group as being detected as unusually cooperative. The same classification is also achieved using the multi-auction bid test.

Overall this group selection method appears to perform reasonably well. The only cartel that has no member in any assigned cluster is cartel D. However, this is a cartel that received the smallest sanction from the court because it concerted actions only sporadically. Although several independent firms are assigned to groups, clusters 1, 4, 5, 7 and 9 have a substantial fraction of members of the same cartel . When clusters do not contain firms from cartels, our tests correctly do not indicate cooperation. The same lack of cooperation evidence occurs when there are two or fewer cartel members in a cluster. In five of the six clusters with three or more firms from a cartel, one or both of our tests rejects non-cooperation. The limits of the procedure are also illustrated here as the tests do not detect cooperation for cluster 5, despite 3 of its 4 members coming from cartel G. However, in this case the reason is specific to the bidding strategies of cartel G. As discussed in the previous Section, this is a large cartel with many fringe firms making piloting bids, but with a very small core of designated winners placing less extreme discounts. The 3 members of cartel G in cluster 5 belong to this subset of designated winners and this is why we fail to detect coordination for cluster 5.

Poor Data Scenario. Many auction datasets often contain information only on bidder identities and bids, thus we are motivated to explore a method for constructing candidate groups with such limited information. Here we examine the performance of a method that forms groups based on participation patterns and then applies our bid test for cooperation. The strategy is to identify firms that could be acting as group leader and then construct a candidate group by collecting firms that frequently participate with them. For our list of potential leaders we use the 10% of firms with highest participation. For each firm in this list, we construct a group with N members by looking at the frequency of joint participation. The first firm attached to the leader is the one that participates most often with this firm. Then we attach the firm that participates most often with the pair created in the previous step. We continue iterating upon this until we reach a group of N members.

We apply this methodology 'in-sample' to the Validation data to assess its performance. Imposing a group size of N=6, we obtain that only cartel B is detected. The reason is that the groups produced are a very poor approximation of the original cartels. These groups are described in Table 9, which reveals that only cartel B is well represented by several groups. Out of all the groups that we obtain,[26] the members of all the other cartels appear very sporadically and, in the case of cartels F and H, their members never appear.

This poor performance with only bid and identity data indicates that our methods will require some a priori knowledge of groups. While there are many applications where this is available, it inherently limits the use of our tests. Thus, for the final Section of this paper we assess the presence of groups in the Main data using the approach shown for the good data scenario.

6 Search for Cooperating Groups in Main Dataset

This Section illustrates our methods by applying them to study our Main data, exploiting the comparable auctions in our Validation dataset. We begin by applying our 'Good Data Scenario' group selection method and both cooperation tests to the ABAs in our Main dataset. We then use the results of our tests to identify a set of unusually cooperating firms. Using these firms as a benchmark we investigate the potential effect of these firms' cooperation on the revenues of the auctioneer and non-cooperating firms. We conclude this

[26] As explained above, we pick the top 10% of firms in terms of participation and for each of them construct a group. Thus, we pick 81 distinct firms. However, because some of the groups generated are identical we end up with the 56 distinct groups described in Table 9.

Section with a brief discussion using our benchmark cooperators to better understand the striking drop in participation when ABAs are replaced by by FPAs.

Group selection begins with a list of 400 potential group leaders comprising the top 10% participants in the Main data. We use the estimates from our Validation data to construct predicted probabilities of cooperative group membership for all potential pairings of each leader with other firms that are connected to them by at least one link based on common ownership/management, subcontracts, or consortia. We end up with a set of 1,848 different firms that our clustering procedure partitions into 289 clusters, most of which composed by a single pair of firms. Next, we prune these clusters by dropping firms that do not have at least a 20% predicted probability of being together with at least one of the other cluster members and then only consider clusters with at least 4 members. This results in 49 pruned clusters which comprise our groups for testing.

We applied our tests to these 49 groups producing the outcomes reported in the top panel of Table 10. The table provides details about those clusters for which one of our tests suggests cooperation/coordination. For each group we conducted the participation test comparing joint frequencies of various numbers of group members to those in a randomly selected control group. Essentially, we replicate the exercise detailed in Section 4 and illustrated in Figure 5 and the typical patterns are similar to those in this figure: common rejections of the no cooperation null for larger sized groups' joint attendance. We label a group as being unusually coordinated in entry if, for its largest participating group size its test statistic is above the 95th percentile of the reference distribution. This results in 42 groups being classified as unusually cooperating with an average size of 10 firms each. This is indicated in the first row of Table 10. In total, there are 408 firms in these 42 groups and their average number of bids, victories, and revenues is reported in the final columns of the table. For comparison these values can be related to those in the whole sample of firms reported in Table 2. Along all these dimensions, the average firm in the 42 groups appears orders of magnitude larger that the average firm in the whole sample.

The second row of Table 10 reports results for bid tests. For each of our 49 clusters we conduct two-sided single-auction bid tests for all auctions in which a subset of its members participate. In most of the cases, the results of these tests unambiguously distinguish groups between those that have very skewed histograms similar to those of the cartels in Figure 6 and those for which there is no evidence of possible cooperation. In particular, there are 6 groups with clearly skewed histograms. However, 2 more groups also appear to engage in cooperative behavior if we decide to label a group as unusually cooperative if the p-values for our single-auction test of non-cooperation are less than 5% in 30% of the group's auctions. Therefore, in total we find that 8 groups (135 firms) meet this standard for unusual cooperation.

The third row indicates that for 5 groups (80 firms) we observe a rejection using the multi-auctions bid test. We report the results of this test in Table A.6 in the web appendix. As with the Validation dataset, we face the problem that we cannot perform the test using all the members of the groups that we want to study because otherwise we would have an extremely small set of control firms. Indeed, even though a group of size 16 might participate in several auctions with all of its members, it is hard in the data to find enough comparable firms that participate together often enough to be a usable control group. Therefore, as done for the groups in the Validation data, for each of the 49 clusters we first fix the size of the subset that we want to pick by taking the size of the subgroup for which the frequency of joint participation is a local maximum and no less than 4. Then, among all possible subsets of this size in the group, we select the subset with the firms that participate together the most. Finally, we calculate the multi-auction bid test for each of these groups. Table A.6 in the web appendix reports the detailed results of the test for the 5 groups for which we observe evidence of cooperative behavior. These 5 groups are a subset of the 8 groups for which we observed a rejection of the single-auction bid test. The largest of these groups has size 9 and the smallest has size 4. For each of them we observe a large number of auctions in which the members of the group jointly enter: this number ranges from 52 to 95 auctions. We investigate sets of 2, 4, 6, and 8 auctions and we compute both the two-sided and the one-sided left version of the test and report the distribution of the p-values across sets of auctions. Using the same ad hoc criterion as with the Validation data to flag unusually cooperative groups if any of the median p-values is .05 or less. Four groups appear unusual with the two-sided test and one more via the one-sided test (rejecting only when the winning discount appears piloted up).

Given these definitions of coordinating/cooperating groups, we can quantify the number of auctions potentially impacted by firms engaging in coordinated behavior. However, it is not obvious what criterion to use when labeling an auction as suspected of being influenced by such behavior. Near one extreme, we could classify an auction as suspect if a minimal number of participants belong to a group identified by our participation test as unusually coordinated. Towards the other extreme, we could insist on only labeling auctions whose bidders include a group whom the single-auction bid test rejected no cooperation in that auction and who were part of a group that routinely failed bid tests in many other auctions. We could also adopt intermediate criteria involving both tests[27], e.g., checking whether an auction had bidders whose group was unusual according to our participation test and who the bid test indicated were unusual in that particular auction. Alternatively, we could require an indication of unusual behavior from the multi-auction bid test in addition to the participation test.

We prefer not to rely solely on a participation test to flag suspect auctions and so use bid test outcomes combined with the participation test. Of the 802 ABAs in our Main data, 61% of the auctions have among their bidders at least 3 bidders from a group that both is labeled unusual by our participation test and has single-auction bid tests routinely rejecting no cooperation, operationally defined as the 8 groups having p-values less than .05 in 30%+ of auctions. Instead, 43% of the auctions have at least 3 bidders from one of the 5 groups labeled unusual via our multi auction test described above. Finally, if we impose a stricter criterion that at least 5 members from a group must be present, then, respectively, 48% and 34% of the auctions qualify when using the 8 groups detected by the single-auction bid test and 5 groups detected by the multi-auction test.[28]

[27]Note that rejections under the participation test do not imply rejections under the bid test for a given group. The bid test does condition on a participation pattern, but such patterns need not be unusual from the participation test point of view.

[28]A formal test of whether an auction has suspect behavior from one of a set of groups is an alternative

6.1 Potential Effect on Cooperation on Revenues

The set of unusually cooperative groups detected by our tests captures a significant share of the revenues in this market.[29] Their members win 333 out of 802 ABAs. Nevertheless, contrary to typical cases of collusion in auctions, this is not necessarily an indication that the PAs could have paid a lower procurement price were these firms not engaged in bid coordination. In the unique equilibrium without cooperating groups all firms bid zero discounts and the auctioneer pays the reserve price: the highest procurement cost. Regulations mandate that this reserve price cannot be set based on the PAs' expectations about bidder behavior.[30] This makes the observed reserve prices reasonable values for their counterparts in a counterfactual thought experiment without cooperating firms. This gives us a clear benchmark for this counterfactual scenario: all PAs would have paid an amount equal to the observed reserve price. Therefore, in the Main data, at average reserve price of €312,000, the average winning bid of 13.4% implies that the auctioneer's savings due to firm cooperation is about €42,000 per auction.

The cooperative activity of groups surely results in both winners and losers. The cooperating group members piloting the winning discounts upwards are intending to increase their chance of winning at the cost of getting a lower payoff if they do win. Clearly this can be beneficial to them if the increase in the win probability is large enough compared to the cost of lower payoffs for a win. In contrast, the non-cooperating firms are surely worse off. Their winning probabilities are reduced due to being crowded out by cooperators and when cooperators force down the winning discount this obviously reduces the payout when the non-cooperators win.

Consider an example scenario in which we can assess the relative importance of win probability reduction versus win payoff reduction in expected revenues for non-cooperators resulting from cooperation. A typical auction in our main data has about 51 bidders, 17 of whom are members of our detected cooperating groups.[31] Consider a hypothetical auction with 34 independent firms and 17 colluders. In the no cooperation equilibrium, each of the 51 bidders would have a 1.96% chance of winning the auction. Suppose that with cooperation the 17 colluders can increase the probability that one of them wins to our sample group win frequency of 333/802 and independent firms all have the same, lower probability of winning. In this scenario, the win probability of the 34 independents when there is cooperation among the colluders drops to (1 - 333/802)/34 = 1.70%. Thus in

approach here and straightforward to implement. Testing a null that more than one group of specified sizes have the same distribution as a comparably sized random set of groups can be done via randomization inference in the same fashion as our tests. Test statistics determined by the set of groups outcomes can be compared to a reference distribution determined by randomly choosing sets of groups.

[29]For the 802 auctions in the Main data, the cumulative reserve price is €370 million. The 42 groups in the first row win 333 auctions, with a cumulative reserve price of €143 million. The analogous values for the 8 and 5 groups in the following rows are: 135 and 86 auctions and €54 and 37 million.

[30]The reserve price is obtained by applying an official menu of prices, common across PAs in the same region, to the estimated input quantities required by the work. Although the PAs could try to manipulate these estimates, for the simple roadwork contracts that we study, this manipulability should be rather limited.

[31]The average entry in the 802 ABA is 50.7. Considering as cooperating firms those belonging to the 42 groups detected by the participation test, these firms are on average 33.3% of the entrants.

this example, there is a 13.2% decrease in the win probability for independent firms due to cooperation among their competitors with a corresponding 13.2% decline in expected revenues. As above, we take our sample's 13.4% winning discount as representing the effect of cooperation upon winning discounts. Insofar as this example is a reasonable benchmark for firms in our Main data, the effect of cooperation upon independents' win probabilities appears to be as important as its effect discounting a winning payoff in impacting expected revenues.

However, one good reason such a simple calculation may not be a good counterfactual scenario is it fixed the auction participants and so does not account for the greater entry that would likely occur were all auctions awarded at the observed reserve price. A structural analysis in the spirit of Asker (2010) would be needed to properly pin down the counterfactual revenues for independent firms. Although this is beyond the scope of the current paper, it seems plausible that sample of FPA auctions could be useful to permit the estimation of firms' entry cost necessary for this type of analysis. Similarly, our above calculation of the PAs' savings due to bid coordination does not consider that, were all bids to converge to zero, PAs would probably change their auction format. It is known that the legislators introducing the ABA were not expecting all the bids to go to zero, but, were this to happen, it is not known which changes to the rules they would have made. In any case, our analysis has shown that a strong incentive to bid coordination operates against such convergence.[32]

6.2 Drop in Participation with FPA Introduction

In this Section we use our cooperation tests to better understand the drastic drop in participation that occurred in our Main data when the ABA was replaced by the FPA beginning in 2006. We use the full version of the Main data including both the 802 ABAs and the 232 FPA. In Section 2, we discussed the striking difference in the number of bids in ABAs and FPA. Although all the auctions described in Table 2 are rather similar, the ABAs receive on average 51 bids and, frequently, more than 100. Instead, in FPAs the average number of bids is 7 and the auction with most bidders has 48 bids.[33] As explained in Section 2, starting from 1999 ABA was the major procurement system until a series of reforms introduced FPA. The drastic change in the number of bids associated with these switches to FPA is well illustrated by Figure 7. In this figure, the blue circles mark the ABAs and the red cross mark the FPA. The top panel reports the number of bidders in ABAs and FPAs held by four PAs in the Main data that switched to the FPA in 2006. The systematically lower values of the crosses (the FPA) relative to the circles (the ABA) is evident.

There are at least two causes for the drop in participation with introduction of FPA:

[32] An auctioneer may exploit these incentives to break an all inclusive coalition. If the auctioneer is using a mechanism weak versus collusion (like a second price auction) and is limited in the choice of an alternative mechanism by a high default cost (so that a first price auction would not work), then the use of an ABA might induce the formation of groups that break the grand coalition (without exacerbating the default risk). For the Italian case, the introduction of the ABA was unrelated to the concern about all inclusive coalitions.

[33] Notice that the values for the FPAs are in line with those observed in other countries. In the US, road construction contracts awarded by the largest DoT via the FPA receive on average between 3 and 7 bids.

the disappearance of shills who provide their creator only with extra bids that are useless in FPA and the exit of inefficient firms that have too little chance of winning FPA.[34] In the Main data, about 4,000 firms bid at least once in ABAs, but about 3,000 of these never bid once in FPAs. Focusing on firms that were qualified and nearby prospective FPAs we examine a subset of 1482 firms who attended at least 3 ABAs in counties where subsequently at least three FPAs for which these firms legally qualified to bid were held. The 1482 firms contain 298 members of our 42 cooperating groups and 1184 non-grouped firms.[35] Of the 298 cooperators about half (159) do not participate in an FPA and likewise about half of the 1184 non-cooperators also do not participate in an FPA. Referring to those not participating in an FPA as exiters, the frequency of exiters does not depend on cooperating status.

Characteristics for these firms are reported in Table 11. We anticipate that shill firms will be predominately located in our detected cooperating groups rather than among our noncooperating firms. (If we had perfect measures of cooperation, shills would only be present in cooperating groups.) Thus the composition of exiters in terms of shills versus inefficient firms should vary according to whether the firms are cooperators and this should show up in firm characteristics.

There are substantial differences in the characteristics of exiters according to whether they were labeled cooperators or non-cooperators. Among the non-cooperators, exiters have smaller capital and labor force relative to those who participate in FPA despite being slightly older firms, possibly signaling their relative inefficiency. Cooperating group exiters also have less capital and workers than FPA participants but these gaps are much smaller than for non-cooperators.

An important caveat to the interpretation of the ownership and management characteristics reported in Table 11 is that there are serious missing response issues. We do not have the data to address this issue and necessarily proceed to interpret these statistics as though non-response was random.[36] With this caveat in mind, there do appear to be female ownership and management differences according to cooperation status. For noncooperators exiters have lower or nearly the same frequency of female ownership and management presence. In contrast for cooperating firms there is modest evidence of exiting firms having more women owners and more female managers versus those that stayed and participated in FPAs. This is in line with the legal case in Turin where shill firms were often formally owned and managed by the mothers, sisters or wives of the men convicted for collusion. The presence of shills is also suggested by some ad hoc comparisons of the firms in the 5 groups detected by our multi-auction bid test. For instance, we have a few instances of pairs of firms registered at the exact same street address that bid together in almost all the ABAs in which they participate, but that have only a single member of the pair bidding in FPA. In part because of the large scale market shakeout that we have discussed in this Section, the

[34] Although from an economics standpoint shills are 'fake' firms, from a legal standpoint they must be perfectly legitimate firms, otherwise they would not be allowed to bid in public auctions.

[35] The 42 groups in the top row of Table 10 are used to classify group firms. Qualitatively, the results do not change if one of the two more stringent classifications is used.

[36] With more complete data, the exogenous shock given by the switch to FPA could have been exploited to more rigorously trace out the connections between firms, in the spirit of Bertand et al. (2002).

Italian Parliament is currently discussing a new program to offer subsidies to small business bidding for public contracts. However, our results suggest there would be benefits to a more in-depth analysis of the shakeout and of which firms might truly deserve subsidies.

As a final note, we want to remark that a similar drop in participation occurred also with the switch from ABA to FPA of Turin in 2003. Indeed, when the collusion case that we discussed became public, both the county and the municipality of Turin abandoned ABA in favor of FPA. Although, our Validation data contains only ABAs and our Main data starts in 2005, we used the data from the Italian Authority for Public Contracts to report in the bottom panel of Figure 7 the evolution of the number of bids under the two formats.[37] As in our Main data, the drop in participation following the switch to FPA is evident.

7 Conclusions

In this paper, we document that the ABA gives strong incentives to bidders to coordinate their entry and bidding choices. We propose two statistical tests to investigate bidder cooperation and show that they work well in a Validation dataset where 8 cartels have been identified by a court. These are tests for whether groups of firms participate or bid differently than other comparable groups of firms. Our metrics for describing participation and bidding patterns are motivated by a conjecture that cooperating firms employ a strategy of jointly participating in large enough numbers to pilot winning bids in these auctions. This is a reasonable equilibrium strategy and supported by anecdotal empirical evidence. Finally, we apply these tests to a different dataset of ABAs in which the presence of groups has not been previously known and show that the tests suggest the presence of several groups influencing numerous auctions. Thus, although no statistical test is a final proof, a natural application of our tests could be of help to courts evaluating cases of coordinated bidding. In this respect, a good feature of our tests is that they are somewhat 'inspector proof' in that even if firms knew of them, avoiding detection would require foregoing, at least in part, the benefits of cooperation.

We are optimistic that our tests could be adapted to detect cooperation in other environments where similar incentives to manipulation thresholds exist. For example, in public works auctions, like those referenced in Table 1, or in the Medicare system both for the medical equipment DEMPOS auctions and for the determination of the low income subsidy in Part D.

Importantly, our results also indicate that it is not obvious that bidder cooperation should always be sanctioned. Indeed, we present the case of an important market in which bidder cooperation reduces the procurement cost for the auctioneer. Therefore, our results argue against any automatism in antitrust activity. Instead, we see a role for the use of an accurate economic analysis of bidder behavior as a guide to the quantification of the effects of bidder agreements.

[37]This dataset contains information about the winning bid and the number of bids but not the identity of all bidders and, hence, cannot be analyzed through our tests.

Finally, our most general contribution is to clearly show how sophisticated and quantitatively important bidders' responses can be to auction rules. Therefore, our study joins the recent literature on market design in arguing in favor of a careful design of incentive schemes in auctions and procurement mechanisms.

8 Bibliography

Abrantes-Metz, R., P. Bajari, (2010). "Conspiracies Detection and Multiple Uses of Empirical Screens," mimeo.

Albano, G., M. Bianchi and G. Spagnolo, (2006). "Bid Average Methods in Procurement," *Rivista di Politica Economica*, 2006, (1-2): 41-62, reprinted in *Economics of Public Procurement,* Palgrave-MacMillan.

AGCM, (1992). "Appalti Pubblici e Concorrenza," Italian Antitrust Authority report.

Asker, J., (2010). "A Study of the Internal Organization of a Bidding Cartel," *American Economic Review*, 100, 3, 724-762.

Bajari, P., L. Ye, (2003). "Deciding Between Competition and Collusion," *The Review of Economics and Statistics*, November, 971-989.

Bertrand, M., P. Mehta, S. Mullainathan, (2002). "Ferreting Out Tunneling: An Application To Indian Business Groups," *Quarterly Journal of Economics*, 117, 1, 121-148.

Conley, T. G., (1999). "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 92, 1, 1-45.

Coviello, D., M. Mariniello, (2011). "The Role of Publicity Requirements on Entry and Auction Outcomes," mimeo.

Cramton, P., S. Ellermeyer, B. E. Katzman, (2011). "Designed to Fail: The Medicare Auctions for Durable Medical Equipment," mimeo.

Decarolis, F., (2011). "When the Highest Bidder Loses the Auction: Theory and Evidence from Public Procurement", mimeo.

Decarolis, F., (2012). "Pricing and Incentives in Publicly Subsidized Health Care Markets: the Case of Medicare Part D," mimeo.

Decarolis, F., C. Giorgiantonio, V. Giovanniello, (2010). "The Awarding of Public Works in Italy," Bank of Italy, QEF n. 83.

Gordon, A. D., (1999). *Classification*, 2' edition, Chapman & Hall/CRC.

Graham, D. A., R. C. Marshall, (1987). "Collusive Bidder Behavior at Single-Object Second Price and English Auctions," *Journal of Political Economy*, 95, 1217-1239.

Haberbush, K. L., (2000). "Limiting the Government's Exposure to Bid Rigging Schemes: A Critical Look at the Sealed Bidding Regime," *Public Contract Law Journal*, 30, 97-122.

Harrington, J., (2011). "A Theory of Tacit Collusion," mimeo.

Hendricks, K., R. Porter, (1989). "Collusion in Auctions," *Annales d'Economie et de Statistique*, 15/16, 218–230.

Ioannou, P., S.S. Leu, (1993). "Average-Bid Method. Competitive Bidding Strategy," *Journal of Construction Engineering and Management*, 119, 1, 131-147.

Ishii, R., (2006). "Collusion in Repeated Procurement Auctions: A Study of Paving Market in Japan", mimeo.

Katzman and McGeary (2008). "Will Competitive Bidding Decrease Medicare Prices?," *Southern Economic Journal*, 74 (3), 839-856.

Klemperer, P., (2004). *Auctions: Theory and Practice*. Princeton University Press, NJ.

Liu, S., K.K. Lai, (2000). "Less Averge-Bid Method. A Simulating Approach," *International Journal of Operations and Quantitative Management*, 6, 4, 251-262

Mailath, G., P. Zemsky, (1991). "Collusion in Second Price Auctions with Heterogeneous Bidders," *Games and Economic Behavior*, 3, 467–486.

Marshall, R. C., L. M. Marx, (2006). "Bidder Collusion", *Journal of Economic Theory*, 133, 374-402.

McAfee, R. P., J. McMillan, (1992), "Bidding Rings," *American Economic Review*, 82, 579–599.

Pesendorfer, M., (2000). "A Study of Collusion in First-Price Auctions," *Review of Economic Studies*, 67.

Porter, R., D. Zona, (1993). "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, 101, 3, 518-538.

Porter, R., D. Zona, (1999). "Ohio School Milk Markets: An Analysis of Bidding," *RAND Journal of Economics*, 30.

Radner, R., (1980). "Collusive Behavior in Oligopolies with Long but Finite Lives," Journal of Economic Theory 22, 136-156.

Robinson, M. S., (1985). "Collusion and the Choice of Auction," *RAND Journal of Economics*, 16, 141–145.

Rosenbaum, P., (2002). *Observational Studies*, Springer, 2nd edition.

Spulber, D.F. (1990). "Auctions and Contract Enforcement," *Journal of Law, Economics and Organization*, 6, 2, 325-344.

9 Appendix: Proofs

Proof of Proposition 2: If the coalition is all inclusive, offering a zero discount is the best that can be done. Therefore, if $N^g = N$ all equilibria have the winning discount equal to zero. Instead, for coalitions that are not all inclusive, the relevant "minimum breaking coalition" in defined as $N^* = 2 + N'$.[38] Any group that can submit at least N^* bids has profitable deviations when all other discounts are equal to zero. One such deviation is to place $N^g - 1$ identical bids, all equal to ε, for small $\varepsilon > 0$, and the remaining bids equal to $\varepsilon/2$. This strategy gives to the group (approximately) the highest payoff in case of victory and a probability of winning of one (prior to the deviation the probability of winning was N^g/N). However, if the group does not reach a size of at least N^* it cannot profitably deviate from the zero-discount equilibrium: all its bids away from zero would have a zero probability of winning due to the trimming and the fact that the winner has to be below A2.

Proof of Proposition 3: Proving the second part of the proposition is simple. The claim is that given any pair (b_f^I, N^I) for any any small $\varepsilon > 0$ we can find a value N^{g**} such that if the size of the group, N^g, is $N^g \geq N^{g**}$, then there is an ε-equilibrium in which all group's bids are clustered below $b_f - \eta$. To see why this is the case, suppose that $N^g = (9)N^I$. Then consider a strategy profile for the group bids such that: (i) exactly $(.1)N$ bids are equal to zero and (ii) all the remaining bids are extremely close together and randomized in $(0, \varepsilon)$. Together this strategy profile and b_f^I constitute a ε-equilibrium. In fact, the group is winning with probability one and it is doing so at a discount that is less than epsilon above zero. The independents make a zero profit. However, their expected gain from a deviation can be made arbitrarily small because their probability of winning can be made arbitrarily small. This happens becuase the location of the interval where the winning bid lies, $[A1, A2)$, is governed exclusively by the group's bids. Since they are closely clustered together and randomized, the probability that an independent wins is negligible. Hence, we have shown that for $N^g = 9N^I$ an ε-equilibrium exists. For any larger N^g the same strategy profile is also a ε-equilibrium. This argument implies that N^{g**} always exists: it is at most equal to $9N^I$ and it is possibly smaller, but this depends on the exact profile of bids in b_f^I.

To prove the first part of the proposition, we want to show that there is an N^{g*}, with $N^{g**} \geq N^{g*}$, such that if the group size is at least N^{g*} but less than N^{g**}, then there is an ε-equilibrium. By allowing for the possibility that $N^{g*} = N^{g**}$, we are saying that there might be no group size that allows a ε-equilibrium where its bids are clustered above b_f and in this case only downward clustering as defined above occurs. In this case we define N^{g*} to be equal to N^{g**}. Instead, to see what are the conditions determining $N^{g*} < N^{g**}$ and so the clustering of bids above b_f, we use a constructive proof for a specific b_f^I. The same logic can then be applied for other b_f^I. Therefore, consider the b_f^I such that: (i) a subset ofN' independents bids $b_f - \eta$ and (ii) the remaining independents bid b_f. Given this bids profile, we start by looking at a group of size N^* (see the previous proposition) with bids clustered above b_f according to the profile: (i) two bids, b_1^g,and b_2^g, such that $b_2^g = b_1^g + \delta$ with a very small $\delta > 0$ and with $b_i^g \in [b_f, b_f + \varepsilon]$, $i = 1, 2$ with small $\delta < \varepsilon$ and (ii) the remaining $N^* - 2$ bids all identically equal to some value $b_h \in (b_f + \varepsilon, 100]$. To find the conditions under which

[38]Throughout this section we defineN' as $\lceil (.10)|N|) \rceil$.

this bids profile together with b_f^I constitutes a ε-equilibrium we proceed in steps.

Step 1: First we show when the group's bids constitute an ε-best response to b_f^I. Regardless of the exact values of b_h and b_f, any $b_1^g > b_f$ implies that the group wins with probability one at a price of b_1^g. Could this group do better by bidding b_f or less? If the group could place all its bids in $(b_f - \eta, b_f)$ it would again win probability one and with a better price. However, this gain is bounded by η so that at most the group could gain $(b_f + \varepsilon - \delta) - (b_f - \eta) = \varepsilon - \delta + \eta$. Since the only restriction on ε is $\varepsilon > 0$, by selecting appropriately small δ and η we can make ε small. As regards placing bids below $b_f - \eta$, placing less than N^* of them below b_f leads to a zero probability of winning. However, even clustering all bids below (downward clustering) $b_f - \eta$ might never lead to a positive profit if b_f is low and N is large relative to the group size. The reason is that a downward clustering strategy is profitable iff $A2 \leqslant b_f - \eta$, otherwise one of the independents win. However, dragging down $A2$ cannot be achieved by placing all bids equal to zero: in this case the group minimizes $A1$ but loses all its influence on $A2$ which would then be commanded only by the independents bids resulting in the victory of one of them at the price $b_f - \eta$. To maintain any influence on $A2$ the group must keep at least one bid strictly greater than $A1$. We next show that sometimes this is impossible. Let's indicate by $b_{N^*}^g$ the highest bid that the group submits. Since the first $N^* - 2$ bids are trimmed in the first stage of the AB alorithm and since among the 2 remaining bids the lowest will always be strictly less than $A1$, the best the group can do is to place: (i) $N^* - 1$ bids equal to zero and (ii) $b_{N^*}^g \in (0, b_f - \eta)$. However, if such bids profile has to achieve $A2 \leqslant b_f - \eta$, then it must be that $b_{N^*}^g \leqslant b_f - [N(.7) - 1]\eta$. But since $b_{N^*}^g > A1$ requires $b_{N^*}^g > [(N(.8) - 2)b_f - N(.1)\eta]/(N(.8) - 1)$, then there is no $b_{N^*}^g$ that can satisfy both conditions at the same time whenever: $b_f \leqslant [N(N(.56) - 1.6) + 1]\eta$. Similar conditions to the ones found here for a group of size N^* can be derived for larger groups to check whether there is a downward clustering strategy achieving at the same time $b_{N^*}^g > A1$ and $A2 \leqslant b_f - \eta$. If that is not the case, then only through upward clustering the group ε-best responds to b_f^I. However, from part two of proposition 3 we also know that at a certain point the coalition size will be so large to allow only for ε-equilibria with downward clustering.

Step 2: To close the proof, we need to show that with the proposed profile of bids for the group, no independent bidder can deviate and gain more than ε. An individual independent deviating to a bid below b_f loses with probability one and the same is true for any deviation above $b_f + \varepsilon$. However, a deviation to a bid $b' \in (b_f, b_f + \varepsilon)$ might be profitable if there is a high enough probability that b' wins. With a group of size N^* this is the probability that $b' \in (A1(b_2^g, b_f^I), b_2^g - 2\delta] \bigcup [b_2^g - \delta, b_2^g - \delta/2)$. Depending on the deviant firm's valuation and on the group's randomization, the expected gain for the deviant might be contained within ε. Nevertheless, it is always possible to arbitrarily shrink the gain of the deviant by increasing the group size: if we increase the group size above N^* and place each additional bid equal to b_2^g, then, as the group size grows, the probability that the deviant wins goes toward zero. If the group size needed to achieve this is smaller than N^{g**}, we have obtained an ε-equilibrium with upward clustering and we define N^{g*} as the smallest group for which this is the case.

Proof of Proposition 4: It follows from the rules of the ABA that a bid of b_h cannot win unless all bids are equal to b_h. However, this latter scenario cannot be an equilibrium since $b_h > 0$ implies that there is always a unilateral profitable deviation by bidding zero.

Tables and Figures

Table 1: Rules for Identification and Elimination of Abnormal Bids

Automatic Elimination	Only Identification	Rule Not Disclosed
Chile	Belgium	USA - California DoT
China	Brazil	
Colombia	Germany	
Italy	Portugal	
Japan	Romania	
Peru	Spain	
Switzerland	Turkey	
Taiwan	UK	
USA - Florida DoT		
USA - NYS Proc. Ag.		

Source: Decarolis (2011). Classification based on the rules for public procurement for works, goods and services. Left: countries having auctions with automatic (i.e. algorithmic) elimination of some bids. Center: countries having auctions with algorithms to identify abnormal bids but without the automatic elimination of these bids. Right: the California DoT is believed to use an algorithm to identify abnormal bids but it does not publicly disclose it. For additional details on the regulation of each country see Decarolis (2011).

Table 2: Summary Statistics - Main Data

Auctions for roadwork contracts below €1 million, Nov 2005 - May 2010													
Statistics by Auction							Statistics by Firm						
	Mean	SD	Med	Min	Max	Obs		Mean	SD	Med	Min	Max	Obs
ABAs							Entry	13.1	22.1	4	1	205	4005
HighBid	17.4	5.4	17.4	1.6	37.4	802	Wins	.31	.87	0	0	18	4005
WinBid	13.4	5.2	13.5	.51	36.8	802	Pr.Win	.03	.12	0	0	1	4005
W.-2Bid	.24	.68	.07	0	9.4	802	Reven	170	1081	0	0	$4e^{04}$	4005
With.SD	2.9	1.4	2.7	.14	9.2	802	Age	22.3	13.8	21	1	106	3611
No.Bids	50.7	34.3	43	5	253	802	Capital	447	2411	52	10	$8e^{04}$	2484
Res.Price	312	204	250	11	999	802	Subct	.65	2.9	0	0	53	4005
							Miles	159	234	47.8	0	1102	4005
FPAs													
WinBid	28.9	9.9	29	1.2	53.4	232	Frims that ceased activity						3.4%
W.-2Bid	4.5	5.0	3.0	.01	41	232	Location of firms headquarter:						
With.SD	6.9	3.1	6.6	.07	19.1	232	North5						69.6%
No. Bids	7.3	5.5	6	2	48	232	Center and other North						18.4%
Res.Price	342	288	215	30	978	232	South and Islands						12.0%

Table 2 in the left panel reports summary statistics for ABAs and FPAs for roadwork contracts procured by municipalities of five Northern regions: Piedmont, Liguria, Lombardia, Veneto, Emilia-Romagna. Top left panel: statistics by auction for the sample of ABAs. The variable HighBid is the highest discount, while WinBid is the winning discount. W.-2Bid is the difference between the winning bid and the bid immediately below it (sometimes referred to as 'money left on the table'). W.-2Bid is frequently equal to zero. In these cases ties are broken with a fair lottery. Across all bids within the same ABA, ties are frequent: in 209 ABAs at least two bids are identical, for a total of 720 couples and 38 triplets. With.SD is the within-auction standard deviation of bids. No.Bids is the number of bids. Res.Price is the auction reserve price. The bottom left panel reports for comparison the same statistics for FPAs held by the same PAs. The HighBid is (almost) always WinBid and so is not reported. Right panel: statistics by firm. The variables reported are the number of auctions attended (Entry), the number of victories (No.Win), the probability of winning in the sample (Pr.Win), the total revenues earned (Reven), the age (Age, measured in years in 2010) and the capital (Capital, measured in 2005), the number of subcontracts received (Subct), the miles between the firm and the work (Miles), whether the shuts down between 2005 and 2010 (Closed) and whether it is located in the same five regions in the North where also the auctions were held (North5), in other northern or central regions or in the southern regions or the islands. Revenues and capital are in thousands of Euro.

Table 3: Known Cartels of the Validation Data

Cartel Name and ID	No. Firms	No. Victories	No. Auctions
1 - Torinisti (B)	17	83	247
2 - San Mauro (C)	13	35	234
3 - Coop (G)	16	73	240
4 - Pinerolesi (A)	11	1	110
5 - Canavesani (E)	11	7	155
6 - Settimo (D)	6	10	220
7 - Provvisiero (F)	7	11	73
8 - Tartara-Ritonnaro (H)	14	1	62

Table 3 reports data on the 8 cartels of the Validation data. The first column reports the name of the cartel and, in parenthesis, the capital letter that we use to identify the group. These capital letters also pinpoint the cartels in the map in Figure 3. The last three columns of the table report the size (i.e., the number of firms) of the cartel, the total number of auctions its members won and the total number of auctions attended by at least one member of the cartel (out of the 276 auctions of the Validation data).

Table 4: Summary Statistics - Validation Data

	Mean	SD	Med	Min	Max	Obs		Mean	SD	Med	Min	Max	Obs
Statistics by Auction													
HighBid	22.8	5.6	22.1	12.5	47.5	276	With.SD	3.6	3.9	1.7	.34	10	276
WinBid	17.4	5.0	17.3	6.7	37.7	276	No.Bids	73.3	37.1	70	6.0	199	276
W.-2Bid	.09	.23	.05	0.0	2.9	276	Consort	3.0	4.8	1.0	0.0	24	276
Independent Firms							Firms in the 8 Cartels						
Entry	17.2	22.3	9.0	1.0	186	717	Entry	82.9	71.1	54	1.0	263	95
Wins	.13	.42	0.0	0.0	3	717	Wins	1.9	3.1	1.0	0.0	19	95
Reven	51.8	19.6	0.0	0.0	2319	717	Reven	822	1466	327	0.0	$1e^{04}$	95
Miles	237	284	101	0.0	1071	504	Miles	101	207	15	0.0	991	86
Age	27.1	14	25	2.0	106	559	Age	29.6	14.1	30	1.0	72	91
Subct	1.8	5.0	0.0	0.0	53	717	Subct	6.8	8.6	4.0	0.0	44	95

The variables used to describe the auctions are the same of those in Table 2. The only additional variable is Consort which measures the number of (legal) bidding consortia present in the auction. Each consortium places one single bid. The type of jobs and the reserve price of contracts is similar to those in Table 2. The set of 717 independent firms contains 24 firms that share part of their owners and managers with the cartel firms. Their presence makes the summary statistics of the independent firms slightly closer to those of the cartels. The missing values for miles and age are due to the impossibility of identifying with certainty some firms.

Table 5: Multi-Auction Bid Test - Cartel C Results

	Cartel C			
	Subgroup: Size = 5; Auctions Entered = 51			
	Two-sided Test			
K	10^{th}	50^{th}	90^{th}	No. Repetitions
2	0.03	0.11	0.38	531
4	0.01	0.06	0.20	399
6	0.01	0.03	0.14	311
8	0.00	0.01	0.04	278
10	0.00	0.01	0.03	209

Table 5 shows the distribution of the result of the multi-auction bid test for cartel C. The results are based on a subgroup of cartel C composed by the 5 members of this cartel that participate together the most often. These 5 members jointly entered 51 auctions. In the first column, K=2, ..., 10 indicates how many touples of auctions were used to perform the test. Thus, K=2 means that the test is for a pair auctions, K=3 that it is for a triplet auctions and so on. In principle we could look to K as large as K=51 but if we were to do so the number of control firms that jointly entered all these auctions would be extremely limited. Columns 2 to 4 of the table report the distribution of the result of the test. We have a distribution of results because when we take any K<51 there are various combinations of auctions that could be used to perform the test. For instance, if K=2 every 2 auctions among the 51 jointly entered by the subgroup could be used. Instead of arbitrarily picking a single pair, when K=2 we could perform the test on every pair of 2-auction out of the 51 auctions or, when the number of pairs is large, randomly draw pairs of auctions. We do the latter but also require that the touple of auctions has at least 30 firms in common so that we can have enough control groups. So, for instance, the results reported for K=2 are based on 531 pairs of auctions that are drawn at random out of all the pairs that could be constructed from the 51 auctions and that also have at least 30 firms in common. The results of the three central columns should be read as p-values: the null stating that the group is not different from the control groups is rejected if the reported p-value is less than a critical level. Unless stated otherwise, all comments in the main text are based on the 50^{th} percentile of the distribution of the results and a critical level of 5%. See the Web Appendix for more details about how the test statistic is constructed and for the complete results for all the 8 cartels.

Table 6: Multi-Auction Bid Test: All Cartels

Known Group	Subgroup Size	Auctions Won	Rejection of Null	N-touple for Rejection
B	5	53	Yes	7
C	6	29	Yes	6
G	4	17	No	-
A	7	1	Yes	2
E	5	5	Yes	10
D	4	8	No	-
F	5	6	Yes	2
H	5	0	Yes	2

Table 6 summarizes the results of the multi-auction bid test for the 8 groups. The full set of results of this test are reported in Table A.4 in the Web Appendix. The eight rows of Table 6 report the outcome of the multi-auction bid test for each of the groups in the Validation data. The first column reports the group identifier, the second the size of the subgroup used for the test, the third the number of auctions won by this subgroup, the fourth report a Yes if the null of the two-sided multi-auction bid test is rejected at the 5% level using at most a 10-touple of auctions. The last column reports the smallest touple of auctions that results in a rejection of the null at the 5% level.

Table 7: Probit Regression - Validation Data

Probability a Pair of Firms Has both Members in the Same Cartel	(1)		(2)	
Common Personnel	0.94	(0.21)***	1.67	(0.32)***
Common Owner	0.07	(0.46)	-0.04	(0.50)
Common Manager	-0.67	(0.49)	-0.48	(0.38)
Common Zipcode	0.18	(0.27)	0.12	(0.53)
Common Municipality	-0.06	(0.21)	-0.03	(0.20)
Common County	0.33	(0.19)*	0.35	(0.20)*
Subcontract	0.88	(0.15)***	1.89	(0.40)***
Winning Consortium (All Piedmont Contracts)	0.46	(0.23)**	1.66	(.76)**
Bidding Consortium (Validation Data)	1.01	(0.14)***	-2.15	(.94)**
(1 - Common Personnel) x Common Zipcode			0.01	(0.53)
(1 - Common Personnel) x W.Consortium			-0.59	(0.75)
(1 - Common Personnel) x B.Consortium			1.41	(0.61)**
(1 - Common Zipcode) x W.Consortium			-0.48	(0.55)
(1 - Common Zipcode) x B.Consortium			0.07	(0.26)
(1 - Subcontract) x W.Consortium			0.94	(0.45)**
(1 - Subcontract) x B.Consortium			0.97	(0.50)*
(1 - W.Consortium) x B.Consortium			1.85	(0.59)***
Constant	-2.23	(0.17)***	-3.29	(0.42)***
Prob. Chi2	0.000		0.000	
Obs.	775		775	

Significance level: * is 10%; ** is 5%; *** is 1%. Table 7 presents probit coefficients and, in parenthesis, their standard errors corrected following Conley (1999) for the correlation across any pairs that share firms. The dataset consists of pairs of firms from the Validation data. The dependent variable equals one if the pair belongs to the same cartel and zero otherwise. As regards the independent variables, they are all dummy variables. The first three variables listed in Table 7 are equal one if the couple shares, respectively, any white collar worker, any owner (regardless of the shares owned) or any top manger (regardless of his exact role). The following three variables equal one if the firms' headquarters are located, respectively, at the same zip code, in the same municipality or in the same county. Subcontract equals one if the couple ever exchanged a subcontract. Winning Consortium equals one if the couple has won as a legal temporary bidding consortium at least one contract for public works held in Piedmont between 2000 and 2003. Bidding Consortium, instead, equals one if the pair of firms ever bid in the Validation data as a legal temporary bidding consortium. All pairs of firms considered are linked by at least one variable. However, we drop all couples that are only linked by location (either Zipcode, Municipality or County). Model (2) differs only in that it includes interactions.

Table 8: Assigned Groups - Validation Data

Good Data Scenario						
Assigned Group	Known Cartel	Members Cartel	Members Other Groups	Non Suspects	Auctions Won	Detection
1	B	13	5	11	106	Both
2	B	1	0	3	6	No
3	B	1	1	2	5	No
4	C	4	0	3	15	Both
5	G	3	0	1	12	No
6	A	3	1	7	7	Part
7	E	10	0	7	6	Bid
8	F	2	0	2	4	No
9	H	3	0	2	0	Both
10	-	0	0	4	3	No
11	-	0	0	3	2	No
12	-	0	0	2	1	No
13	-	0	0	2	1	No
14	-	0	0	4	0	No

Table 8 shows the groups obtained by applying the 3-step procedure described in the text. The firms for which we construct their full network of connections are those in the top 10% of participation of the Validation data auctions. The first column in the table reports an identifier for the group generated by the clustering algorithm. The second column reports the identifier of the cartel to which most of the firms in the assigned group are affiliated. The third column reports the number of firms belonging to the group in column 2. The following two columns describe who are the other members: the fourth column reports the number of members belonging to some cartel different from that in column 2 and the fifth reports the number of members not belonging to any of the 8 cartels. The sixth column reports the number of victories by the members of the group. The last column reports whether detection occurs only via the participation test (Part), only via the bid test (Bid), through both of them (Both) or whether no detection occurs (No). All tests are at the 5% level.

Table 9: Assigned Groups - Validation Data

Poor Data Scenario		
Cartel	No. of Groups	Avg. No. Cartel Firms
B	13	4.53
C	8	2.37
G	5	1.20
A	3	1.00
E	3	1.30
D	3	1.00
F	0	0.00
H	0	0.00
-	21	-

Table 9 shows the ‘in-sample’ results of our group construction method for the poor data scenario. Using the Validation data, we start by selecting the top 10% of firms in terms of participation. We fix a size of the groups equal to 6 and we use each of the top 10% of firms as a group-head. The remaining 5 members of each group are found using the iterative method described in the text. We obtain 56 different groups. The first column of the table reports the cartel to which the group-head belongs. The second column reports how many groups have the group-head belonging to this cartel. The last column reports the average number of firms across these groups that belong to the same cartel of the group-head. For instance, for cartel B the second column of the table indicates that there are 13 groups for which a member of this cartel is the group-head. The last column, instead, reports that across these 13 groups there are on average 4.53 firms belonging to cartel B. Notice that only cartel B is well represented by the groups constructed with this methodology. Moreover, there are 21 groups whose group-head does not belong to any cartel.

Table 10: Detection Results in the Main Data

	Detected Groups				
Rejected Test	N.groups	Group size	Entry	No. Wins	Revenue
Participation	42	10	45.2	0.82	350,231
Part.+Bid (Single-A.)	8	16	53.2	1.00	398,296
Part.+Bid (Multi-A.)	5	16	59.0	1.08	462,914

Table 10 reports the groups detected. Using the participation test at the 5% level, a rejection is found for 42 groups. Among these groups, 8 are such that a rejection of the single-auction (either one or two-tailed) bid test at the 5% level is recorded in at least 30% of the auctions entered. When, instead, in addition to a rejection of participation we consider a rejection of the multi-auction bid test, we detect 5 groups. The exact details of this latter test are presented in Table A.6 in the web appendix. "Group size" is the average of the size of the groups. The last three columns report means calculated across all firms in the groups for: entry, the number of victories and the revenues.

Table 11: Firms' Size and Gender Composition

	Not Entering FPA			Entering FPA		
VARIABLES	Mean	SD	N	Mean	SD	N
Independent Firms:						
Capital	216.7	777.7	585	336.0	1,052	599
Revenues	6,296	13,185	433	8,652	28,012	423
Profits	115.3	1,184	430	116.2	461.1	427
Number of Workers	28.23	47.79	527	30.18	58.12	532
Firm Age	23.64	13.56	583	21.32	14.57	593
Proportion of Women	0.145	0.206	582	0.151	0.212	593
Number Female Owners	0.143	0.452	582	0.140	0.458	593
Proportion Female Owners	0.032	0.104	582	0.035	0.108	593
Number Female Managers	0.475	0.957	582	0.499	0.947	593
Proportion Female Managers	0.077	0.957	582	0.079	0.163	593
Firms Belonging to the 42 Groups:						
Capital	313.8	584.1	159	882.9	2,280	139
Revenues	7,313	5,375	127	14,786	19,454	115
Profits	88.40	264.8	127	186.8	485.7	115
Number of Workers	32.18	27.29	147	49.16	59.74	134
Firm Age	27.84	14.62	158	28.81	15.82	136
Proportion of Women	0.157	0.189	158	0.155	0.187	136
Number Female Owners	0.113	0.409	158	0.105	0.352	136
Proportion Female Owners	0.025	0.095	158	0.025	0.082	136
Number Female Managers	0.619	1.103	158	0.550	0.982	136
Proportion Female Managers	0.069	0.138	158	0.065	0.142	136

Table 11 reports statistics for 4 groups of firms: (i) independent firms that never bid in FPAs (top left), (ii) independent firms that bid in FPAs (top right), (iii) group members that never bid in FPAs (bottom left) and (iv) group members that bid in FPAs (bottom right). Firms are classified as group members if they belong to any one of the 42 groups described in the top row of Table 10. Firms are classified as entrants in FPA if they bid at least in one FPA. Instead, a firm is classified as not entering FPAs if the firm never bid in any FPA but bid in at least 3 ABAs held in counties where at least 3 FPAs for which the firm was qualified to bid were held. For each of the 4 groups, the columns "Mean" and "SD" are the average and standard deviation taken across all firms in the group. The column "N" is the number of firms considered. "Firm Age" is the number of years between the beginning of activity and 2010. All other variables are averages over the years 2006-2010. "Capital", "Revenues" and "Profits" are in €1,000. "Number of Workers" is the number of all dependent workers. "Proportion of Women" is the fraction of female white collar workers over all white collar workers. "Number Female Owners (Managers)" is the number of female owners (managers). "Proportion Female Owners (Managers)" is the ratio of the number of female owners (managers) to that of the total number of owners (managers).

Figure 1: An Illustration of the Italian ABA

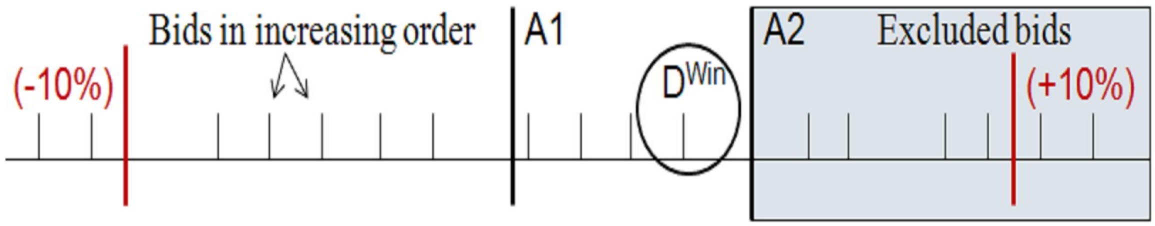

Figure 1 is taken from Decarolis (2011) and shows an example of an ABA with 17 bids. Bids, which are discounts over a reserve price, are represented by the 17 small vertical bars. Discounts are ordered in increasing order. The trim mean (A1) is calculated disregarding the 10 percent of the lowest and highest bids. In the figure, these discounts disregarded to compute A1 are the two discounts to the left of the bar marked "-10%" and the two discounts to the right of the bar marked "+10%". The threshold (A2) is calculated as the mean of all the bids greater than A1 but lower than the highest top 10% of bids. The highest discount below A2 wins: this discount is indicated as D^{win} in the figure. All discounts equal or greater than A2 are excluded.

Figure 2: Localization of the ABA

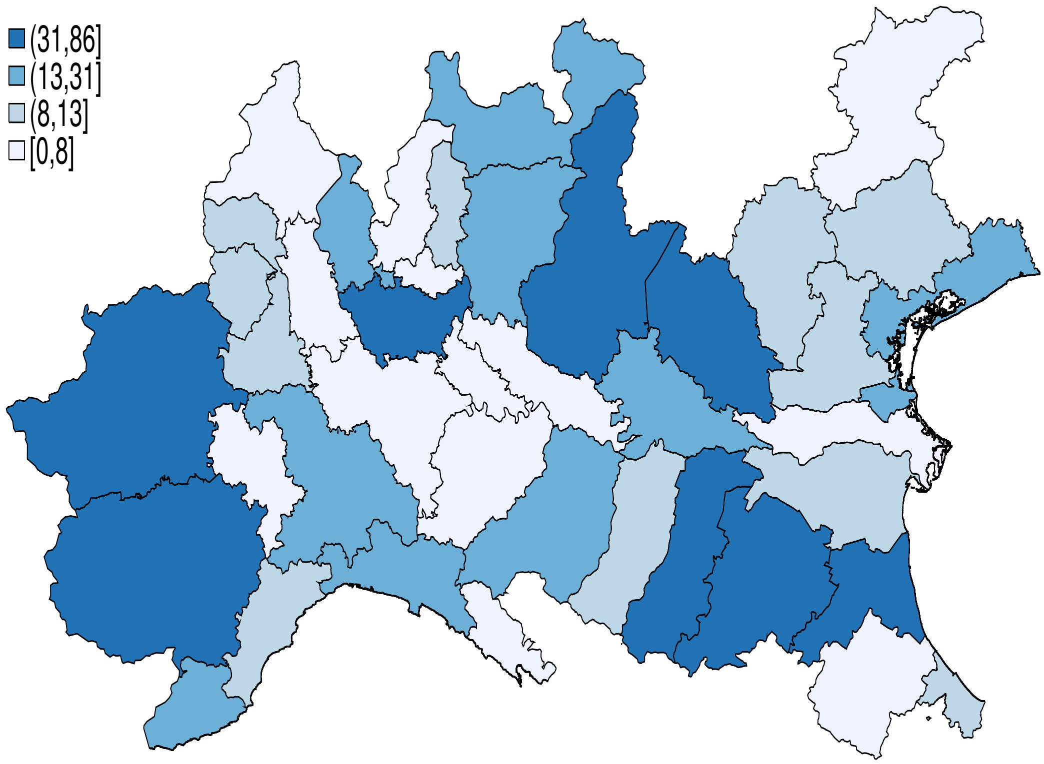

Figure 2 shows the location of ABAs in the Main data. The map divides the five Northern regions studied into their counties. The darker areas indicate a greater number of ABAs in the data and correspond to the more densely populated areas.

Figure 3: Localization of the 8 Cartels

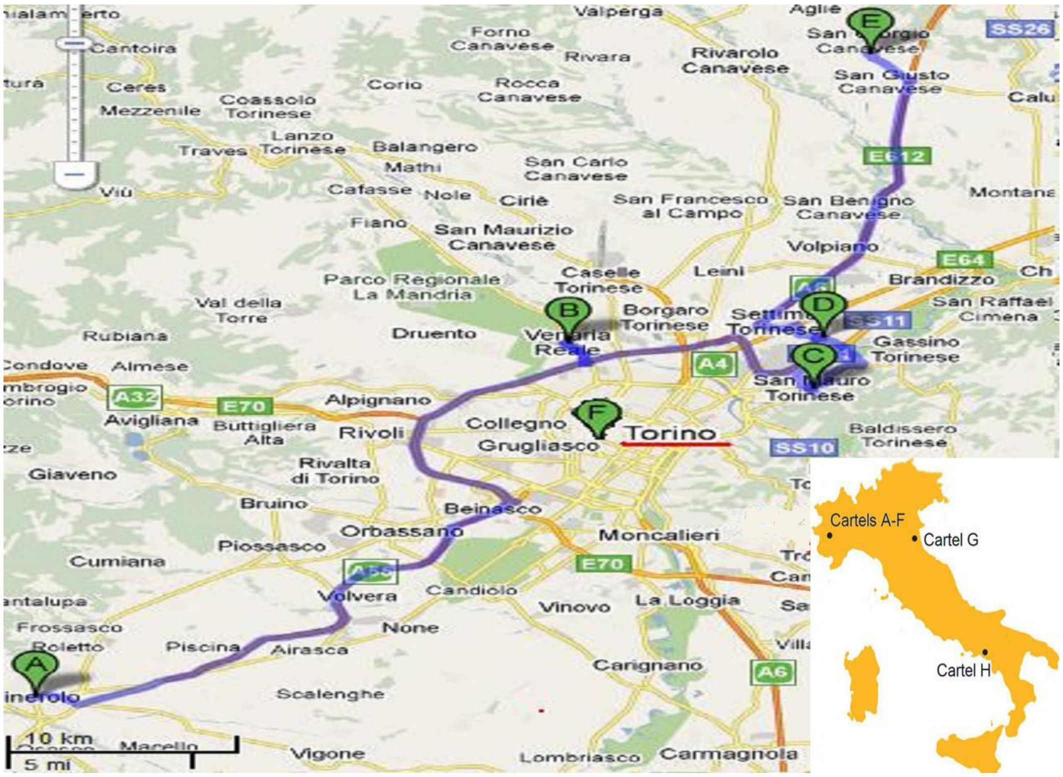

Figure 3 shows the location of the 8 cartels of the Validations data. As in Table 3, the capital letters from A to H indicate the different cartels. The map of Italy in the bottom right corner shows that 6 cartels are located in the North-West, one in the North-East and one in the Center. The large map shows the location of the 6 cartels in the North-West, which are all based within 50 kilometers from Turin.

Figure 4: Example of an ABA in the Validation Data

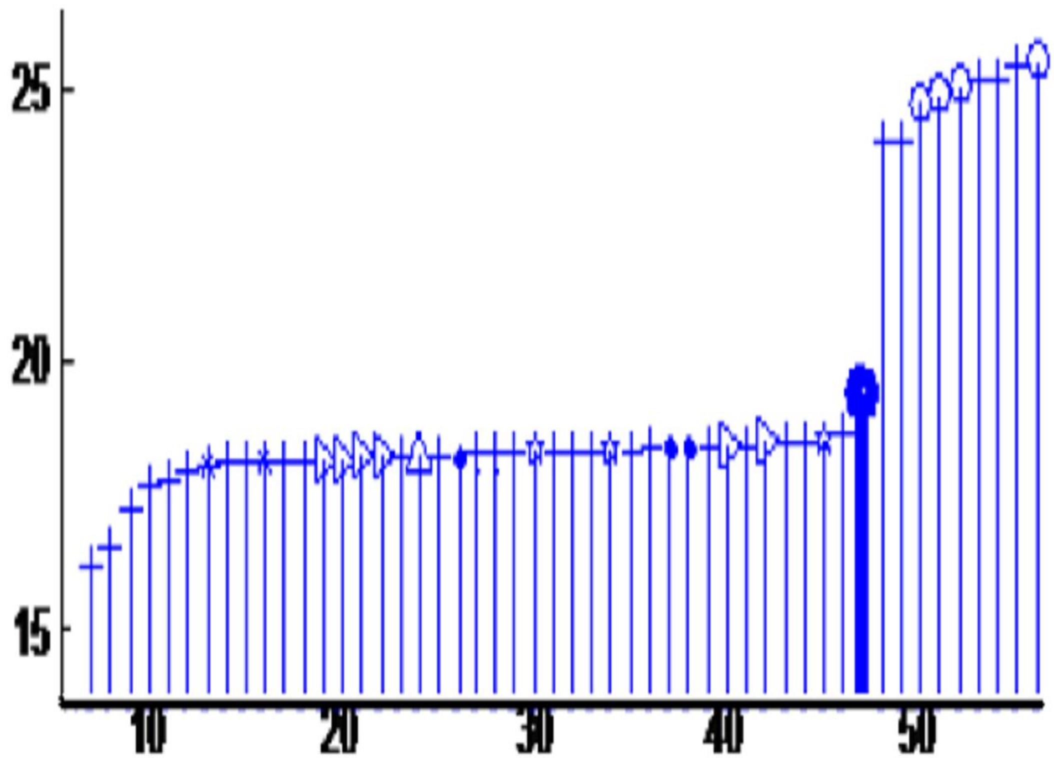

Figure 4 shows an example of one ABA in the Validation data. On the vertical axis there is the discount (bid). The horizontal axis lists all bidders in increasing order of their discount. The different symbols mark different cartels, but the cross indicates independent firms. The thick blue line marks the winner. The majority of bids lie close to the 18% approximate mode. The nine highest bids comply with the description of 'supporting bids' offered by the convicted firms and reported in the text.

Figure 5: Participation Test - Validation Data

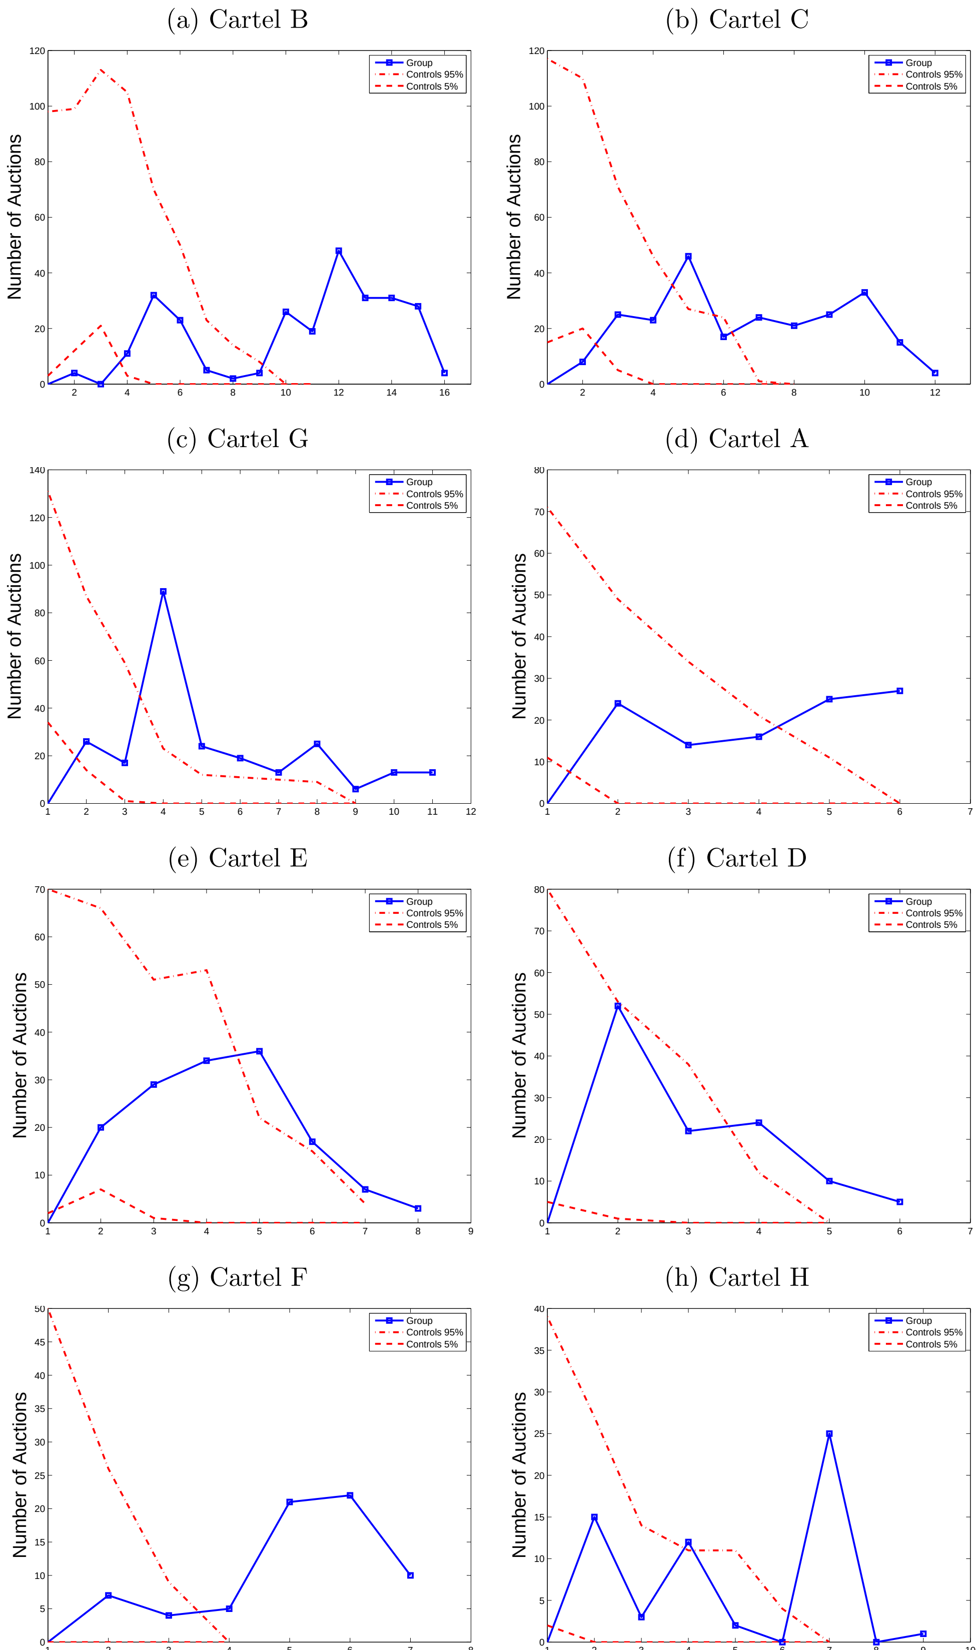

Figure 5: Participation test for all cartels and all of their possible subgroups.

Figure 6: Single-auction Bid Test - Validation Data

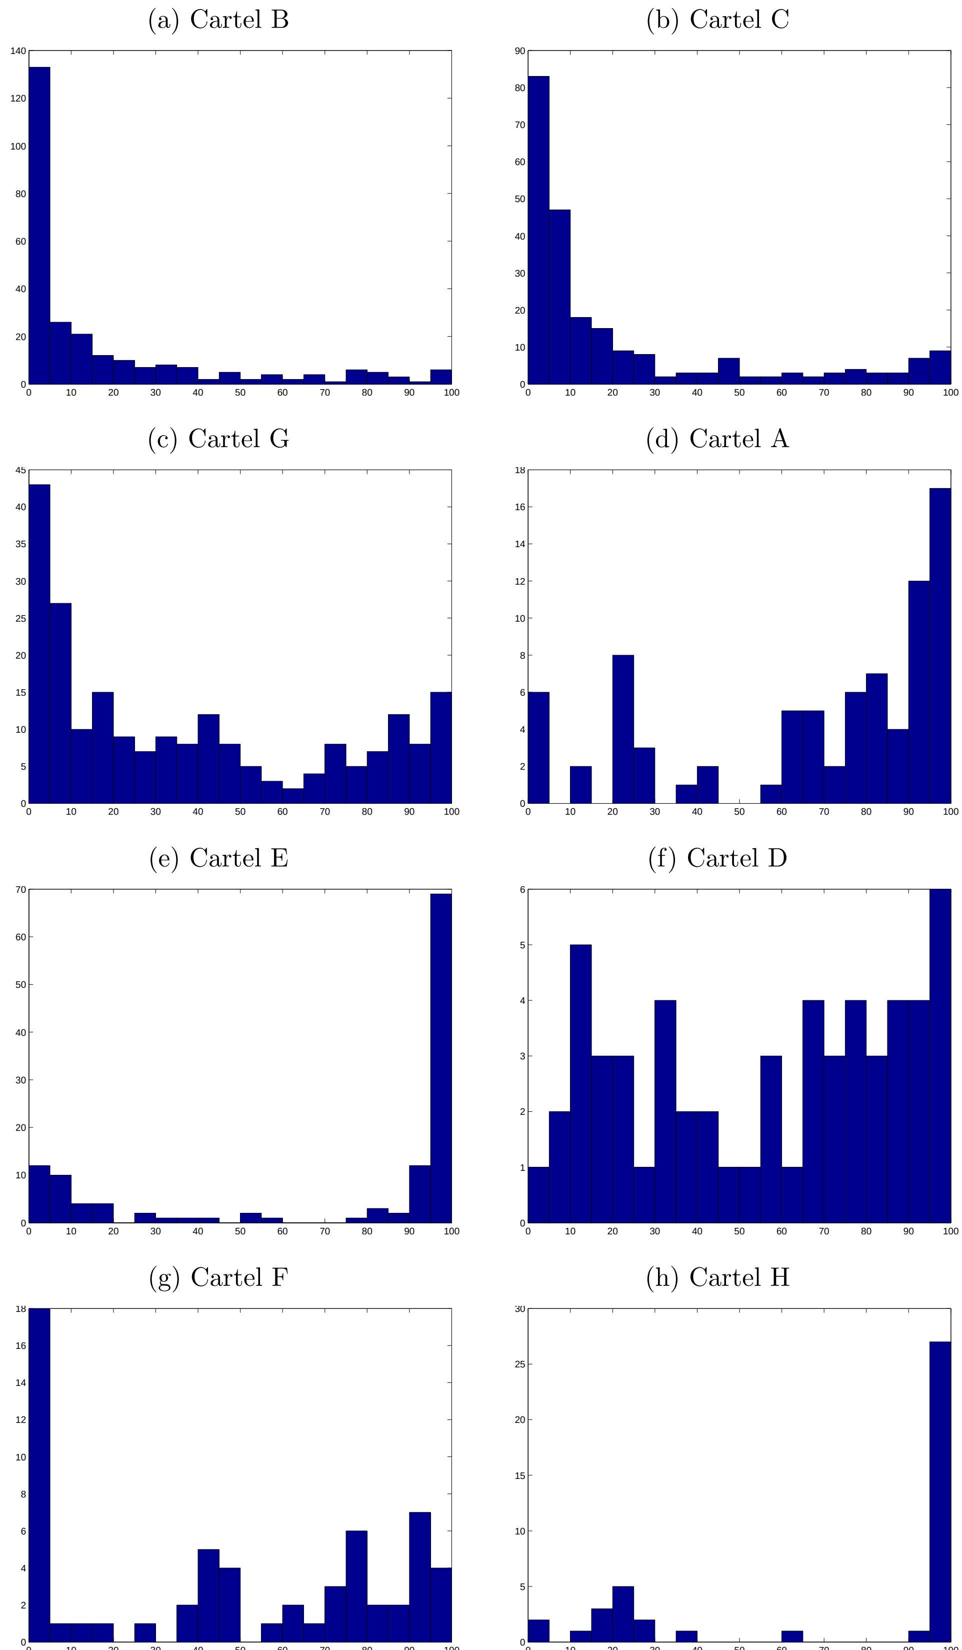

Figure 6: Histograms of P-*values* of single-auction bid tests for all cartels and all auctions.

Figure 7: Number of Bids in ABAs and FPAs

- The 2006 Reform in Four PAs -

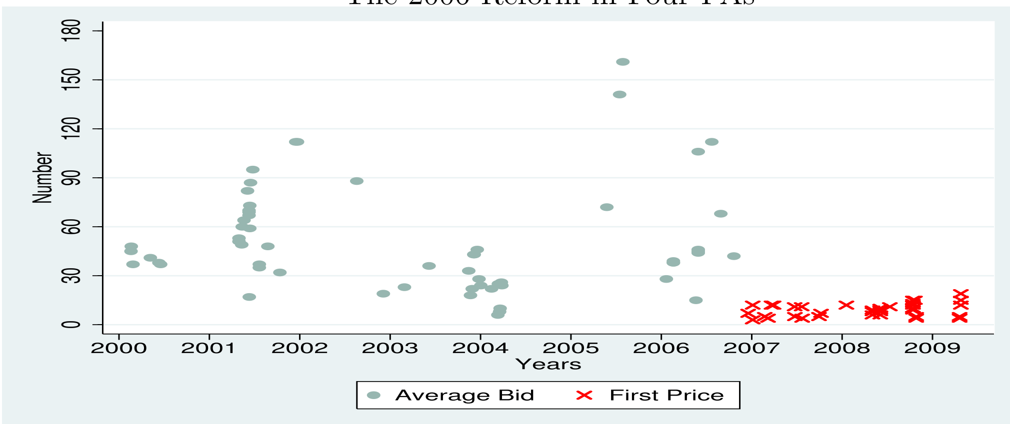

- The 2003 Reform in Turin -

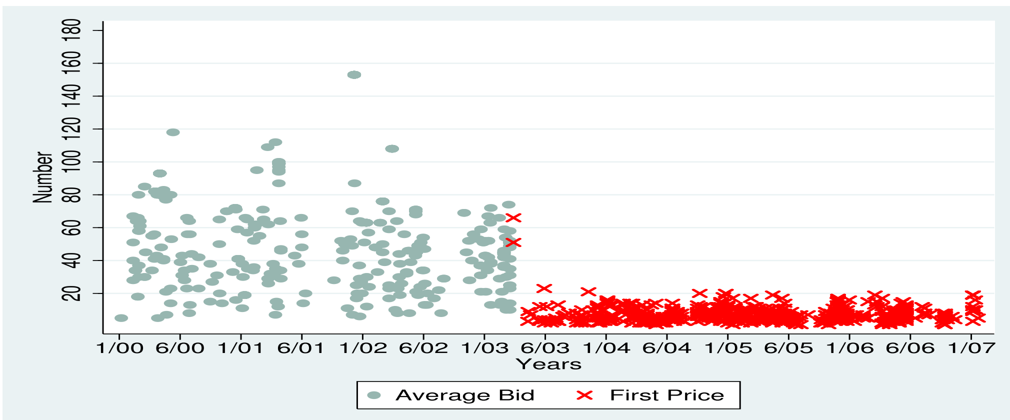

Both panels of Figure 7 report time on the horizontal axis and the number of bidders on the vertical axis. ABAs are marked with a blue circle and FPAs with a red cross: each circle (cross) represents one auction. The top panel plots the number of bidders in the auctions held by four PAs in the Main data: Padova, Varese, Sondrio and Cremona, which all switched to FPA once in 2006 the European Union mandated the liberalization of the use of this format. The bottom panel, instead, describes the evolution in the number of bidders in the auctions held by the county and the municipality of Turin: in 2003, after the collusion case became public, they switched from ABA to FPA for all of their auctions of public works. The figure in the bottom panel uses the records of the Italian Authority for the Vigilance on Public Contracts (AVPC) for all the auctions involving roadwork jobs affected by the reform. The Validation data are a subset of the ABAs recorded by the AVPC.

For Publication on the Authors' Web Page

Detecting Bidders Groups in Collusive Auctions

Web Appendix

Data

The Main data was assembled using the information released by a private company:

`http://www.telemat.it/`

This is one of the two largest information entrepreneur (IE) in the Italian market. It resells to construction firms both the "auction notice" (a document describing of the job features) and the "auction outcome" (a document reporting bids and bidders identities) that it collects from all Italian public administrations. We consider only ABAs for roadwork contracts held between 2005 and 2010 by counties and municipalities in five Northern regions.
As regards the Validation data, it comprises all the auctions of the municipality of Turin:

`http://www.comune.torino.it/en/`

cited in the court case (Turin Court of Justice, 1st criminal Section, sentence N. 2549/06 R.G., 04/28/2008). These are auctions for roadwork jobs held between 1999 and 2003.
The source for the data about public administrations is Italy's National Statistical Institute:

`http://demo.istat.it/index_e.html`

In particular, we used the freely available data on geographic location and demographic characteristics of Italian counties and municipalities. The single year of data employed is 2006. Finally, the data on firms' characteristics comes from the Italian Registry of Firms:

`http://www.infocamere.it/eng/about_us.htm#`

The version of this dataset used is that compiled by the Bank of Italy which keeps track also of those firms that the Registry cancels once they cease activity. The data is organized as a panel with four years starting from 2006. The "Capital" variable used in Table 10 is the value of the firm's subscribed capital at the date of the panel closest to that of the auction. Finally, all measures of distance between firms and PAs were obtained using their zip code as input for the freely available API of:

`http://classic.mapquest.com`.

Table A.1: Probit Regression for the Probability of Entry

	FPA	FPA	ABA	FPA	ABA
	(1)	(2)	(3)	(4)	(5)
Log(Miles Firm-Work)	-0.840***	-0.811***	-0.863***	-0.845***	-0.861***
	(0.021)	(0.027)	(0.007)	(0.023)	(0.008)
Log(Firm Capital)	0.062***	0.047***	0.052***	0.062***	0.052***
	(0.012)	(0.001)	(0.004)	(0.012)	(0.004)
Backlog	0.123	0.073	0.069	0.126	0.071
	(0.211)	(0.301)	(0.046)	(0.211)	(0.047)
Unlim. Liability Firm	0.448	0.281	0.044**	0.455**	0.046**
	(0.193)	(0.357)	(0.022)	(0.153)	(0.023)
No. of Workers (x100)	0.004	0.012	-0.066	-0.002	-0.067***
	(0.055)	(0.068)	(0.017)	(0.055)	(0.017)
Firms Links	No	No	No	Yes	Yes
R-squared	0.41	0.40	0.31	0.41	0.32
Observations	11,810	9,249	80,275	11,806	80,274

Significance level: * is 10%; ** is 5%; *** is 1%. Table A.1 reports the result of probit regressions using the Main data and where the dependent variable is 1 if the firm bids in the auction and zero if the firm does not bid but is a potential participant. The latter condition of potential participant is satisfied if at the same time: (i) the firm has the legal qualification to bid, (ii) it has bid at least once in the county where the auction is held and (iii) it has bid at least once in the region where the auction is held in the same year of the auction. All regressions include: a constant, six dummies for the categories of value of the reserve price and dummies for each year, the PA region and the firm region. Backlog is constructed following Jofre-Bonet and Pesendorfer, "Estimation of a Dynamic Auction Game," *Econometrica*, 2003, 71(5), 1443-1489. In the first three columns the specification is the same but the sample differs: all FPAs are used in (1), only the FPAs held in the county and municipality of Turin are used in (2) and all ABAs are used in (3). The last two columns extend, respectively, model (1) and (3) by including 10 additional variables. For each firm i in auction j, these variables count, how many other firms bidding in auction j are linked to firm i separately by each one of the 9 links in Table 7 (common personnel, common owner, common manager, common zip code, common municipality, common county, subcontracts, winning consortium and bidding consortium). Analogously, the tenth variable counts for each firm i and each auction j how many other firms entering auction j are registered at the same street address of firm i.

Table A.2: OLS Regressions for the Bids

	FPA	FPA	ABA	ABA
	(1)	(2)	(3)	(4)
Log(Miles Firm-Work)	-.647**	-0.301*	0.263***	0.045
	(0.250)	(0.174)	(0.077)	(0.043)
Log(Firm Capital)	0.054	0.122**	0.006	-0.002
	(0.053)	(0.046)	(0.015)	(0.005)
Backlog	0.882	-0.438	0.056	0.179*
	(.944)	(1.168)	(0.148)	(0.098)
Unlim. Liability Firm	-3.087*	-0.053	-0.095	-0.055
	(1.651)	(1.171)	(0.157)	(0.078)
No. of Workers (x100)	0.022	0.234	-0.164**	-0.017
	(0.867)	(0.798)	(0.074)	(0.039)
Auction Fixed Effects	No	Yes	No	Yes
R-squared	0.21	0.55	0.13	0.65
Observations	2,182	2,182	45,513	45,513

Significance level: * is 10%; ** is 5%; *** is 1%. Standard errors are clustered by PA and year. Table A.2 presents OLS estimates obtained using the discount offered in the Main data by each firm in each auction as the dependent variable. All regressions include: a constant, six dummies for the categories of value of the reserve price and dummies for each year and region of the auction. The first two column report the results using a dataset consisting of only FPAs, while the latter two columns report results based on the sample of only ABAs. The difference between (1) and (2) (and between (3) and (4)) is that the latter includes auction fixed effects while the former does not.

Table A.3: Constructing Groups - Validation Data (95 Suspect Firms)

	Variable $\hat{p}$					$\hat{p} = 0.2$			
Assigned Group	$\hat{p}$	Known Cartel	Group Member	Auctions Won	Detect	Known Cartel	Group Member	Auctions Won	Detect
1	0.45	E	(8,0,1)	5	Bid	E	(10,1,7)	8	Bid
2	0.1	A	(2,0,6)	3	No	A	(2,0,4)	2	No
3	0.3	H	(7,0,2)	1	Both	H	(8,0,3)	1	Both
4	0.15	G	(5,0,1)	13	No	G	(5,0,1)	13	No
5	0.25	C	(4,0,4)	20	Both	C	(7,0,6)	25	Both
6	0.55	B	(8,0,0)	66	Bid	B	(11,4,8)	85	Bid
7	0.25	C	(3,1,1)	19	Bid	C	(3,2,10)	21	No
8	0.15	D	(3,1,5)	11	Bid	D	(3,1,5)	11	No
9	0.1	F	(2,1,7)	4	Part	F	(2,0,4)	4	Both
10	0.15	G	(2,1,4)	2	Both	G	(2,1,3)	2	Both
11	0	H	(1,0,6)	3	No				
12	0	A	(1,1,8)	7	No				
13	0	G	(2,0,8)	2	No				
14	0	C	(1,0,8)	4	No				
15	0	C	(1,2,3)	3	Part				
16	0.1	H	(1,0,7)	0	No				
17	0	H	(1,0,7)	3	Part				

Table A.3 shows the groups constructed under our "Good Data Scenario" when we apply our 3-step method. We apply this method to the Validation data with the intent to show the performance of the proposed method and to illustrate a possible alternative. The right panel is the analogue of Table 8 with the only difference that the list used to construct the network of connected firms is not the top 10% participants in the Validation data but the official list of suspect firms. Notice that, relative to Table 8 we get only 10 groups, instead of 14, and that they match more precisely the cartels. As in Table 8, the first column of the right panel reports the identity of the most represented cartel. Next, the triplet in the parenthesis reports the number of members of this cartel, the number of members of all other cartels and the number of independent firms. The last two columns report the number of victories and whether the groups is detected as such by our bid test (Bid), participation test (Part.), both of them (Both) or it is not detected (None). Detection is defined as in Table 8. The left panel is analogous to the right panel with the only exception that to get the final assignment to groups we do not impose a fix cutoff which excludes firms that have a predicted probability of being together in a cartel below 20% (i.e., $\hat{p} = 0.2$), as done in the right panel and in Table 8. Instead, we impose that the assigned group has at most a size equal to 10 and so, whenever the clustering algorithm returns a group larger than 10, we trim it by increasing the $\hat{p}$ until the size is no greater than 10. The first column in the left panel reports the $\hat{p}$ we used to trim each group. The structure of the rest of this panel is the same of the right panel.
Remarks: (1) Step 1 of the procedure estimates the probit using 7 cartels, excluding cartel D which was not convicted for repeated violations as the others; (2) Step 3 uses the default hierarchical clustering algorithm of Stata (cluster) using the average distance between clusters as the aggregation criterion and a cutoff (cutt) of .997; (3) to refine the clusters, we pick the most connected firms (using a fixed $\hat{p} = 0.2$); (4) we test the validity of the groups using the Monte Carlo approach described in chapter 7 of Gordon (1999) rejecting that the groups are identical to random groups of firms at a 5% level.

Table A.4: Muti-Auction Bid Test Results - 8 Cartels (Validation Data)

	Cartel B - size 5 (184 auctions)							Cartel C - size 5 (51 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.02	0.13	0.47	0.01	0.09	0.43	739	0.03	0.11	0.38	0.02	0.08	0.33	531
4	0.00	0.06	0.20	0.00	0.05	0.19	574	0.01	0.06	0.20	0.01	0.07	0.36	399
6	0.00	0.06	0.19	0.00	0.05	0.16	354	0.01	0.03	0.14	0.01	0.05	0.26	311
8	0.00	0.04	0.13	0.00	0.04	0.13	727	0.00	0.01	0.04	0.00	0.03	0.16	278
10	0.00	0.03	0.09	0.00	0.02	0.08	56	0.00	0.01	0.03	0.00	0.02	0.12	209
	Cartel G - size 4 (68 auctions)							Group A - size 7 (10 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.03	0.22	0.72	0.02	0.18	0.56	728	0.00	0.00	0.04	0.98	1.00	1.00	45
4	0.03	0.23	0.58	0.04	0.26	0.63	831	0.00	0.00	0.02	0.98	1.00	1.00	206
6	0.03	0.15	0.54	0.06	0.19	0.46	621	0.00	0.00	0.02	0.99	1.00	1.00	207
8	0.04	0.18	0.43	0.07	0.21	0.47	455	0.00	0.00	0.02	0.99	1.00	1.00	45
10	0.04	0.18	0.37	0.09	0.24	0.41	313							
	Cartel E - size 5 (20 auctions)							Cartel D - size 4 (19 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.01	0.11	0.47	0.02	0.23	0.98	999	0.20	0.43	0.64	0.39	0.75	0.98	160
4	0.00	0.08	0.17	0.88	0.97	1.00	466	0.17	0.39	0.63	0.60	0.83	0.99	482
6	0.02	0.05	0.14	0.96	1.00	1.00	615	0.21	0.39	0.56	0.70	0.84	0.98	280
8	0.02	0.06	0.14	0.95	1.00	1.00	427	0.26	0.38	0.52	0.73	0.84	0.96	127
10	0.02	0.05	0.11	1.00	1.00	1.00	66	0.27	0.40	0.53	0.76	0.85	1.00	40
	Cartel F - size 5 (21 auctions)							Cartel H - size 5 (25 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.00	0.05	0.45	0.00	0.51	0.88	199	0.00	0.04	0.32	0.82	0.97	1.00	289
4	0.00	0.02	0.22	0.03	0.52	0.92	822	0.00	0.03	0.12	0.89	0.97	1.00	965
6	0.00	0.02	0.10	0.13	0.47	0.87	938	0.00	0.02	0.09	0.95	0.98	1.00	997
8	0.00	0.01	0.06	0.17	0.44	0.82	972	0.00	0.01	0.04	0.95	0.99	1.00	999
10	0.00	0.01	0.04	0.26	0.50	0.79	984	0.00	0.01	0.04	0.96	0.99	1.00	999

Table A.4 shows the distribution of the result of the multi-auction bid test for subgroups of the eight groups of the Validation data. The table reports the size of the subgroup and the number of auctions in which the firms in the chosen subgroup jointly participated. The leftmost column indicates whether the test was conducted looking at a pair, triplet, quadruplet, etc. of auctions. The next three columns report three percentiles (10^{th}, 50^{th} and 90^{th}) of the distribution of the two-sided version of the test. The following three columns report the same information for the one-sided (left) version of the test. All numbers should be read as p-values of a test with a null stating that the group analyzed is not different from the control groups. The distribution of the test's results is reported because there is a multiplicity of pairs, triplets, quadruplets, etc. for which the test can be conducted. The number of these pairs, triplets, quadruplets, etc. used is reported in the eight column. In all these pairs, triplets, etc. of auctions, we observe at least 30 firms in common and, hence, that allows the construction of a large number of random groups.

Table A.5: Muti-Auction Bid Test Results - Independents as Controls (Validation Data)

	Cartel B - size 5 (184 auctions)							Cartel C - size 5 (51 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.01	0.13	0.43	0.00	0.02	0.07	625	0.02	0.10	0.47	0.00	0.00	0.06	407
4	0.00	0.07	0.27	0.00	0.00	0.04	318	0.00	0.04	0.12	0.00	0.00	0.03	212
6	0.00	0.03	0.11	0.00	0.00	0.04	158	0.00	0.02	0.08	0.00	0.00	0.00	102
8	0.00	0.02	0.12	0.00	0.00	0.04	31	0.00	0.02	0.05	0.00	0.00	0.00	74
10								0.00	0.02	0.06	0.00	0.00	0.00	42
	Cartel G - size 4 (68 auctions)							Cartel A - size 7 (10 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.05	0.32	0.82	0.00	0.08	0.71	676	0.00	0.00	0.03	0.96	0.99	1.00	45
4	0.01	0.20	0.77	0.00	0.05	0.43	464	0.00	0.00	0.02	0.98	1.00	1.00	206
6	0.03	0.21	0.74	0.00	0.06	0.45	200	0.00	0.00	0.02	0.98	1.00	1.00	207
8	0.01	0.14	0.51	0.00	0.04	0.26	105	0.00	0.00	0.02	0.98	1.00	1.00	45
10	0.02	0.12	0.40	0.00	0.04	0.17	48							
	Cartel E - size 5 (20 auctions) - OLD							Cartel D - size 4 (19 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.02	0.12	0.42	0.28	0.95	1.00	94	0.23	0.50	0.70	0.08	0.50	1.00	114
4	0.00	0.10	0.26	0.87	0.97	1.00	468	0.22	0.46	0.69	0.11	0.59	1.00	451
6	0.00	0.03	0.16	0.95	1.00	1.00	623	0.21	0.43	0.60	0.27	0.68	0.93	265
8	0.00	0.02	0.14	1.00	1.00	1.00	428	0.25	0.41	0.56	0.30	0.64	0.98	120
10	0.00	0.00	0.08	1.00	1.00	1.00	63	0.28	0.39	0.50	0.34	0.64	0.89	31
	Cartel F - size 5 (21 auctions)							Cartel H - size 5 (25 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.00	0.07	0.44	0.00	0.43	0.80	169	0.00	0.05	0.36	0.61	0.96	1.00	288
4	0.00	0.03	0.27	0.00	0.37	0.98	766	0.00	0.04	0.18	0.74	0.95	0.99	943
6	0.00	0.02	0.13	0.05	0.32	0.91	875	0.00	0.02	0.11	0.83	0.95	1.00	946
8	0.00	0.02	0.09	0.12	0.46	0.89	891	0.00	0.03	0.08	0.88	0.96	1.00	910
10	0.00	0.02	0.07	0.18	0.48	0.79	886	0.00	0.03	0.06	0.88	0.96	1.00	949

Table A.5 has the same structure of Table A.4. The only difference relative to that table is that here the control groups used to test the 8 cartels of the Validation data are constructed using exclusively firms outside the list of 95 suspect firms. For most of the cartels this implies that we achieve detection (using the median of the two-sided test results) with a smaller touple of auctions relative to Table A.4, but also that less auctions have enough independent bidders to be usable in the analysis.

Table A.6: Muti-Auction Bid Test Results - Assigned Groups (Main Data)

	Group 1 - size 5 (63 auctions)							Group 2 - size 5 (83 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.00	0.08	0.29	0.00	0.21	0.86	808	0.00	0.09	0.68	0.00	0.36	0.89	987
4	0.00	0.03	0.16	0.00	0.05	0.46	999	0.00	0.06	0.48	0.02	0.36	0.91	999
6	0.00	0.03	0.16	0.00	0.02	0.36	254	0.01	0.09	0.38	0.13	0.58	0.93	165
8	0.00	0.03	0.15	0.00	0.05	0.43	502	0.00	0.05	0.27	0.25	0.61	0.92	414
	Group 3 - size 9 (95 auctions)							Group 4 - size 5 (52 auctions)						
	Two-sided			One-sided Left				Two-sided			One-sided Left			
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}		10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}	
2	0.00	0.02	0.57	0.06	0.49	0.97	999	0.00	0.13	0.58	0.00	0.20	0.89	499
4	0.00	0.02	0.27	0.12	0.33	0.94	999	0.00	0.08	0.33	0.00	0.20	0.69	499
6	0.00	0.02	0.10	0.44	0.45	0.80	552	0.00	0.05	0.40	0.00	0.16	0.38	499
8								0.00	0.05	0.21	0.00	0.04	0.08	36
	Group 5 - size 4 (75 auctions)													
	Two-sided			One-sided Left										
	10^{th}	50^{th}	90^{th}	10^{th}	50^{th}	90^{th}								
2	0.05	0.20	0.65	0.01	0.08	0.44	999							
4	0.04	0.15	0.53	0.02	0.06	0.34	999							
6	0.03	0.15	0.39	0.02	0.06	0.29	999							
8	0.04	0.11	0.31	0.00	0.05	0.23	999							

Table A.6 shows the distribution of the result of the multi-auction bid test for the 5 groups of the Main data for which this test detects cooperation. The table has the same structure of Table A.4. See the note to that table for the description of the table content.