Detecting Bidders Groups in Collusive Auctions

Timothy G. Conley and Francesco Decarolis*

February 15, 2013

Abstract

We study entry and bidding in procurement auctions where 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 cooperation tests motivated by how a coalition bids to manipulate the mechanism. We show that our tests perform well in a validation dataset with known cartels. We also use them to investigate cooperation in a larger 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, Boston University. *E-mail address:* fdc@bu.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, Ali Hortacsu, Jakub Kastl, Jonathan Levin, Greg Lewis, Nicola Persico, Gustavo Piga, Rob Porter, Phil Reny, 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* [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 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 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 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 are used in several countries to procure public works.[1] ABAs are characterized by the fact that the winner is decided through an algorithm that eliminates all bids that are deemed 'too good to be true.' The ABAs we study have been used in Italy since 1999 and have a mechanism similar to that used in other countries. These ABAs are conducted by a public administration (PA) who 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. In contrast, 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 can reduce the risk of poor ex post performance, bidders' incentives are deeply distorted relative to a standard FPA.

Auctions like these are 'collusive auctions' because the fact that the awarding rule is 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. In this paper, we document that ABAs give strong incentives to bidders to coordinate their entry and bidding choices.

We begin by describing how competing firms should bid in an ABA and we show that in the unique equilibrium of the game all firms offer a discount of zero over the reserve price. Since the data reveals a very different behavior, we enrich our description allowing for the presence of coalitions of firms. Relative to non-cooperating firms, firms within the same cooperating group should gain by manipulating the average that determines the winner. By concentrating bids on the same side of the bid distribution, cooperating firms can pilot the

[1] A non-exhaustive list of these countries includes: Chile, China, Colombia, Italy, Japan, Peru, Malaysia, Switzerland, Taiwan and the US (specifically the Florida DoT and the New York State Procurement Agency).

thresholds that determine the winner and thus significantly increase the chances that one of them wins. This also implies that cooperating firms employing such an average-piloting technique have an incentive to coordinate their entry into the same auctions in order to have a large enough number of bids to manipulate the mechanism. Albano et al. (2006) and Engel et al. (2006) are the only other studies mentioning the risk of collusion that ABAs pose. Relative to these earlier theoretical studies, our paper presents the first empirical study of collusion in ABAs.

A main contribution of this paper is two statistical tests for cooperation based on these incentives induced by the ABA mechanism. Our entry test is motivated by the simple idea that, in order to engage in average-piloting, a group of firms must be present together in sufficient numbers in an auction. In contrast non-cooperating firms have no such incentive to jointly participate. Our test compares the frequency of joint entry for a suspect group with a collection of control groups whose members are comparable to firms in the suspect group in terms of various determinants of entry. Our bid test is motivated by a search for groups employing a cooperative strategy to pilot the relevant average bid toward one of their members. We exploit the exact rules of these ABAs 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 'mean piloting' for the suspect group versus that of a set of comparable, control groups. These tests are directly helpful for courts evaluating cooperation in ABAs or, with slight modifications, coordination mechanisms with similarly manipulable awarding rules.

When we apply our tests to known groups of cooperating firms, they perform well in detecting these groups. 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 Court of Justice ruled that these auctions had been rigged by 8 groups made up of 95 firms.[2] Each group strategically submitted bids to affect the awarding of the contract. These groups were identified as cartels and their members fined, with some of them even being 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 strongly support the capability 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 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

[2]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.

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 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 relative to the competitive case. This is because in the competitive case all firms offer a discount of zero versus observed discounts of about 13% of the reserve price on average. However, firms outside the cooperating groups are harmed. They are less likely to win and when they do, they get a worse price than under competition.

Finally, we present an illustration of how quantitatively important the bidders' reaction is to a change of the auction incentives. We analyze a change in the regulation that replaced ABAs with FPAs for certain types of contracts. We document that this change resulted in the exit of hundreds of firms from the market. Firms may exit because they are too inefficient to compete in FPAs or because they are 'shill' firms who existed for the sole purpose of allowing a controlling firm to place multiple bids. Clearly, placing multiple bids in very valuable in ABAs, but less so in FPAs. It is important to understand the relative frequency of these two motivations for exit to evaluate, for instance, the benefits of augmenting the FPAs with a system of subsidies for small firms. We investigate whether a classification into cooperating groups based on our cooperation tests can be useful to understand the frequency of shill firm among exiting firms. Our findings suggest that, among the 774 exiting firms, 159 of them (or 21%) belong to groups. We show that exiting firms belonging to groups display characteristics consistent with being shill firms.

Our study contributes to the empirical literature on collusion in auctions.[3] 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).[4] Both approaches have led to the flourishing of a literature 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. Thus, in essence, our approach implements the idea of Hendricks and Porter (1989) that collusion is intrinsically tailored to the specific rules of the environment where it takes place.

This paper is also closely related to a vast literature analyzing how firms and individuals respond to mechanisms similar to the ABAs. For instance, Abrantes-Metz et al. (2012)

[3]In auction design, collusion is generally regarded as a first order concern (Klemperer, 2004) and has received substantial attention from the theoretical literature. The seminal studies in the theoretical literature include Robinson (1985) addressing the strength to collusion in 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]See also Haberbush (2000) for a review of cases of collusion in U.S. public procurement auctions. Porter and Zona (1993) and Ishii (2009) analyze specifically collusion in auctions for roadwork contracts.

study the case of the LIBOR. This rate, to which contracts worth $300 trillion are linked, is a trimmed mean of bank quotes for interest rates. Evidence that several banks coordinated their quotes to manipulate this trim mean emerged in 2012. In health care markets, Scott Morton (1997) and Duggan and Scott Morton (2006) study how drug manufacturers distort prices in response to a regulation setting the mandatory rebate for Medicaid as an average of the drug prices faced by non-Medicaid enrollees. For Medicare Part D, Decarolis (2012) studies how insurance companies use the multiple plans that they offer to increase the subsidy paid by Medicare which, in turn, is a function of the average of plan premiums. For the Medicare auctions for durable medical equipment (DEMPOS), Cramton et al. (2011) study how firms respond to an awarding rule based on the median price offered. Even in the context of compensation schemes for agricultural workers similar rules exist. For example, the studies of Bandiera et al. (2005 and 2006) on a U.K. farm using such contracts have shown that workers learn how to cooperate to manipulate the average on which their payments depend.

Finally, this paper also offers two more general contributions. The first is to show that firms' response to incentives in these Italian procurement auctions is both highly sophisticated and quantitatively very large. This represents important evidence in favor of the growing literature in market design that advocates the use of accurately designed mechanism to achieve publicly desirable goals. The second, more general contribution is to 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.[5] Therefore, our results are useful for the design of antitrust regulations because they argue against the usage of automatic sanctions punishing all types of cooperation and in favor of a careful economic analysis of the markets.[6]

The outline of the paper is as follows: Section 2 provides a description of the market and our data, Section 3 presents a model of bidding in an ABA, 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 Main data and, finally, Section 7 concludes.

2 Description of the Market

In this Section, we describe both the institutions and our datasets. 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 because for the PA of Turin we have access to what we call Validation data as a result of legal cases where several firms were convicted for collusion in these auctions. These data are comparable in various aspects to

[5] As discussed in Harrington (2011), the main difference between the legal and economic definitions of collusion consists in the fact that while the former typically indicates collusion as every action that firms take to coordinate prices, the latter generally requires that prices resulting from firms coordinated choices are higher than the ones achievable under competition.

[6] A related example where the U.S. Supreme Court favored a careful analysis of the market instead of a rigid application of antitrust laws is Broadcast Music v. Columbia Broadcasting System, 441 U.S. 1.

the remainder of our data, which we refer to as our Main data.

For these roadwork contracts, PAs are typically required to select the contractor through sealed bid price-based auctions. A small fraction of these auctions are of the well known first price auction (FPA) type, but the vast majority are average bid auction (ABA). The regulations of ABAs and FPAs are identical in everything except for how the winner is identified.[7] 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 FPAs the highest discount wins, in ABAs 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.[8] Figure 1 offers an example with 17 bids: the winner is denoted D^{Win} and, in this case, is the 7^{th} highest discount. Notice that in ABAs the winner is paid his bid to complete the work.

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 awarded using ABAs, resulting in a total reserve price of the auctioned contracts of approximately €10 billion. Between July 2006 and May 2011, a series of reforms required by the European Union temporarily limited the use of ABAs and extended the use of the FPAs. However, even after these reforms ABAs remained the most frequently used procurement format. In this paper, we do not consider the auctioneer problem of choosing among auction formats, but we focus on the ABA to study firms 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.

Our Main data consists of 802 ABAs and 232 FPAs. Table 1 presents some summary statistics separately for the two types of auctions. Comparing the statistics for the two sets

[7] A detailed discussion of the regulation is contained in Decarolis (2009) and Decarolis et al. (2010).

[8] Ad hoc rules exist to deal with the special cases that can occur. First, if all bids are equal, the winner is selected with a fair lottery. Second, if there are no bids strictly greater than $A1$ and less than each of the highest 10% of bids, then the winner is the bidder with the highest discount among those not higher than $A1$. Third, a random draw is used to ensure that exactly 10% of the top/bottom bids are disregarded when, due to ties at the minimum/maximum values of these two sets of bids, more than 10% of bids would be in these sets. Finally, special rules apply when $N \leq 4$, but we ignore them since this never occurs in the data.

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 1, 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 higher capital 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. In April 2008 the Court of Justice of Turin convicted the owners and managers of numerous construction firms. The court documents identify a network of 95 firms that operated in 8 cartels.[9] 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 data. 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.[10] Proximity to Turin surely provides cost advantages for execution of road construction contracts. In Table 2 and throughout the remainder of the paper, to indicate each cartel, we use a capital letter, from A to H. Two cartels, G and H, despite having all members close to each other, are the only cartels 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, Table 2 shows that the 8 cartels are quite heterogenous in their size, entry and victories.

In addition to the asymmetries across cartels, there are also significant asymmetries

[9]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.

[10]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.

within cartels. The bottom panel of Table 2 reports summary statistics for both the firms inside and outside the cartels. Given that this sample was assembled to compare alleged colluders with non-cooperating 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 middle panel of Table 2 suggests that these auctions are similar to those in the Main data described in Table 1 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 1, 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 non-cooperating firms. Four 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 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 17 to 18%. Court documents report the cases of various defendants claiming that it was known to all players in this market that most of the discounts would be near this range. Figure 2 illustrates this for one Validation data auction. Individual bids are plotted in increasing order with discounts on the vertical axis. There is a clear mode in the distribution around 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 modes and winning bids. This basic pattern occurs in the ABAs in the Main data as well. For example, both the difference between the winning discount and the next discount and the within-auction standard deviation are similar in Main and Valdidation datasets (See Table 1 and Table 2). Furthermore, 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.[11] Decarolis (2009) finds that, across PAs in our Main data, in their ABAs there is a strong tendency for the winning bids to remain nearly identical across the auctions of the same PA.

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 competitively and in favor of there being cooperative groups. The predictability of the range for winning bids motivates our discussion of equilibria where non-cooperating firms can predict the range where winning bids will lie.

[11] The sources of information are both public and private: Regulations require the publication of auction outcomes on the PAs notice board. Moreover, an active market exists for firms reselling information on auctions. Coviello and Mariniello (2011) study the effects of these sources of information on auction outcomes.

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 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 2, 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 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). Many 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 data. 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 third 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 Validation data reveal that entry by closely connected firms is 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 data, we know that in both datasets it is extremely common to find several closely connected firms entering the same auction. 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 7 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 firms acting jointly as if they were all subsidiaries of a mother company.

C.4) Entry and Bid Regressions for the Main Data

The last piece of descriptive evidence that we present exploits the fact that in the Main data we observe both FPAs and ABAs. Therefore, we can analyze separately correlations for the probability of entry and the discount offered under the two formats. The results in Table 3 reveal an interesting difference: While for the entry regressions the sign and significance of all independent variables is the same (with the only exception of the number of workers) for both ABAs and FPAs, the opposite is true for the bid regressions. For the FPA bid regressions, our estimates conform to those in the literature: Firms further away from the location of the work offer lower discounts, while firms with a higher capital offer higher discounts. For the ABA bid regressions, instead, both variables are not significant and have the 'wrong' sign once auction fixed effects are controlled for. Overall, these bid regression complement our description of the Validation data: ABA bidding appears to be disconnected from all observable measures of firm costs.[12] Entry, instead, is associated with observable cost measures in both ABAs and FPAs. This difference will be relevant to guide our choice of how to construct the control groups for our two tests.

Overall, common features between the Main and Validation data discussed above strongly suggest they are comparable and lessons learned from our Validation data will be valuable in analyzing the Main data. We begin this analysis from a basic model of bidding in ABAs.

3 ABA Bidding and Incentives for Cooperation

This Section presents a stylized model of bidding behavior in ABAs. The model shows why ABAs have incentives for bidders to cooperate. We discuss a simple method of cooperation in which a subset of cooperators bid in order to pilot the trimmed means that determine winners in an ABA. Use of this method of cooperation creates observable patterns in firms' bids and their participation frequency. Our tests for cooperation are aimed at detecting precisely these bidding and participation patterns induced by firms engaged in mean-piloting cooperation.

Bidding in ABAs

First we focus on the case with competitive bidders and then look at how the game changes when a subset of bidders cooperates. Thus, suppose it is known that N firms submit a bid. Each firm j has a privately observed cost $c_j \in [c^l, c^h]$ for completing the job. These c_j are independent with an absolutely continuous marginal distribution $F_C(.)$. The expected profit for firm j offering b_j is: $[(1-b_j)R - c_j]\Pr(b_j\ wins)$, where R is a commonly observed reserve price (i.e., the highest price that the auctioneer is willing to pay).[13] Thus, b_j is a sealed bid between 0 and 1 representing a discount over R. The winner is determined according to the ABA rule: b_j wins if it is the highest discount strictly below $A2$.

Proposition 1: When all firms are non-cooperating, there is a unique Bayesian Nash equilibrium in which all firms bid a discount of zero percent (zero-discount equilibrium). Proof in Appendix.

The intuition for this result is straightforward: If the highest discount is offered by a single firm, this discount will surely be above $A2$ and so this firm will never win. Therefore, the

[12]Notice that the backlog variable measures the amount of unfinished work across the stock of contracts won at the time of the entry/bid decision. The positive sign appearing in column (8) is the opposite of what capacity constrained firms should show (and of what Jofre-Bonet and Pesendorfer (2003) find in their data). It is, however, compatible with collusion in ABAs: Firms that rig $A2$ upward are more likely to win often.

[13]We also assume that $R > c^h$. This implies that even the least efficient firm strictly prefers winning at the reserve price. This assumption serves only to rule out some uninteresting cases in the equilibrium analysis.

strategy profile in which all firms bid zero is an equilibrium because no firm has individually any gain from offering a positive discount. Moreover, it is the only equilibrium because the ABA rule requires the winning bid to lie strictly below $A2$ and the lowest discounts are not disqualified like the highest discounts. So, even if all discounts were identical but greater than zero, a single bidder deviating to $b = 0$ would certainly win, earning the highest possible rent. In Decarolis (2009), this zero-discount equilibrium is shown to characterize ABAs even in a more complex environment where firms are asymmetric and can default on their bids.

However, a sufficiently large group of bidders coordinating their bids can break the zero-discount equilibrium. Suppose a group of cooperating bidders consists of N^g firms that commit to submit bids that maximize the sum of the expected profits of the group members.[14]

Proposition 2: The zero-discount strategy profile is not an equilibrium unless all bidders are non-cooperating, or they all belong to the same group, or there is no group as large as the size N^* =2+(10% of N rounded to the next highest integer). Proof in Appendix.

The proposition says that a strategy profile in which all firms bid a discount of $b = 0$ is *not* an equilibrium when there is a group of cooperating firms that is 'large enough', but not as large as including all the N bidders. We define N^* as 'minimum breaking coalition' size as it is the smallest group size allowing the group to profit by breaking the zero-discounts equilibrium. Starting from a situation where all firms bid $b = 0$, a group of N^* firms can profitably deviate by submitting positive bids: For instance, the group ensures victory if it submits $N^* - 1$ discounts equal to $\varepsilon > 0$ and one equal to $\varepsilon/2$. Although the winning firm earns $R - \varepsilon/2$ instead of R, there is always an ε small enough to make this strategy more profitable than bidding $b = 0$ and winning with probability $1/N$. To get a rough measure of the strength of the incentive to cooperate, consider an auction with 51 bidders and a reserve price €312,000, which are the average values in the Main data. If the zero-discount equilibrium is played, each firm has an expected revenue of €6,000. However, if a coalition of size $N^* = 8$ forms and adopts the deviation suggested above, the expected revenues would be about €39,000 per coalition member. The ABA format provides clear incentives for firms to cooperate using this type of mean-piloting method.

It is difficult to completely characterize equilibrium behavior of bidders groups in ABAs because of the complexity of the strategy space. We present in Proposition 3 a simple example equilibrium in which a group of firms engages in an average-piloting manner. This example captures the idea of a predictable winning range presented in Section 2 and illustrates the basic method of cooperative behavior that our tests will be designed to detect. Consider a simple setting where there are N^I non-cooperating bidders and a single group of N^g cooperating bidders. Let the discounts of the non-cooperators be denoted by a vector b_f^I. The support of the elements of b_f^I is common knowledge and equal to $[b_f - \eta, b_f]$ with $b_f \in (0, 1)$ with $b_f > \eta$ and $\eta > 0$ but small. Proposition 3 shows there is an equilibrium where the group clusters its discounts in the tails of the anticipated discount distribution in order to pilot the trimmed means that determine the ABA winner.

[14] Thus, we model the group as a single player who submits N^g bids. We abstract from how cooperation is sustained between the firms in the group and make this simplifying assumption of full commitment because our data is complete on how group firms bid, but not on how they enforce their agreements.

Proposition 3: For any b_f^I and N^I and any small $\varepsilon > 0$, there exists a value N^{g**} such that if the group size is at least N^{g**}, there is a mixed strategy ε-equilibrium[15] in which all group discounts are below $b_f - \eta$. If the group size is less than N^{g**}, but at least N^* (the minimum breaking coalition), there are values of b_f^I and N^I such that for any small $\varepsilon > 0$ an ε-equilibrium exists in which all group discounts are above b_f. Proof in Appendix.

This proposition describes an example equilibrium in which non-cooperating firms bid within an interval $[b_f - \eta, b_f]$ and the group, when it is not too large, places all 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. Bid randomization is essential to avoid non-cooperating firms outguessing where the group will push $[A1, A2)$. Although pushing the interval $[A1, A2)$ downward and winning with a discount lower than b_f seems preferable, for such a strategy to be an equilibrium the group needs to have a large size, at least N^{g**}. This is because the ABA rule implies that moving A2 upward requires less bids than moving it downward (see the Appendix for details).

The average-piloting strategy in this example equilibrium is a likely suspect for a methodology followed by cooperating firms in our application. In thinking about likely methodologies for non-cooperating firms we think it is important to acknowledge that such firms are surely aware that they are likely to compete against multiple groups that are average-piloting. Moreover, if multiple groups are attempting to manipulate thresholds towards the tails of the bid distribution then we conjecture that best responses for non-cooperating firms are likely to involve randomized bidding near the center of the equilibrium bid distribution. Non-cooperating firms will have a reasonable chance to win when the lack of coordination among the competing groups results in their attempts to pilot trimmed means largely offsetting each other, preventing $A2$ from being piloted as the groups intended. Although the complexity of the strategy space prevents us from formally proving this conjecture, we believe that this basic method of non-cooperating firms randomizing near the center of the equilibrium bid distribution is a reasonable description of how non-cooperating firms should best respond.

The use of such a mixing strategy could make non-cooperating bids uncorrelated and, hence, potentially distinguishable from those of group members. At least two caveats limit the applicability of this idea. First, correlation in costs might induce correlation among non-cooperators' bids. This makes it important for our testing procedures to control for cost determinants. Second, information on how a cartel intends to rig an given auction might leak to some non-cooperators and induce them to adjust their bid up/down to match those of the cartel. To the extent such events are idiosyncratic across auctions we may be able to distinguish cooperators and non-cooperators by looking at bids across multiple auctions.

[15] In an ε-Nash equilibrium (Radner, 1980) no player has a deviation leading to gain more 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 empirical facts: (i) in ABAs, the winning discount is much less than what the most efficient firm would be willing to offer (indeed, the average winning discount more than doubled after FPAs replaced ABAs) and (ii) even an inefficient firm can profit by reselling the contract via subcontracts (indeed, the average value of subcontracts declined by a third after FPAs replaced ABAs).

Finally, a group of firms utilizing an average-piloting strategy must jointly participate in sufficient numbers for their strategy to work. In contrast, non-cooperating firms have no such incentive for joint participation. Conditional on costs of entry this incentive for joint participation should still be present for cooperators but absent for non-cooperators. This motivates our use of a comparison between participation patterns, conditional on observable entry costs, to detect cooperating groups.

4 Econometric Tests

4.1 Participation Test

Our participation test compares the participation patterns of a group of firms g comprised of firms suspected of cooperation with participation patterns in a reference (or control) set of groups that we call H. Choice of this reference set H reflects our conditioning on observable determinants of cost structures for the firms in g. For example, suppose costs can be either high or low and group g has 5 members total, 3 with high and 2 with low costs. H will consist of all groups comprised of 3 high and 2 low cost firms. Our test asks whether participation patterns in g are unusual relative to those for groups in H.

Formally, we test whether a statistic reflecting g participation pattens is a tail event relative to a reference distribution induced by randomly selecting a group from H, a uniform distribution over the groups in H. Define T as the total number of auctions and use the indicator $d_{it} = 1$ to indicate that firm i attends auction t. Then, for group g having size N^g, the fraction of auctions participated in by $K \leq N^g$ members of g is:

$$f_K^g = \frac{1}{T} \sum\nolimits_{t=1}^{T} \mathbf{1}\{K = \sum_{i \in g} d_{it}\}$$

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

$$f_K^h = \frac{1}{T} \sum\nolimits_{t=1}^{T} \mathbf{1}\{K = \sum_{i \in h} d_{it}\}$$

Our test decides whether firms in g have unusually coordinated entry by determining whether f_K^g is a tail event relative to the distribution of f_K^h induced by the random selection of group h from H. This is commonly referred to as randomization or permutation inference (see Rosenbaum 2002). At the 5 percent significance level, a one sided test of our null that g is not unusual relative to the set of groups in H corresponds to the following decision: reject if f_K^g is greater than the 95th percentile of the f_K^h distribution. The f_K^h distribution can be exactly calculated or approximated via simulation.

The choice of the set of comparable groups H will be a key decision for implementation of our participation test. For the Italian roadwork 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. 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 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.

We implement this test for a range of values for K. Participation in sufficiently large numbers is essential for average-piloting cooperation and coincidental attendance of a large group of non-cooperating firms will be unlikely. In addition small values for K are also potentially good choices since participation in small numbers would be counter to an average-piloting cooperative strategy but will coincidentally occur for non-cooperators. Thus we anticipate our test will perform best for values of K will tend to be relatively large or small.

It is important to note that in typical (non-validation) datasets, regardless of how well we use firms' characteristics to determine H, H is very likely to contain both non-cooperating firms and undetected cooperating firms. We are testing the participation patterns of a group g compared to the groups in H. We are not testing g compared to a representation of the conduct of non-cooperating firms. We cannot implement this ideal comparison of g to known non-cooperating firms since in a typical non-validation style dataset the researcher cannot know which firms are non-cooperators. This composition issue for H is unavoidable without validation data and can be an important consideration when choosing conditioning information used to construct H. We first present our main results for the Validation data and, then, we discuss how they are affected by this composition issue.

Validation Data Results: The first step to apply the test to the 8 known cartels in our Validation data is to choose H. As illustrated in Section 2, there are clear correlations between firms' entry in an auction and distance to the place of work as well as their amount of capital. In addition, firms in the dataset have a legal qualification to bid: Each firm has its own qualification, allowing it to bid for certain types of contracts, but possibly not for all of them. Therefore, we construct H as the set of all groups of firms whose composition of distance, capitalization and legal qualification match the given cartel. Matching is determined by categorizing subscribed capital and distance measured as the miles between the zip code of the work and that of the nearest establishment of the firm. We divide each characteristic into small, medium, and large categories (a third of all firms in each) and match firms based on the joint distribution of these distance and capital categories. For example, consider a cartel with 8 members who are small distance and small capital and 2 members with medium distance and medium capital. This cartel will have an H that contains all groups of 10 firms in our dataset that have the cartel's distance/capital configuration of 8 small/small and 2 medium/medium firms and that have a configuration of the legal qualifications to bid equivalent to that of the cartel.[16]

[16] In addition, we are implicitly conditioning on several factors that disciplined our dataset construction. 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.

We report the results obtained for the Turin cartels in Figure 3. 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. 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. Therefore, the evidence presented in the remaining 7 panels of Figure 3 also shows an entry behavior compatible with cooperation between cartel members. A second relevant aspect for cartels B is that small subsets, of size 2 and 3 have joint participation frequencies that are *lower* than the 5^{th} percentile of the reference distribution. The same is true for cartel C for the subset of size 2. Thus, firms in B and C exhibit behavior consistent with a cartel considering minimum breaking coalition size when coordinating entry. At end of this Section, we discuss the robustness of these results to the issue of the composition of H.

4.2 Bid Test

Our bid test is based on detecting the difference between firms cooperating via an average-piloting strategy and non-cooperators. We exploit the details of our ABA mechanism to construct a test statistic that should be sensitive to exactly the kind of average-piloting behavior that will influence winning bids.

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. Consider 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[17] 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. We compare $A1^g$ to its analogs for a set H of comparable groups. The trimmed mean without $h \in H$ is:

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

We consider how $A1^g$ compares to the distribution of $A1^h$, induced by a uniform draw from H.

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

Specifically, we compute the percentile of this distribution that corresponds to $A1^g$ and call it p^g. If H is too large to compute the percentile exactly, we approximate it via simulation.

In principle we could construct H using a conditioning approach in a manner analogous to our participation test. If the number of firms in the auction is large enough relative to the desired level of conditioning, H can be constructed so that its groups have the same composition in terms of observed characteristics as the firms in g. However, we anticipate that in many applications conditioning on cost determinants like we did for our participation test within an auction may be problematic because there might be auctions with too few participants usable as controls. Furthermore, firm characteristics may not be available to the researcher. Therefore we focus on the case where H is constructed by forming all groups of bidders that are the same size as g without conditioning on firm characteristics, within an auction. We do however employ a conditioning strategy when we examine the distribution of the percentiles, p^g, across auctions.

We now combine these percentile measures across multiple auctions. First consider a bid test across two auctions for a suspect group g that participated in both. Use the notation p_1^g for the percentile of $A1^g$ in the first auction. For the second auction, use the notation p_2^g for the analogous statistic. Our joint test statistic J^g describes the extent to which these percentiles are extreme, either small or large, across the two auctions. 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 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 and many other firms may also bid in both auctions, so the p_1^g and p_2^g statistics could have substantial dependence. In order to capture dependence across auctions in our bid test statistics we condition on participation by constructing a reference set M using only groups m that participate in auctions one and two. Our reference distribution for J^g under the null hypothesis of no cooperation is the distribution of:

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

implied by a uniform draw of m from the set M. Again, when M is too large for an exact calculation, we approximate this distribution via simulation. This joint test is trivially extended in principle to any number of auctions by redefining J^g and J^m to depend on percentiles of $A1^g$statistics from all the auctions.

Conditioning on participation patterns not only allows us to conduct inference properly with dependent p_1^g and p_2^g, it is also a control for firm costs. Firms' costs are of course an important determinant of their bidding behavior and any detection of a group being unusual in its bidding could be rationalized by its unobserved costs being unusual. The same costs relevant for bidding likely determine firms entry decisions. Thus, a firm's attendance patterns are a summary statistic for its costs that is potentially very informative. When conditioning on participation patterns, we are effectively controlling for many of the costs that drive the choice of bids. To better understand how well this conditioning works, we experiment with an alternative definition of M that not only conditions upon participation patterns but also explicitly constructs groups conditioning on the legal qualifications to bid and the capital and distance-to-job combinations as we use for our participation test. Previewing results, we find that just using participation patterns works well compared to the more data intensive conditioning on participation, capital, and distance-to-job. This is very important since many applications may not have good data on costs.

It is important to note that the number of auctions used in our multi-auction test will impact its properties due to the same undetected cooperating vs. non-cooperating firms composition issue we mentioned with our participation test. The set M will always contain non-cooperating and cooperating firms. As we increase the number of auctions jointly attended, there will be a change in the composition of the groups in M and hence the distribution of J^m. The proportion of undetected cooperating firms relative to non-cooperators grows as attendance is required at an increasing number of auctions. For example, if we conditioned upon all the members of a group attending dozens of auctions, the large majority of firms left would be those in undetected cooperating groups. Our statistical test would often (correctly) indicate that a collusive group g was not unusual relative to groups in M but this would not be an indicator of a lack of cooperation. Thus, as the number of auctions jointly considered increases, there is a cost in terms of cooperation detection performance eventually decreasing due to this composition effect.

Single-auction illustration To illustrate the usefulness of our $A1$ statistic to detect unusual bidding behavior, the distributions of p^g values for groups whose members are in cartels B and D are illustrated in Figure 4. Consider first cartel B. The histogram describes the percentile of the reference distribution of $A1^h$ to which $A1^g$ corresponds for all the auctions in the Validation data where at least 3 members of cartel B were present. The test group g for each auction is comprised of all the cartel B firms in attendance. In each auction, H consists of all groups with the same size as g. 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. Thus, for cartel B it appears that, in the large majority of the ABAs participated, its behavior is consistent with upward average-piloting. In contrast, the histogram for cartel D suggests that this cartel bids in a way that is not dissimilar from that of its control groups. Thus, we do not have indications that cartel D manipulates $A1$. For the remaining six cartels, the histograms suggest a rather clear tendency of either upward or downward (or both) manipulation of $A1$. As discussed below, cartel D was the only cartel that the court sanctioned less heavily because it judged that it only rarely colluded. Hence, these results seem supportive of the possibility of detecting cartels from their joint bidding behavior. We now discuss our multi-auction bid test which, by accounting for common cost determinants, should avoid the detection of too many false positives due to firms facing common cost shocks.

Multi-auction bid test results The first choice needed to implement this test is the group g. We use a group that is a strict subset of the cartel since too large a group will result in too few non-cooperating firms jointly attending auctions. We choose the size of g to be four or the most frequent size in which cartel members participate in an auction (see Figure 3), whichever is greater. When there are multiple groups of this size, we choose g to be the one with the highest frequency of joint participation. As described in the note to Table 4, for most of the cartels we end up using groups of five firms.

Table 4 reports the results of our multi-auction bid test for two sets of conditioning information applied to sets of two, four, six, and eight auctions. The columns labeled 'Firm Controls' report results using tests which condition upon legal qualifications, firm distance to job, and capital, while the columns labelled 'No Controls' report results where the only conditioning occurs through participation in the same auctions. For each cartel, there are two rows of entries. The first reports the median p-value over sets of two to eight auctions and the number below reports a count of these sets. We require that the auctions in these sets have at least thirty participants in common. The sets of auctions reported are randomly chosen from among all the potential combinations of auctions. The number of selected sets was chosen by imposing a time limit of one month for the Matlab routine searching for the elements to be included in the set, or 1,000 elements, whichever was reached first. For example, in the first column the entries of .13 and 739 indicate that among 739 sets of pairs of auctions attended by the group from cartel B, the median p-value of our test was .13.

It is important to note that the results in the columns labeled 'No Controls' and 'Firm Controls' are similar to each other. This suggests that conditioning on participation in the same auctions can account for relevant determinants of firms behavior. It also implies that our multi-auction bid test can be applied even in the absence of data of firm characteristics.

Since the p-values reported upon in Table 4 are not from independent sets of auctions, their distribution needs to be considered along with prior information/assumptions about the strength of dependence. Our strong prior beliefs are that this dependence is weak enough for substantial fractions of small p-values to be taken as evidence against the no-cooperation null. Therefore, our conclusion from the findings reported in Table 4 is that our multi-auction bid test is successful at detecting 6 of the 8 cartels. In fact, when considering sets of up to 8 auctions, the median p-value reaches a value below .10 (for the 'Firm Controls' case) and below .05 (for the 'No Controls' case) for all cartels with the only exceptions of D and G. It is interesting to explore more in details these latter two cartels.

The results for cartel D are not surprising because also the court had established that firms in this cartel only sporadically coordinated their actions. As regards cartel G, failing to detect it is a bit surprising as the individual auction test provided clear 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, while 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. What allows this cartel to evade detection in the multi-auction bid test is that the designated winners are the only groups that frequently participate together, while 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 of our bid test: its performance can be sensitive to the choice of group g.

Finally, it is interesting to discuss the robustness of both the participation and the multi-auction bid tests to the use of control groups that contain both cooperators and non-cooperators. Both tests should be less capable of detecting cooperation when multiple groups of cooperating firms are active. To evaluate this phenomenon, we used the Validation data to repeat all the previous tests, but with the difference that only firms indicated by the court as non-cooperating were included in the control groups. To summarize the results, which are fully documented in a Web Appendix, we do find an improvement in the detection capability of both tests. However, the results are qualitatively not different from those reported in this Section. In the same Web Appendix, we also document a series of experiments conducted to assess the robustness of our findings to the presence of correlation in firms entry/bid driven by common observable characteristics. The results broadly support the idea that our tests capture a coordination in behavior that is not driven merely by common firm characteristics.

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 for situations like that in our Main data with hundreds of bidders. Feasible strategies for selecting groups of firms will of course depend on the available information. In this Section, we describe an approach that is 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 data is comparable to the Validation data allows us to use this estimated model to predict links and groups in the Main data. We examine both the 'in-sample' performance of this method using the Validation data itself as the target, as well as its performance using our Main data. We make no claim that this group selection method is optimal, leaving the question of optimal group selection for future research. Our group selection method has three steps:

Step 1: In both our Validation and Main data, we observe measures of firms' association along three dimensions: common ownership and management, formation of temporary bidding consortia and exchange of subcontracts.[18] Using the Validation data, 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

[18]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% level a null of equal distributions.

and zero otherwise. Table 5 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 5 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 jointly in a consortium. In our favorite specification, model (2), we also use interactions between the links to improve the model predictive capacity.

Step 2: We use our estimates from the cartel membership probit model (Step 1) to generate predicted cartel membership probabilities for pairs of firms from the Main data. We will refer to these predictions as predicted cooperative group membership probabilities. To form a set of firm pairs, we begin by selecting the top 10% of firms in terms of participation, a set of likely suspects for group leaders. Each one of these firms is paired with the other firms in Main sample with which they have at least one linkage due to common ownership and management, formation of temporary bidding consortia, or exchange of subcontracts. For each of these pairs, we construct a predicted probability of cooperative group membership using the estimates of model (2) of Table 5. 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 this procedure entails arbitrarily chosen tolerance parameters, we provide its exact details in the Web Appendix in the note to Table A.5. We experimented with different parameters and settled with those reported in the note to Table A.5.

The 'in-sample' performance of this group selection method with our Validation data is reported in Table 6. 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 reports a letter from A to H that identifies the cartel most often represented in the cluster. The following column reports the size of this set of cartel members. The following two columns report the number of members from different cartels and the number of non-cooperating 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 detection of unusual cooperation. Our participation tests use the largest jointly participating set of firms within each cluster for g and detection coincides with the frequency of joint participation for g being above the 95th percentile of the reference distribution. For our multi-auction bid test we treat each cluster exactly as we treated cartels in the Validation data. The size of g is four or the most frequent size in which cluster members participate in an auction, whichever is greater. When there are multiple groups of this size, we choose g to be the one with the highest frequency of joint participation. We compute two sided p-values for all combinations of 2 auctions in which g participated and compute the median p-value across these auctions. We then repeat this for sets of 4, 6, and 8 auctions in which g participates. A detection of unusual cooperation is recorded for a cluster if any of these median p-values is .05 or less.

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, as discussed above, this is the cartel whose members cooperated only sporadically. Although several non-cooperating firms are assigned to clusters, nevertheless 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 members of the same cartel 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. Table 6 also shows the limits of the procedure: Our 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 different method for constructing candidate groups with such limited information. We examine the performance of a method that forms groups based on participation patterns and then applies only our bid test to analyze cooperation. However, when we apply this methodology 'in-sample' to Validation data, this method works poorly generating groups that allow us to detect cartel B only. Therefore, we leave a more in depth discussion of this method in the Web Appendix and, instead, proceed in Section 6 with our 3-step method for the good data scenario.

6 Search for Cooperating Groups in Main Data

This Section illustrates our methods by applying them to study our Main data. We begin by applying our 3-step group selection method and both cooperation tests to the ABAs in our Main data. We then use the results of our tests to identify a set of unusually cooperative 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 Section with a brief discussion using our benchmark cooperators to better understand the striking drop in participation when ABAs are replaced by FPAs.

Group selection begins with a list of 400 potential 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 it 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 are 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 apply our tests to these 49 clusters producing the outcomes reported in the top panel of Table 7. The table provides details about those clusters for which at least one of our tests suggests cooperation/coordination. We replicate the exercise detailed in Section 4 and illustrated in Figure 3 and and we find that the typical patterns are similar to those in this figure. We label a cluster as being unusually coordinated in entry if the test statistic for its largest number of jointly participating firms is above the 95th percentile of the reference distribution. This results in 42 clusters being classified as unusually cooperating with an average size of 10 firms each. This is indicated in the first row of Table 7. In total, there are 408 firms in these 42 clusters and their average number of bids, victories, and revenues are 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 1. Along all these dimensions, the average firm in the 42 clusters appears orders of magnitude larger that the average firm in the whole sample.

The second row of Table 7 reports results for bid tests. For each of our 49 clusters we conduct a multi-auction bid tests by treating the cluster in the same manner as we treated cartels in the Validation data with the group g selected based on joint participation as detailed above. We conducted one and two-sided multi-auction tests for all sets of 2,4,6, and 8 auctions in which the group g participated. Table A.4 in the Web Appendix reports the median, 10^{th} and 90^{th} percentile of the resulting distributions of p-values. Four clusters show clear indications of cooperation having a median two sided p-value for at least one auction-set size that are less than .05. A fifth cluster shows some evidence of cooperation having a median p-value for the two sided test of .11 and one sided test of .05. We label these five clusters as being detected to have unusual cooperators according to our bid test. They are a subset of the 42 cluster detected as unusual by our participation test.

Given these definitions of 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.[19]

A basic measure of the volume of auctions impacted by cooperation is the share of auctions receiving bids from at lest 3 members from at least one of the clusters of cooperating firms. When using the 42 clusters detected by the participation test, this definition implies that 79% of the 802 ABAs in the Main data are afflicted. This share is 43% when using the 5 clusters detected by the bid test (and also by the participation test). Using a more

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

conservative measures that requires bids from at least 5 members instead of 3, the share of afflicted auctions becomes 64% and 34% using the clusters detected, respectively, by the participation and bid tests.[20]

6.1 Potential Effect of Cooperation on Revenues

The set of unusually cooperative clusters detected by our tests captures a significant share of the revenues in this market. For instance, considering the 5 clusters detected by the bid test, their members win 333 out of 802 ABAs, corresponding to a cumulative reserve price of €143 million out of a total of €370 million. 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.[21] 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. Thus, in the Main data, at an average reserve price of €312,000, the average winning bid of 13.4% implies that the PA savings due to firm cooperation is €42,000 per auction.

The activity of groups surely results in both winners and losers. 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 up the winning discount this obviously reduces the payout when 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. A typical auction in our main data has about 51 bidders, 17 of whom are members of our detected cooperating groups.[22] Consider a hypothetical auction with 34 non-cooperating 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

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

[21] 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.

[22] The average entry in the 802 ABAs 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.

and non-cooperating firms all have the same, lower probability of winning. In this scenario, the win probability of the 34 non-cooperating firms when there is cooperation among the colluders drops to (1 - 333/802)/34 = 1.70%. Thus in this example, there is a 13.2% decrease in the win probability for non-cooperators 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 win probabilities of non-cooperating firms appears to be as important as its effect on discounting a winning payoff in impacting expected revenues. However, one good reason such a simple calculation may not be a good counterfactual scenario is that it fixes 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 non-cooperating firms, but this is necessarily beyond the scope of the current paper.

6.2 Drop in Participation with FPA Introduction

In a window of time between 2006 and 2011, the introduction of new regulations deriving from the European Union forced Italian PAs to replace ABAs with FPAs. This switch from ABAs to FPAs was accompanied by a drastic drop in participation. In our Main data, this drop can be seen by comparing the statistics in Table 1. However, Figure 5 offers an even clearer image of this phenomenon. In Figure 5, the black triangles mark the ABAs and the hollow grey circles 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 FPAs. The systematically lower values of the circles (FPAs) relative to the triangles (ABAs) is evident.[23]

There are two main causes for the drop in participation with the introduction of FPAs: The exit of inefficient firms that have too little chance of winning FPAs and the disappearance of shills who are useless in FPAs.[24] From a policy perspective, distinguishing between the two reasons might be a major concern because the regulator might want to foster participation of some less efficient firms, but most likely not of shills. In the Main data, about 4,000 firms bid at least once in ABAs and only about 1,000 bid one or more times in an FPA. However, not all of the 4000 firms were necessarily potential FPA participants. Focusing on firms that were qualified and near to prospective FPAs, we examine 1482 firms who attended at least 3 ABAs in counties where subsequently at least three FPAs for which these firms were legally qualified to bid were held. The 1482 firms contain 298 members of our 42 cooperating groups and 1184 non-group firms.[25] Of the 298 cooperators about half (159) do not participate in

[23]The bottom panel of Figure 5 documents that a similar drop in participation occurred also with the switch from ABAs to FPAs of Turin in 2003. When the collusion case that we discussed became public, both Turin abandoned the ABA in favor of the FPA. Given that our Validation data contains only ABAs and our Main data starts in 2005, the figure is based on data from the Italian Authority for Public Contracts (APC). APC data cannot be used to conduct our tests because they do not contain information on losing bidders.

[24]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.

[25]The 42 groups in the top row of Table 7 are used to classify group firms. Qualitatively, the results do

an FPA and likewise about half of the 1184 non-cooperators also do not participate in an FPA. Referring to those not participating in FPAs as exiters, the frequency of exiters does not depend on cooperating status.

Characteristics for these firms are reported in Table 8. We anticipate that shill firms will be predominately located in our detected cooperating clusters rather than among our noncooperating firms (with a perfect measure of cooperation, shills would *only* be present among cooperators). 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. We find clear differences in the characteristics of exiters according to whether they are labeled cooperators or not. Among non-cooperators, exiters have smaller capital and labor force relative to those who participate in FPAs despite being slightly older firms, possibly signaling their relative inefficiency. Exiters among cooperators also have less capital and workers than FPA participants but these gaps are smaller than for non-cooperators.

An important caveat to the interpretation of the ownership and management characteristics reported in Table 8 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.[26] 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.

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 data 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 how cooperating firms can coordinate their bids to pilot the thresholds that determine the awarding of the contract. 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.

not change if one of the two more stringent classifications is used.

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

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 manipulate thresholds exist. Similar types of manipulable mechanisms are fairly common in numerous relevant markets ranging from the procurement of public works to financial markets (the LIBOR being the most striking case), health care markets (like the subsidies awarded to insurers in Medicare D) and even labor markets.

Importantly, our results also indicate that it is not obvious that bidder cooperation should always be sanctioned. Indeed, we present the case of a 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.

8 Bibliography

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

Abrantes-Metz, R., M. Kraten, A. D. Metz, G. Seow, (2012). "LIBOR Manipulation?" *Journal of Banking and Finance*, 36, 136-150.

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.

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.

Bandiera, O., I. Barankay, I. Rasul, (2005). "Social Preferences and the Response to Incentives: Evidence From Personnel Data," *Quarterly Journal of Economics*, 120, 917-62.

Bandiera, O., I. Barankay, I. Rasul, (2006). "The Evolution of Cooperative Norms: Evidence From a Natural Field Experiment," *Berkeley Electronic Journals in Economic Policy and Analysis*, 6, 1-28.

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., (2009). "When the Highest Bidder Loses the Auction: Theory and Evidence from Public Procurement," University of Chicago, Department of Economics, PhD Thesis.

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.

Duggan, M. and F. Scott Morton, (2006). "The Distortionary Effects of Government Procurement: Evidence for Medicaid Prescription Drug Purchasing," *Quarterly Journal of Economics*, 121, 1, 1-31.

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.

Ishii, R., (2009). "Favor exchange in collusion: Empirical study of repeated procurement auctions in Japan,"*International Journal of Industrial Organization*, 27, 2, 137-144.

Jofre-Bonet, M., M. Pesendorfer, (2003). "Estimation of a Dynamic Auction Game,"*Econometrica*, 2003, 71(5), 1443-1489.

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

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.

Scott Morton, F., (1997). "The Strategic Response by Pharmaceutical Firms to the Medicaid Most-Favored-Customer Rules," *Rand Journal of Economics*, 28, 2, 269-290.

9 Appendix: Proofs

Proof of Proposition 1: Decarolis (2009) proved the result within a model allowing for ex post bidder default, but considered only symmetric pure strategy equilibria. In our model, uniqueness mechanically follows from the ABA rule because, for any strategy profile, any bid above zero that is the maximum in the bid profile is necessarily strictly dominated.

Proof of Proposition 2: If the coalition is all inclusive, $N^g = N$, the winning discount equals zero in all equilibria. All bids equal to zero is one these equilibria. For coalitions that are not all inclusive, define the "minimum breaking coalition" size as $N^* = 2 + N'$.[27] Any group of size at least N^*, but less than N, has profitable deviations relative to the strategy profile in which all discounts equal zero. In fact, a group of size N^g can, for instance, 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 profile since, in this case, any bid above zero that it places loses for sure.

Proof of Proposition 3: Proving the first 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 $N^g = 9N^I$. 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 an 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 non-cooperating bidders 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 a non-cooperating bidder 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 an ε-equilibrium. This argument implies that N^{g**} always exists: it is at most equal to $9N^I$, but it is possibly smaller, depending on the game parameters. For instance, when $N = 5$, a group of size $N^g = N' = 3$ suffices to guarantee that an ε-equilibrium with downward clustering exists. The group places one bid at zero. It clusters its other two bids within a tiny interval of size δ and randomizes them within a small interval $(0, \varepsilon)$. An independent trying to win would then need to pace its bid within this 2-bibs cluster of cartel bids. However, by reducing δ the probability of this event can be made arbitrarily small.

For the second part of the proposition, 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 of N' non-cooperating firms bids $b_f - \eta$ and (ii) the remaining non-cooperating firms bid b_f. Given this bids profile, we start by looking at a group of size N^* with bids clustered above

[27]Throughout this Section we define N' as $\lceil((.10)N)\rceil$ (i.e., 10% of N rounded to the next highest integer).

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 $b_h \in (b_f + \varepsilon, 1]$. To find the conditions under which this bids profile together with b_f^I constitutes an 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 with 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 non-cooperating firms 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 non-cooperating firms 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. To simplify the exposition suppose $N' = .1N$. 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 ABA algorithm 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. But from part one of proposition 3 we also know that as the coalition size grows, eventually it 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 non-cooperating bidder can deviate and gain more than ε. An individual non-cooperating bidder 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. 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 to zero because the probability that it can outguess the group and place its bid within the narrowly clustered group bids goes to zero. If the group size needed to achieve this is smaller than N^{g**} and the two conditions above for the incentive to cluster downward ($b_{N^*}^g > A1$ and $A2 \leqslant b_f - \eta$) are not simultaneously satisfied, we have obtained an ε-equilibrium with upward clustering.

Tables and Figures

Table 1: Summary Statistics - Main Data

Panel (a): Statistics by Auction							Panel (b): 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
$\Delta W2_{nd}$	.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	Firms that ceased activity						3.4%
$\Delta W2_{nd}$	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%

Panel (a): 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. HighBid is the highest discount. WinBid is the winning discount. $\Delta W2_{nd}$ is the difference between the winning bid and the bid immediately below it. 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 the same statistics for FPAs. The HighBid is (almost) always WinBid and so is not reported.

Panel (b): 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 firm shuts down between 2005 and 2010 (Closed) and whether it is located in one of the five regions in the North where 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 2: Summary Statistics - Validation Data

Panel (a): Statistics by Cartel

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

Panel (b): Statistics by Auction

	Mean	SD	Med	Min	Max	Obs		Mean	SD	Med	Min	Max	Obs
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
$\Delta W2_{nd}$	.09	.23	.05	0	2.9	276	Res.Price	.51	.40	.46	.05	3.7l	276

Panel (c): Statistics by Firm

Non-cooperating 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	3	717	Wins	1.9	3.1	1.0	0	19	95
Reven	51.8	19.6	0	0	2319	717	Reven	822	1466	327	0	$1e^{04}$	95
Miles	237	284	101	0	1071	504	Miles	101	207	15	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	53	717	Subct	6.8	8.6	4.0	0	44	95

Panel (a): 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. 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).
Panel (b): Summary statistics by auction. The definition of the variables is that given in Table 1. The reserve price is expressed in million of euro.
Panel (c): Summary statistics by firm, distinguishing between the firms in the 8 cartels and all the remaining firms. The definition of the variables is again that given in Table 1.

Table 3: Regressions for the Probability of Entry and the Discount Offered

	Probability of Entry				Discount Offered			
	FPA	FPA	ABA	ABA	FPA	FPA	ABA	ABA
	Probit	Probit	Probit	Probit	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(Miles	-0.84***	-0.85***	-0.86***	-0.86***	-.65**	-0.30*	0.26***	0.05
Firm-Work)	(0.02)	(0.02)	(0.01)	(0.01)	(0.25)	(0.17)	(0.08)	(0.04)
Log(Firm	0.06***	0.06***	0.05***	0.05***	0.05	0.12***	0.01	-0.01
Capital)	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)	(0.01)	(0.02)	(0.01)
Backlog	0.12	0.13	0.07	0.07	0.88	-0.44	0.06	0.18*
	(0.21)	(0.21)	(0.05)	(0.05)	(.94)	(1.17)	(0.15)	(0.10)
Unlimited	0.45**	0.46**	0.04**	0.05**	-3.09*	-0.05	-0.10	-0.06
Liability	(0.19)	(0.15)	(0.02)	(0.02)	(1.65)	(1.17)	(0.16)	(0.08)
Number of	0.04	-0.02	-0.66***	-0.67***	0.22	2.34	-1.64**	-0.17
Workers	(0.55)	(0.55)	(0.17)	(0.17)	(8.67)	(7.98)	(0.74)	(0.39)
Firm Links	No	Yes	No	Yes	No	No	No	No
Auction FE	No	No	No	No	No	Yes	No	Yes
Prob. Chi^2	0.00	0.00	0.00	0.00	-	-	-	-
R^2	-	-	-	-	0.21	0.55	0.13	0.65
Observations	11,806	11,806	80,274	80,274	2,182	2,182	45,513	45,513

Significance level: * is 10%; ** is 5%; *** is 1%. Sample: Main data. Columns (1)-(4) report probit regression 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. A firm is a potential participant if it has: (i) the legal qualification to bid, (ii) submitted a bid at least once in the county where the auction is held and (iii) submitted a bid at least once in the region where the auction is held in the same year of the auction. All probit 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. Relative to model (1) and (3), models (2) and (4) include "firm link" variables: For every firm and auction, we count how many other bidders in the auction are linked to the firm along each one of the links described in Table 5 (common personnel, common owner, common manager, common zip code, common municipality, common county, subcontracts, winning consortium and bidding consortium).
Columns (5)-(8) report OLS regressions for the discount offered. Standard errors are clustered by PA and year. 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. Relative to model (5) and (7), models (6) and (8) also include auction fixed effects.

Table 4: Multi-auction Bid Test: Median P-Values

	2-auction		4-auction		6-auction		8-auction	
	No Controls	Firm Controls	No Controls	Firm Controls	No Controls	Firm Controls	No Controls	Firm Controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cartel B	0.13	0.11	0.06	0.10	0.06	0.09	0.04	0.07
	739	815	574	859	354	902	727	445
Cartel C	0.11	0.11	0.06	0.09	0.03	0.08	0.01	0.07
	531	608	399	610	311	628	278	676
Cartel G	0.22	0.36	0.23	0.16	0.15	0.19	0.18	0.11
	728	992	831	981	621	992	455	989
Cartel A	0.00	0.00	0.00	0.02	0.00	0.02	0.00	0.02
	45	45	206	206	207	207	45	45
Cartel E	0.11	0.09	0.08	0.05	0.05	0.02	0.06	0.02
	190	190	466	177	615	119	427	39
Cartel D	0.43	0.41	0.39	0.37	0.39	0.41	0.38	0.40
	160	134	482	67	280	15	127	3
Cartel F	0.05	0.07	0.02	0.03	0.02	0.02	0.01	0.01
	199	210	822	761	938	902	972	956
Cartel H	0.04	0.04	0.03	0.04	0.02	0.03	0.01	0.03
	289	300	965	965	997	997	999	999

For each one of the K-auction tests, the table reports for each cartel in the top row the *median p-value* of the two-sided multi-auction tests across all the combinations used and, in the bottom row, the actual number of combinations used. More in detail, for every cartel, we start by selecting the subgroup on which we conduct the test in the way described in the text. For cartels B, C, G, A, E, D, F and H the subgroups used have size, respectively, 5, 5, 4, 7, 5, 4, 5 and 5. The number of auctions jointly entered by all members of these subgroups are (in the same order): 184, 51, 68, 10, 20, 19, 21 and 25. Thus, for instance, for cartel C there are "51 choose 2" combinations of 2 auctions that could be used to conduct the 2-auction bid test (K=2 test, using the notation in the text). We could perform the test on each of these combinations or, when their number is too large, on a random subgroup of them. We do the latter, but also require that the auctions considered have at least 30 firms in common, so that enough other firms could be used to form the control groups. This implies that we have an entire distribution of results and, hence, we report in the table the *median p-value* of the two-sided multi-auction tests across all the combinations used (and, in the bottom row, the number of combinations used). We interpret low values of the median p-values as a rejection of the null of no cooperation. Even numbered columns report the results when using control groups that, like the ones used for the participation test, match the suspect cartel in terms of legal qualifications to bid, capital and distance to the place of the work. Odd numbered columns, instead, report results without conditioning on firm observable characteristics. In the Web Appendix, Table A.1 and A.2 report the 10^{th}, 50^{th} and 90^{th} percentile of the result distributions, as well as the results of the one-sided left test.

Table 5: Probit Regression - Validation Data

Probability that for a pair of firms both firms belong to 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%. The dataset consists of all pairs of firms (from the Validation data) that share at least one owner (manager) or exchanged subcontracts or bid at least once as a legal temporary bidding consortium. The table presents probit coefficients and, in parenthesis, their standard errors corrected following Conley (1999) for the correlation across any pairs that share firms. The dependent variable equals one if the pair belongs to the same cartel and zero otherwise. All independent variables are all dummy variables. The first three variables listed in Table 5 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. Model (2) differs only in that it includes interactions.

Table 6: Clusters in the Validation Data

3-Step Method						
Assigned Group	Known Cartel	Members Cartel	Members Other Cartels	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

The table shows the clusters 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 cluster. The second column reports the identifier of the cartel to which most of the firms in the cluster are affiliated. The third column reports the number of firms belonging to the cartel 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 (median p-value of the multi-auction) bid test (Bid), through both of them (Both) or whether no detection occurs (No). All tests are at the 5% level.

Table 7: Detection Results in the Main Data

Clusters Detected as Groups of Cooperating Firms					
Rejected Test	Number of Clusters	Cluster Size	Entry	Number of Victories	Revenues
Participation Test	42	10	45.2	0.82	350,231
Bid Test	5	16	59.0	1.08	462,914

The table reports the clusters detected in the Main data. Using the participation test at the 5% level, a rejection is found for 42 clusters. Using the (median of the p-value of the) multi-auction bid test at 5%, a rejection is found for 5 clusters. For this latter test, the whole result distributions are reported in Table A.4 in the Web Appendix. The final four columns report respectively: the average of the size of the cluster and the means (across all firms in the groups) of entry, number of victories and revenues.

Table 8: Firms' Size and Gender Composition

	Not Entering FPA			Entering FPA		
VARIABLES	Mean	SD	N	Mean	SD	N
Non-cooperating 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 Detected Clusters:						
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

The table reports statistics for 4 sets of firms: (i) non-cooperating firms that never bid in FPAs (top left), (ii) non-cooperating 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 clusters described in the top row of Table 7. A firm is in the entering-FPA group if it bids in at least one FPA. A firm is in the not-entering-FPA group if: (i) it never bids in any FPAs and (ii) it bids 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 sets, the columns Mean and SD report the average and standard deviation taken across all firms in the set. The column N reports the number of firms considered. The firm characteristics considered are: the number of years between the beginning of activity and 2010 (Firm Age) and the average value between 2006-2010 of the number of all dependent workers (Number of Workers), the fraction of female white collar workers over all white collar workers (Proportion of Women), the number of female owners (managers) (Number Female Owners (Managers)), the ratio of the number of female owners (managers) to that of the total number of owners (managers) (Proportion Female Owners (Managers)) and (expressed in €1,000) capital, revenues and profits.

Figure 1: An Illustration of the Italian ABA

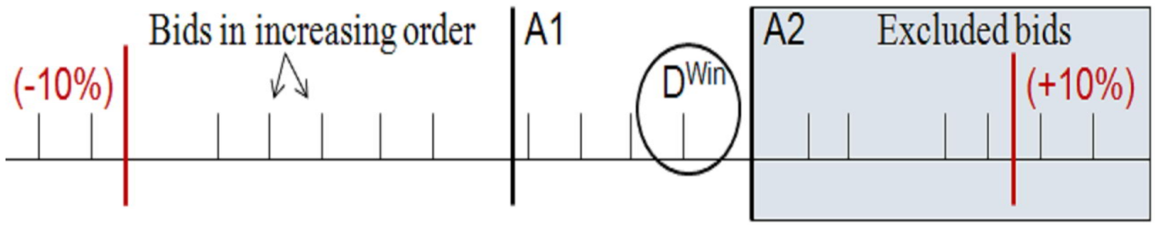

Example of an ABA with 17 bids. Bids, which are discounts over a reserve price, are represented by the 17 small vertical bars. Discounts are ordered in increasing order. The trim mean (A1) is calculated disregarding the 10 percent of the lowest and highest bids, rounding up to the next highest integer. Since there are 17 bids, this means that the 2 lowest and the 2 highest discounts are disregarded (in the figure, two thick vertical bars marked respectively '-10%' and '+10%' separate these discounts from the others). A1 is the mean of the remaining discounts. A2, instead, is the mean of all discounts strictly within A1 and the lowest of the top 10% of discounts calculated in the first step. 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 for being abnormally high.

Figure 2: Example of an ABA in the Validation Data

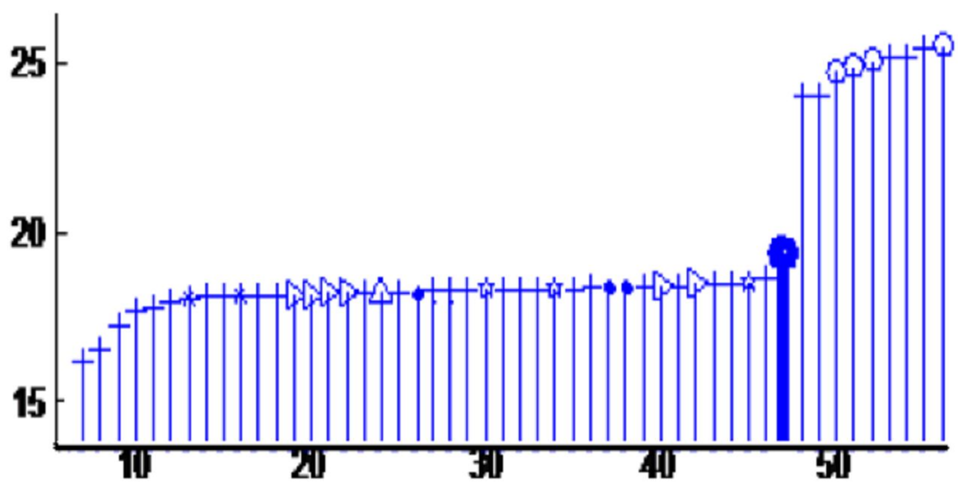

Discounts offered in one of ABAs in the Validation data. The horizontal axis marks all the 56 bids that were submitted in this auction. The bids are sorted to be in increasing order of the discount offered. The vertical axis reports the discount offered. Almost all bidders offered a discount close to 18%. The different symbols mark different cartels, but the cross indicates non-cooperating firms. The thick line marks the bid of the winner. The nine highest bids comply with the description of 'supporting bids' offered by the convicted firms and reported in the text.

Figure 3: Participation Test - Validation Data

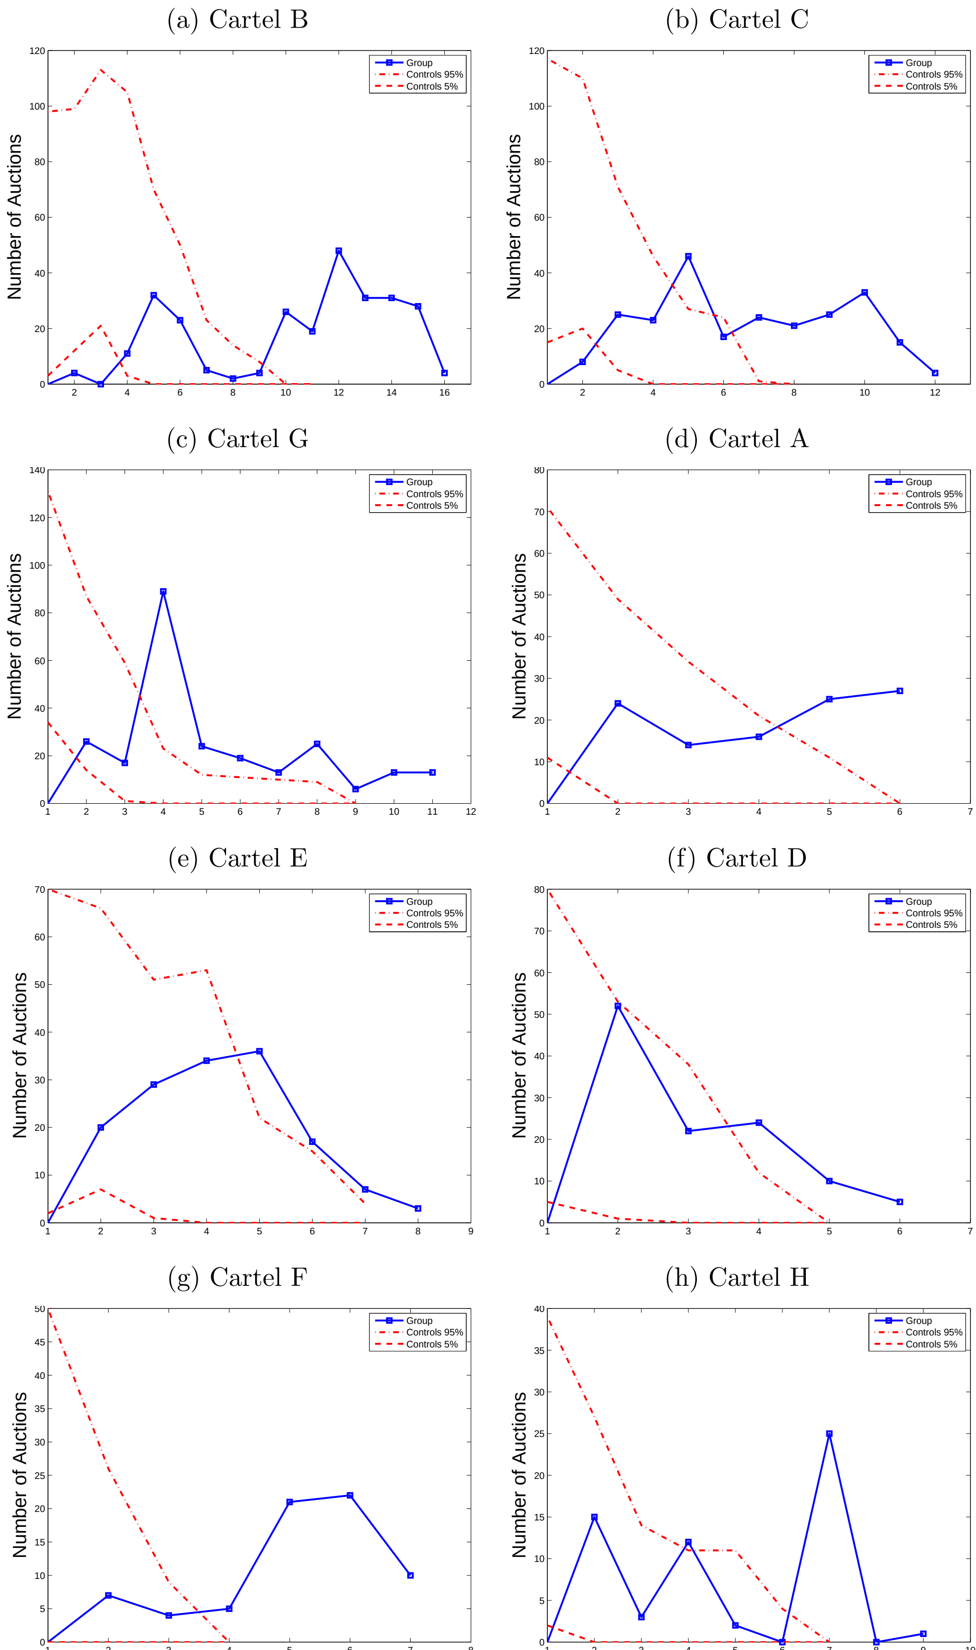

Participation test for all cartels and all of their possible subgroups.

Figure 4: Single-auction Bid Test - Validation Data

(a) Cartel B

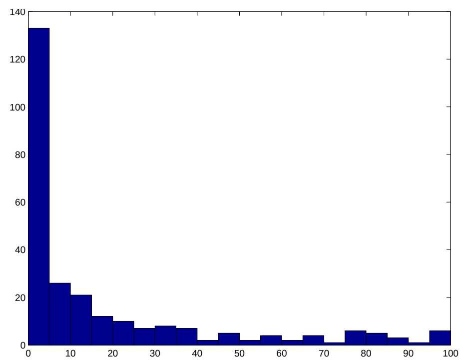

(b) Cartel D

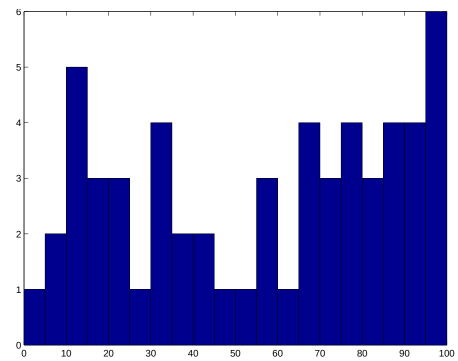

Histograms of the percentiles of the single-auction bid tests for cartels B and D.

Figure 5: Number of Bids in ABAs and FPAs

Panel (a): The 2006 Reform in Four PAs

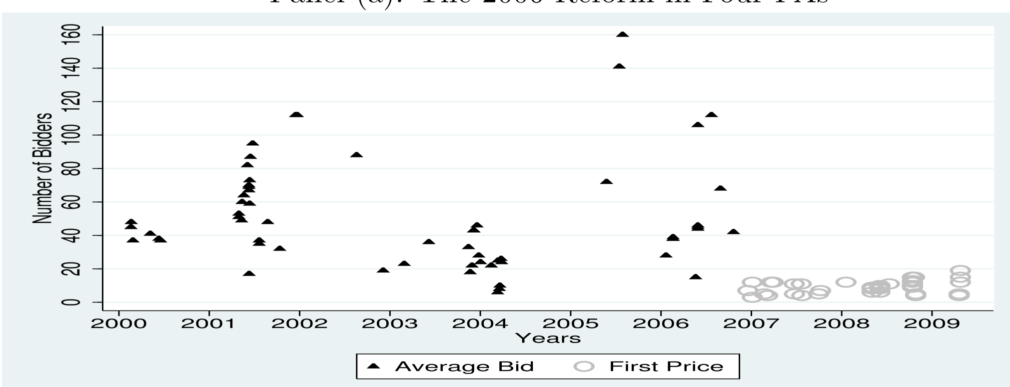

Panel (b): The 2003 Reform in Turin

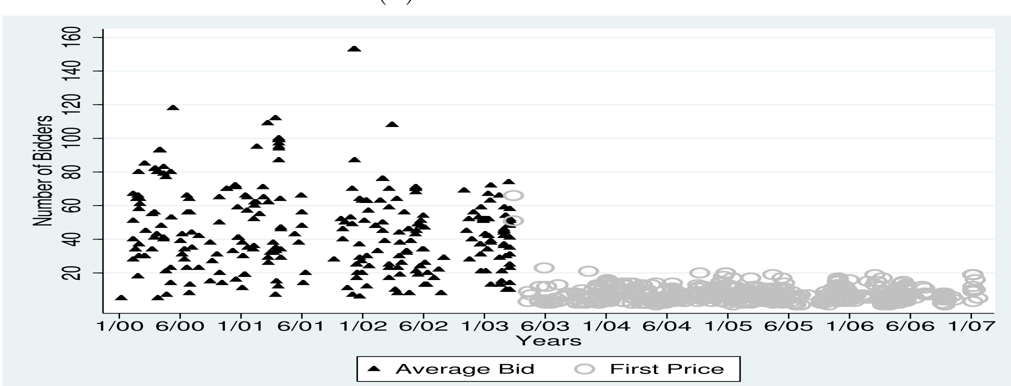

Panel (a): Scatter plot of the number of bidders in ABAs and FPAs held by four PAs in the Main data: Padova, Varese, Sondrio and Cremona, which all switched to FPAs. Panel (b): Scatter plot of the number of bidders in the ABAs and FPAs held by the municipality of Turin. Data source: Records of the Italian Authority for Public Contracts (APC). Validation data are a subset of the APC data.