Detecting Bidders' Groups in Collusive Auctions: Evidence from Average Bid Auctions

Timothy G. Conley and Francesco Decarolis*

27th April 2011

Abstract

In this paper we empirically analyze the behavior of firms participating at average bid (AB) public procurement auctions. AB auctions are characterized by the fact that the winner is the bidder submitting the bid closest to (some funtion of) the average bid. We propose two statistical tests to detect whether firms' bidding and entry decisions are indicative of coordinated behavior. We validate these tests on a subset of auctions where the presence of eight cartels active between 1998 and 2003 has been sanctioned by the judiciary. We then apply the tests to a large set of auctions for road construction works held in the North of Italy between 2005 and 2010 finding strong evidence that multiple groups of not independent bidders are present. We use these results to analyze the effects of the groups' activities on the auctions' revenues and to explain the disappearance of several hundreds of firms after the first price rule gradually replaced the AB rule after 2006.

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

PRELIMINARY AND INCOMPLETE - PLEASE DO NOT CIRCULATE

*Conley, University of Western Ontario. *E-mail address:* tconley3@uwo.edu. Decarolis, University of Wisconsin - Madison. *E-mail address:* fdc@ssc.wisc.edu (corresponding author). The authors would like to thank their colleagues at their institutions, George Deltas, Jakub Kastl, Gustavo Piga, Robert Porter and Giancarlo Spagnolo for the useful suggestions. We thankfully acknowledge the aid of the Bank of Italy, the Italian Authority for Public Contracts and the Legal Office of the municipality of Turin to collect the data.

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

1 Introduction

In Italy, since 1999 most public works are procured through auctions in which the winner is the bidder submitting the bid closest to (a function of) the average of the bids. In particular, sealed bids are submitted in the form of a discount over an announced reserve price and the discount that is the closest from below to a trimmed average of the bids wins. The winner is paid the price he offered to execute the work. In 2008, contracts worth in total about €6 billion were procured using this format. The origin of this type of auction is uncertain, but the civil engineering literature rediscovered it in the '90s and suggested that forms of the Average Bid (AB) auction could outperform the First Price (FP) auction for the procurement of contracts. Indeed, recent studies have shown that an auctioneer might prefer using an AB over an FP auction when there is the risk of a costly default.[1] Nevertheless, these studies show that the AB auction achieves this result at the cost of breaking the link between bids and costs, essentially transforming the auction into a lottery. Furthermore, since the winning price in this auction/lottery depends on the average of the bids, this rule gives strong incentives to bidders to pilot this average by coordinating their bids.[2] This paper seeks to study the behavior of firms bidding in the Italian AB auctions. In particular, it introduces two statistical tests aimed at identifying the presence of groups coordinating their entry and bid decisions to influence the awarding the contract. The application of the tests to a large dataset of auctions for road construction held between 2005 and 2010 reveals the presence of multiple groups and provides a quantification of their effect on the auctions' revenues and on bidders' participation.

Although various forms of AB auctions are used in many countries, little is known about how they work.[3] An interesting feature of most AB rules is that they ensure that the highest discount loses whenever there are no ties of the highest discount. Generally, this implies that all these formats share one equilibrium in which all firms submit the lowest possible discount. In this equilibrium the auctioneer pays the highest price for the execution of the contract, exactly

[1] The original engineering studies, Ioannou and Leu (1993) and Liu and Lai (2000), consider non-strategic bidders. Decarolis (2010) obtains this result in a strategic model with firms that have privately observed costs and asymmetric default types. Burguet, Ganuza and Hauk (2009) and Chillemi and Mezzetti (2009) characterize the optimal procurement mechanism under default risk and show that it has features similar to an AB auction.

[2] This feature of AB auctions has been pointed out first by Albano, Bianchi and Spagnolo (2006).

[3] Various forms of AB auctions have been used to procure public contracts in Chile, China, Italy, Japan, Peru, Taiwan and, in the USA, in Florida and New York, see Decarolis (2009) for a more detailed description.

like when all bidders form one single cartel. Because of this feature, we refer to AB auctions as "collusive auctions". Nevertheless, an equilibrium in which all firms offer the lowest bid is not robust: a subgroup of firms large enough to pilot the awarding threshold can gain by rising their discounts and winning the contract at a slightly higher discount. Although this subgroup of firms rigs the auction, it is not properly a cartel because its behavior benefits the auctioneer.[4] Indeed, the presence of one or more subgroups of bidders is essentially the only form of competition allowed by AB auctions. In this paper we focus on the Italian AB auctions and we try to develop an empirical methodology that could detect the differences in behavior between independent and coordinated firms that attend these auctions.

Our analysis of firms' behavior starts by formulating a model of firms' entry and bidding decisions. We argue that, if firms act independently, the unique equilibrium that should result is not compatible with what we observe in the data. Therefore, we extend the model allowing for the presence of coordinated actions by firms in the same groups. The model indicates that, compared to independent firms, firms that belong to the same group will be more likely to enter together and to bid on the same side of the bid distribution. In fact, clustering bids on the same side of the distribution serves both to shift the average in an unpredictable way and to have the group's bids in the area toward which the average is shifted. Moreover, only a firm within a group would find optimal to bid at the extremes of the support of the bids distribution since such bid will always lose. Finally, joint entry of the group members is essential to ensure that the group has enough bids to influence the average.

These features motivate the two statistical tests, one for bids and one for entry, that this paper develops. These tests seek to identify a group by comparing its behavior to that of comparable sets of randomly grouped firms. In particular, relative to a random group, the test on bids evaluates how much the bids of a suspect group affect the average. We show how having access to a sample of auctions in which the same firms repeatedly participate enhances the development of a multi-auction bid test. The test on participation, instead, essentially compares how likely are firms in a suspect group to enter together relative to similar sets of independent firms in the market. Although the ideas behind the tests are intuitive and have some theoretical foundations, the environment in which the firms operate is quite complex and the capacity of these tests to identify bidders' groups could be a concern.

The way we address this issue consists in evaluating the tests on a "validation" dataset composed by auctions participated by firms with known affiliations to groups. In particular, we use 276 AB auctions for roadworks held by the city of Turin between 1999 and 2002. In 2008, the Turin's Court of Justice ruled that these auctions had been rigged by 8 groups constituted by approximately 95 firms. Each group strategically submitted bids to affect the awarding of the contract. According to the Italian law this activity is a crime, and hence these groups were labelled cartels[5] and their members were fined and, some of them, sentenced to jail. For our

[4] The purpose of a cartel should be to transfer revenues from the auctioneer to cartel's members.

[5] We refer to these 8 groups as cartels. We use the word "groups" for the sets of not condemned firms.

purposes, this is an ideal sample to validate our tests because we can check whether the tests are able to identify the 8 cartels sanctioned by the Court. The results that we obtain are very much supportive about the capacity of our tests to correctly identify groups. Of the 8 groups, the only one for which we do not find systematic evidence of coordination is the one that the Court sanctioned less because its members rarely coordinated bids.

We then turn to the problem of identifying groups when group membership is unknown. In principle, any subset of firms could be taken and analyzed through the test to check whether its members coordinated entry and bid. Nevertheless, given the large number of firms in the market, this approach is not practical. Therefore, we propose two solutions that differ in the amount of information required. The first method entails constructing candidate groups through a clustering algorithm that links firms on the basis of some observable characteristics. We illustrate this approach using information on subcontracts, legal joint bids and ownership. The second method, instead, requires observing only bids and bidders identities. In this case the candidate groups are chosen on the basis of the joint participation with those firms that most frequently win auctions. The groups formed with the first method can be tested with both tests while the ones formed with the latter can be analyzed only with the bid test.

We apply our methods to a dataset of approximately 800 AB auctions held in the North of Italy between 2005 and 2010. In these data, bids are discounts over the reserve price and the reserve price is set homogeneously across public administrations (PA). In each auction, a large number of bidders (57 on average, with peaks of more than 300) bid to win a contract to perform a simple roadwork (like paving a road). Although we have no prior knowledge of groups in these auctions, many of their observed features resemble those of the validation sample. Indeed, the results of the tests strongly suggest that numerous groups are active in these auctions. For instance, when we apply the tests to the auctions in the Piedmont region, we can classify between 20 and 30% of the firms in the market as members of groups. Accordingly, we can argue that between 20 and 50% of the auctions are participated by groups.[6] We then turn to the problem of evaluating the effects of the activity of these groups. As regards the revenues of the auctioneer, we argue that all groups benefit the auctioneer.[7] However, they harm the independent firms both because they reduce their probability of winning and because they prevent the winning discount to be equal to zero. We present some basic estimates of the damage for firms outside the group and discuss how selection could bias these estimates.

The final set of results that we present concerns the large exit of firms that followed the introduction of the FP auction. In 2006, a reform of the procurement regulation required by the European Union gradually induced the substitution of AB with FP auctions. One of the most visible effects of the switch from AB to FP auctions was the drop in participation. Indeed, for the local administrations that changed the format, the average number of firms per auction decreased from 57 to 7, the 99th percentile decreased from 300 to 20. Our results on bidders'

[6] The variation in the number of auctions classified as colluded depends on several choices (see section 6).

[7] The result hinges on the counterfactual used being the zero-bid equilibrium (see section 3).

groups help to address what part of this market shakeout is due to the exit of inefficient firms and what part is due to the disappearance of shill bidders. A shill is a replica of another firm which is created only to have one more bid to affect the auction outcome. Since our tests allow to identify which firms belong to groups, we can argue that exiting firms that are not part of a group are most likely inefficient firms. Instead, for exiting firms that belong to a group it is not possible to distinguish between which are shills and which are inefficient members of the group. Applying this idea to the auctions in the North of Italy, we conclude that about 700 firms that abandon the market do so because they are inefficient.

Therefore, the contribution of this paper is twofold. On the one hand, our tests can be used to identify bidders' groups in AB auctions. Therefore, the tests could be useful to a Court investigating a case or to an inspector choosing which firms to monitor. However, on the other hand our analysis suggests that the activity of these groups is not harmful for the auctioneer's revenues. This result would be useful for a Court evaluating the damages imposed by a group. More generally, we believe that the novel results presented by this paper might help the policy discussion about the opportunity of adopting or continuing using AB auctions which are not limited to the procurement of works but, for instance, characterize also part of the procurement process of Medicare.[8]

Literature: By studying bidders' cooperation, this paper is most closely related to the literature on collusion in auctions. Collusion is generally regarded as a first order concern auction design (Klemperer, 2004). The seminal studies in the theoretical literature include Robinson (1985) addressing the strength to collusion of first price relative to second price auctions and the studies on cartels behavior in second price or English auctions (Graham and Marshall, 1987, and Mailath and Zemski, 1991) and on first price auctions (McAfee and McMillan, 1992). More recent work on auction design in the face of collusion is Marshall and Marx (2006). The theoretical analysis of average bid auctions has been very limited in the economics literature, but the studies analyzing this format have pointed out its susceptibility to bidders' collusion (Albano, Bianchi and Spagnolo, 2006, Engel, Guanza, Hauk and Wambach, 2006). In this paper, we qualify the results of these latter studies for the case of the Italian AB auction and show a rather unusual result: bidders' coordination in a less than all inclusive coalition is beneficial for the auctioneer.

The main strand of the literature to which this paper is related is that concerned with the empirical analysis of collusion in auctions. This literature can be roughly divided into two groups: the studies of collusion practices in markets where the presence of cartels existence has been proved by a court (Asker, 2009, Pesendorfer, 2000, Porter and Zona, 1993 and 1999) and the studies that try to devise methods to distinguish competition from collusion in environments where the presence of collusion is only a possibility (Bajari and Ye, 2003). Both approaches have

[8] Katzman and McGeary (2008) document the use by Medicare of a multiunit median price auction for the procurement of durable medical equipment. The main difference between this rule and the one that we study consists in its multiunit nature and not in the use of the medain instead of the average.

lead to the flourishing of a literature within industrial organization concerned with "screens for collusion" (i.e., statistical tests to detect collusion, see the review by Abrantes-Metz and Bajari, 2010). In this paper, we take an intermediate approach: we use information from auctions where collusion was proved, but we do so in order to devise an empirical methodology that allows assessing the likelihood of groups in markets where their presence has not been proved yet. The motivation of our approach is based on the idea of Hendricks and Porter (1989) who explain that collusion is tailored to the specific rules of the auction and the environment. Therefore, we use data from auctions with collusion to learn about the behavior of groups and then search for evidence of this behavior in other similar auctions.

Finally, our analysis contributes in two ways to the large literature on public procurement auctions. First, it contributes to the study of collusion in public procurement auctions (a review of cases for the US is contained in Haberbush, 2000, for cases in the procurement of roadwork contracts see Porter and Zona, 1993 and Ishii, 2007). Secondly, it contributes to the study of mechanisms similar to the AB auction, which Decarolis (2010) documents being widespread in public procurement. In this respect, our study also contributes to the debate about the problematic features of this type of mechanisms initiated by Cramton, Ellermeyer and Katzman (2011) regarding the proposed use of a multiunit median bid auction for Medicare's procurement of durable medical equipment.

The outline of the paper is as follows: the next section provides a description of the market, section 2 provides a description of the market and our data sources, Section 3 presents a model of firms' entry and bidding, section 4 presents our econometric tests and investigates their performance on the sample of auctions with known groups, section 5 discusses the case of testing with no prior knowledge about groups, section 6 illustrates the results obtained by applying the tests to the auctions in the North of Italy and, finally, section 7 concludes.

2 Description of the Market

Italian public administrations (PAs) seeking to procure a contract for the execution of any kind of construction work are subject to the rules laid down in the national Code of Public Contracts (the Code). The Code allows for essentially four awarding procedures: negotiations (only for small contracts), scoring rule auctions (based on price and other criteria) and auctions based only on price. In the period between 2000 and 2008, the latter category accounted for 82 percent of all the awarded contracts in terms of value and 79 percent in terms of numerosity of the contracts (corresponding, per year, to approximately €13.5 billion and 13,000 contracts). However, of all the auctions based on price, only 40 percent in terms of value (or 3 percent in terms of numerosity) are of the familiar first price (FP) auction type. The remaining auctions are, instead, average bid (AB) auctions. The Code refers to them as "first price auctions with an automatic elimination of abnormal tenders" and specifies the following procedure to identify the winner (and, hence, also the abnormal tenders): a) the sealed bids of all firms are simultaneously opened and ranked from the lowest to the highest bid (bids are discounts over the reserve price reported in the auction notice, so a high bid is a low price for the PA); b) a trim mean (A1) is calculated excluding the 10 percent of the highest and lowest bids; c) a new mean (A2) is calculated as the average of those bids above A1 but below the disregarded top 10 percent of bids; d) the winning discount is the highest discount strictly below A2. Discounts equal or greater than A2 are defined abnormal and excluded, without the possibility of appeal. Ties of winning bids are broken with a fair lottery. The winner is paid his bid to perform the work. Figure 1 offers an example with 17 bids: the winner is denoted Dwin and, in this case, it is the 7th highest discount.

Figure 1: An Illustration of the Italian AB Rule

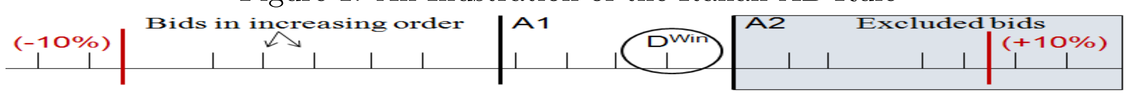

Bids are represented by the 17 small vertical bars. They are discounts and are ordered in increasing order. This figure is taken from Decarolis (2009).

In the period between 1999 and 2006 the AB format was compulsory for all contracts below (approximately) €5 million. However, in 2006 and 2008, two reforms required by the European Union induced a gradual replacement of AB with FP auctions. Nevertheless, in 2008 the AB auctions were still the mechanism used to procure the vast majority of contracts, for a total yearly value of about €6 billion. In a related study, Decarolis (2009) discusses why AB auctions were originally introduced[9] and uses the recent policy changes induced by the EU to quantify the effect of the switch from AB to FP auctions on the PAs' cost of procurement.

Although our focus in this paper is exclusively on firms' behavior in AB auctions, for purely illustrative purposes it is useful to report statistics for both AB and FP auctions for the market that we study. Within the large set of contracts procured through AB auctions, we look at simple roadwork contracts (mostly paving jobs). Moreover, we restrict attention to auctions held by provinces and municipalities in five regions of the North of Italy between November 2005 and May 2010 (more details are given in the endnote of Table 1). These choices allow us to deal with a rather homogenous set of auctions. Moreover, these simple roadwork contracts are the most frequently procured, representing 30 percent of all the contracts awarded per year. Table 1 reports summary statistics for both the sample of AB and FP auctions.

Insert Table 1

The left panel of the table reports several statistics both across and within auctions for the AB and FP samples. Although the contracts in the two samples have a similar reserve price,

[9] For a discussion of the exact evolution of the reulation see Decarolis, Giorgiantonio and Giovanniello (2010).

the bidders' behavior markedly differs in terms of both entry and bidding. As regards entry, the number of bidders is several times larger in AB than in FP auctions: on average there are 7 bidders in FP auctions and 51 in AB auctions. As regards bidding, the winning discount is on average 13 percent in the AB, while it is 30 percent in the FP auctions. Moreover, in the AB auction there is substantially less within-auction variation in the bids than in the FP auctions: this is shown by both the lower within-auction standard deviation of bids and the lower difference between the winning bid and the next highest bid in the AB relative to the FP auctions. This latter variable, sometimes defined as "money left on the table" is on average 4.5 percent of the reserve price in the FP auctions but only .2 percent in the AB auctions.

Our study focuses on the behavior of firms bidding in the AB auctions. Summary statistics for these firms are reported in the right panel of Table 1. There are approximately 4,000 firms that bid at least once and they exhibit substantial differences among many dimensions, which is not surprising given their numerosity. Strong asymmetries appear clearly in firms characteristics (like their capital) and in their performance in the auctions (like their number of victories). The bidders in FP auctions are a subset of these firms and are typically larger than the average firm bidding in AB auctions in terms of capital and workers. An interesting empirical result of Decarolis (2009) is that the variation in firms characteristics explains bids exclusively in the sample of FP auctions but not in that of AB auctions.[10] In the next section, we develop a model that aims to explain why firms' bids could be disconnected from all proxies of firms' costs. At the heart of the explanation there is the idea that the AB auction generates strong incentives to manipulate the awarding rule through coordinated entry and bids.

2.1 Validation Dataset

We conclude this section discussing a dataset of auctions where coordinated entry and bidding is known. In section 4 we will discuss the usefulness of our tests by applying them to this "validation sample". Therefore, this part of our analysis is similar to Porter and Zona (1993) who construct a test for collusion based on some observed relevant features of a known cartel.

Turin's Rigged Auctions

On April 2008 the Court of Justice of Turin convicted the owners and managers of numerous construction firms that rigged the AB auctions for roadwork held in the area of Turin between 1999 and 2002. The sentence identifies a network of about 95 firms that operated trough 8 groups, referred to as cartels.[11] This case is particularly useful for our study because it involves the same type of firms, auctions and contracts that we observe in our main dataset. The Turin's

[10] In particular, firms' covariates are not significant determinants of bids once auction (or auctioneer and contract type) fixed effects are used, see Table 8 in Decarolis (2009).

[11] Turin Court of Justice, 1st criminal section, April 28th, 2008, sentence N. 2549/06 R.G.. Of the 95 suspect firms, the sentence condemns 29. Prescription lead to acquaintance for 2 firms. The judgment of the other firms was scorporated into different court cases. In our study we consider the original full network of 95 firms.

groups 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 in the Piedmont region between 2000 and 2003. Interestingly, the groups were formed mostly on the basis of firms' geographical proximity, as Figure 2 shows. This is likely due to the lower costs of coordinating actions and of exchanging favors.[12] Moreover, two groups, despite having all members close to each other, are located far from Turin. According to the sentence, these groups did not want to win the auctions to perform the jobs, but just to resell them through subcontracts. Finally, as Table 2 shows, these 8 groups are quite different in their size, entry and victories.

Figure 2: Localization of the 8 Cartels

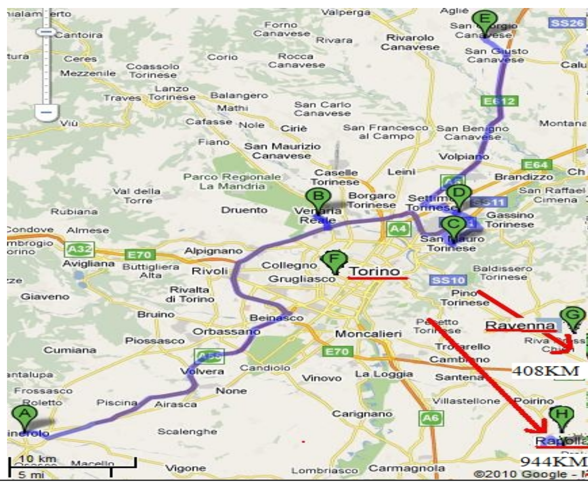

Table 2: Cartels in Turin's Auctions

Cartel ID	No.Firms	No.Win	Entry
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

Capitals letters pinpoint cartels on the map.

Notes: The map gives the exact location of the 6 cartels located around Turin. Cartel G, instead, is centered 408 Km away in the North-East of the country. Cartel H is located in the South, 944 Km away from Turin. However it frequently operates together with some firms located in the same municipality of cartel A. This is the only exception to the fact that firms of the same cartel are all located within few kilometers away from each other.

In addition to the asymmetries across groups, there are also significant asymmetries within groups. The bottom panel of Table 3 reports summary statistics for both the firms inside and outside the groups. Given that this sample was assembled to build a case against the firms accused of collusion, it is normal 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, instead, the top panel of Table 3 suggests that these auctions are rather close to those of our sample reported in Table 1 on the basis of entry and of dispersion

[12] 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 lettings and they take place quite regularly; (4) there are legal forms of joint bidding; (5) the same firms repeatedly interact, (6) ex post the auctioneer discloses the identities and bids of all bidders. These six reason likely played an important role for both the formation and the stability of the groups that we study.

of the bids. Interestingly, the average winning bid is higher in these colluded auctions than in those reported in Table 1, 17.4 compared to 13.7.

Insert Table 3

As both the confessions of some entrepreneurs and the intercepted mails and phone calls reveal, the strategic environment is complex. The cartels compete against each other (although in some occasions some of them form short term agreements) and against numerous independent firms. Nevertheless, there are at least two features of bidding behavior that emerge clearly from both the confessions and the data. The first feature is that across auctions the winning bid appears to be ranging within a narrow interval around 18. Some convicted firms revealed that it was known in this market that most bids were always placed around that value firms. Therefore, it is reasonable to think that both groups and independent firms had acquired enough experience to know which bids would have been to high/low to win. The second feature of the data is that, however, very high/low bids are observed. The explanation is offered by some convicted firms that revealed how sometimes cartels would use what they called "support bids". These consist in extremely high (or low) bids not aimed at winning but at helping another member to win. The definition of what exactly constitutes a support bid is necessarily somewhat subjective. However, two conditions seem necessary: (a) that the bid is discontinuously greater than the bid immediately below it and (b) that the bid is in the top (bottom) end of the bid distribution. Any bid greater than a support bid is also a support bid. Figure 3 illustrates an example drawn from one of Turin's auctions which has these characteristics. In the figure, the horizontal axis lists the bidders which are ordered according to their bid (i.e. discount), which is the vertical axis. Different symbols indicate different cartels with the cross representing firms not in cartels. The vast majority of bids is around 18%. However, several members of the cartel represented by a circle submitted bids that are discontinuously greater than those of all other bidders. In this case, their strategy was successful in making a member of their coalition winning the auction (the thick blue line). Numerous similar cases are present in the Turin's data. Moreover, numerous extreme bids resembling support bids are present throughout the whole set of AB auctions that we observe.

Figure 3: A Cartel Submits "Support Bids" and Wins the Auction

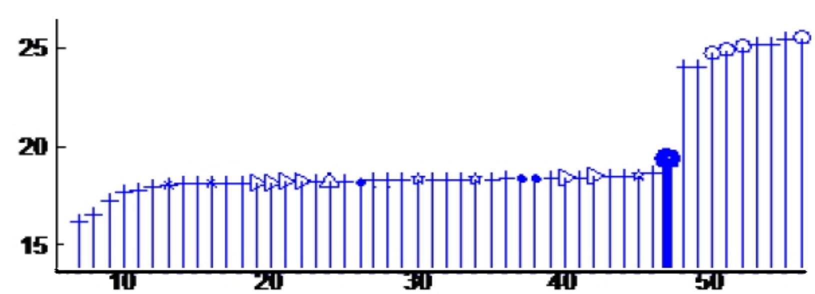

The Court documents also reveal interesting aspect of the entry behavior of cartels. In particular, they disclose that some members of the cartels are shill bidders for some other firm. This means that a firm, in order to have one additional bid has created a replica of itself, a shill, that exists for the sole purpose of allowing the original firm to place multiple bids. The Turin's case reveals that at least seven firms have one or more shills. The relevant aspect of this behavior is that once the fixed cost of constructing a shill is paid, the marginal cost of submitting one extra bids should be negligible compared to the possible increase in the expected payoff. Therefore, neither entry nor bidding choices of firms in a group should be independent and this idea will be at the basis of our tests.

Finally, it is worthwhile to stress why it is reasonable to test for group behavior in the non-validation sample of AB auctions. First, since the environment faced by the firms in the latter dataset is almost identical in every respect to that faced by the firms in Turin, it seems likely that the same incentive to coordinate exist in both cases. Moreover, like in Turin, also in the main dataset both support bids and shill bidders seem to be present. In particular, defining as support bid a bid that belongs to the highest 30% of bids and that is 5 points greater than the preceding one, the data contain 80 auctions with support bids. As regards (possible) shill bidders, in the data there are various firms registered at the same address, sharing some owners and always participating together. The final reason why we expect to find groups active in the auctions is that in 2009 the courts of two cities in the North started a case for collusion in AB auctions. In particular, like in Turin, these courts brought to trial 84 entrepreneurs accused of rigging AB auctions through two cartels.[13] Therefore, it seems relevant to study groups. We illustrate next a benchmark model of competition and some basic predictions about group behavior.

3 Model of Participation and Bidding

This section presents a model of firms' entry and bidding decisions. The model's implications are used to develop our tests for coordinated participation and bidding. We assume that there is a single contract auctioned off and we indicate with M the set of firms that might bid for it. Firms are either independent or part of groups. We model the decision problem of independent firms in two stages: in the first stage firms observe their cost for preparing the bid[14] and then, in the second stage, those firms that decided to pay the preparation cost learn their cost of completing the job and then bid. At the time of bidding, a firm does not know how many

[13] The courts are those of Treviso and Venice. They have not yet established which firms, if any, are guilty. Noticeably, the Antitrust Authority (AGCM, 1992) had expressed concerns about the pro-collusive role of the AB auctions from the very first attempts to introduce these auctions at the beginning of the 90's.

[14] The bid preparation cost might capture very different costs for independent and groups' firms. In particular, for a firm in a group the pure administrative cost of preparing and submitting a bid might be lower because of scale economies but it might also be higher if it accounts for the probability of being sanctioned by a court.

other firms will also bid. However, it is common knowledge in the market that there are $|M|$ potential entrants and that $|M^g|$ belong one or more groups. We abstain from modeling the inner working of the groups. We assume that a group is a collection of firms that delegate to a common mediator their entry and bidding decisions in exchange for a share of the group's joint profits. This mediator observes the costs of the firms in his group and decides their actions to maximize the group's profits.

Characterizing the behavior in the bidding stage

In the auction, N firms place a bid. The independents are indicated by I and the one in group g by N^g, $g = 1, ..., G$. We assume throughout the paper that $N > 4$. Each firm j has cost c_j of completing the job. Assume that $c_j \in [c^l, c^h]$ for all j=1,..,N and that each firm that enters draws its cost from a continuous distribution $F_C(.)$.[15] Before bidding, firms also observe the maximum price, R, that the auctioneer is willing to pay (the reserve price). This price is not binding: $R > c^h$. A firm submits a sealed bid, $b \in [0, 100]$, consisting in a discount over R. Therefore, the expected profit for an independent firm j that entered is: $E_I(\pi) = [(1/100)(100 - b_j)R - c_j] \Pr(b_j \ wins)$. Whether b_j wins is determined according to the Italian AB rule: the discounts' trim mean, A1, is computed as the average bid disregarding the highest and lowest 10 percent (rounded to the highest integer) of bids; then A2 is calculated as the average of the bids greater than A1 and below the disregarded top 10 percent bids; the discount closest from below to A2 wins. The winner is paid his own price and ties of winning bids are broken with a fair lottery. If all bids are equal, the winner is selected with a fair lottery. Finally, if there is a tie at the highest bid among the bottom 10 percent of bids (or at the lowest bid among the highest 10 percent of bids), the bids to eliminate are chosen with a fair lottery. We begin the analysis by looking at the case in which there is no group.

Proposition 1: In the unique Bayesian Nash equilibrium (BNE) all firms bid a discount of zero percent (zero-bids equilibrium).

The idea of the proof[16] is simple: since the highest discount always loses if it is the lone highest bid, then there must be pooling at the top. However, if the highest bid on which bids are pooled is greater than zero, then there exists a unilateral deviation toward a lower discount. The fact that this deviation is profitable is ensured by the trimming of the top 10 percent of bids together with the requirement that the winning discount lies strictly below A2. Therefore, in the unique equilibrium the auctioneer pays the highest price, R, and the allocation is inefficient since each firm wins with probability $1/N$.[17] However, this equilibrium is not robust to the presence of groups as the following proposition illustrates.

[15] Symmetry is not essential for the central results of the bidding model but it greatly simplifies the notation.

[16] The formal proof is an extension of Proposition 3 in Decarolis (2009). All other proofs are in the Appendix.

[17] Despite these undesirable properties, Decarolis (2009) shows that in a richer environment in which there is cost uncertainty and in which firms can default at a cost that is their private information, a revenue maximizing auctioneer may prefer the AB auction over a standard first price (FP) auction. In this paper, instead, we are not concerned with the auctioneer' behavior but only with that of firms given the AB rule.

Proposition 2: Unless all bidders belong to the same group or all groups are smaller than the minimum winning coalition (which is 2 plus 10 percent of N rounded to the next highest integer), then the strategy profile in which all bidders bid zero is not an equilibrium.

The minimum winning coalition, N^*, that is defined in proposition 2 is the smallest group of bidders that can break the zero-bids equilibrium. If all firms are bidding zero, then a coalition of N^* firms can submit all bids strictly greater than zero (for instance $N^* - 1$ bids equal to $\varepsilon > 0$ and one equal to $\varepsilon/2$) and win for sure. Although the winning firm will receive a lower payment form the auctioneer, there is always an ε small enough to make this strategy strictly more profitable for the group than bidding zero. If all groups are smaller than N^*, then no individual firm or group has an incentive to bid more than zero. On the other hand if all firms belong to a single large group, the winning bid must be equal to zero. Completing the characterization of equilibria beyond these cases is complicated by the need to explicitly computing the probability of winning, which is a rather intractable object. Nevertheless, the following propositions characterize some properties of group bidding that capture first order aspects of group behavior. The first one states the intuitive property that groups benefit by manipulating the bid distribution and they most profitable manipulations involve clustering all the groups' bids on the same side of the bid distribution.

Proposition 3: Assume that there is a group N^g and that $N^* \leq N^g < |N|$, then all the strategy profiles in which the group's bids do not alter the location of the trim mean are dominated by at least one strategy in which the group' bids shift A1. If the bid distribution is not degenerate, clustering all bids to the right of A1 dominates any strategy that places at least one bid below A1.

This proposition offers an ex post argument to characterize what any group would do regardless of the opponents strategy. It holds regardless of whether there are other groups present or of whether the other firms are playing according to equilibrium strategy. These latter properties are also true for the following proposition which highlights a main difference between group and independent bids.

Proposition 4: Assume that $F_B(\cdot) \sim [b_l, b_h]$ with $b_h > b_l \geq 0$ and is the bid distribution that independent bidders expect to face. Then there cannot be any Bayesian-Nash equilibrium in which an independent firm bids b_h. There cannot be an equilibrium in which the group submits at least one bid equal to b_h and does not cluster its other bids on the right tail of the bid distribution.

These two propositions contain the essential features of bids' coordination. First of all, since for any non degenerate distribution of bids the winning bid must lie in [A1, A2), then, if the group does not alter A1, it does not take into advantage the possibility of tilting the winning interval toward the area where the group places its bids. The second key aspect is that clustering the bids is the best way to shift A1 because it allows to have the bids exactly where the winning interval is moved. Finally, bidding the highest discount is advantageous only for a group and never for an individual firm. Indeed, this highest discount cannot be used to win but just to support the strategy of a group.

The argument in proposition 4 does not necessarily extend to b_l as b_l might be the only individually rational bid for high enough production costs. Nevertheless, if we assume that firms only care about winning and not about the winning bid, for instance because they can resell the contract, then again only groups would bid b_l.[18] Finally, one additional relevant feature of groups' bidding should be mixing. Essentially, this is due to the matching-penny nature of the game. Indeed, the group wants to move the winning interval but it wants to do so in a way that the other firms cannot anticipate it. A group will be particularly likely to mix when other groups are present since, by proposition 4, only another group might try to outguess the group when it is bidding near b_h.

The behavioral features discussed, if followed by a group, would induce observable differences between its bids and those of independent firms' bids. Indeed, the "bid-test" that we illustrate in the next section aims at capturing these differences. However, since we cannot argue that our analysis fully characterizes the bidders' behavior, it is crucial to assess its empirical relevance. Therefore, in the next section we apply the bid-test to the sample for which the presence of groups is known.

Characterizing the behavior in the participation stage

The above discussion makes clear that in the AB auction a group can improve its expected payoff by coordinating bids. However, a group needs the participation of at least the minimum winning coalition if it wants to gain from bids' coordination. From the Court case concerning the cartels in Turin, we know that groups play complex participation strategies sometimes involving bribing independent firms to bid with them for a single auction. However, since our data is not rich enough to measure precisely this phenomenon, we abstain from modeling the inner working of the cartel entry choice. Instead, we simply assume that a group's mediator trades off the benefit of one additional bid to manipulate the average and increase the probability of winning against the cost of the additional bid preparation. Therefore, we expect that a firm in a group is more likely to enter if also other $N^* - 1$ firms from the same group enter because of the greater expected payoff. For the independent firms, instead, we assume that every firm j independently draws a participation cost $q_j \tilde{} F_Q(.)$. Therefore, if we define the expected profit before independent firm j observes its cost as $E_I(\pi_I)$, then such a firm follows the following cutoff rule: enter if $q_j \leq E_I(\pi_I)$ and stay out otherwise.

Since independent firms are not aware of how many groups and of how many firms per group will enter, their expected profit from participating is constant. Therefore, the independent firms' entry decision is independent across these firms. On the contrary, the entry of a group member is more likely when at least $N^* - 1$ other members enter too. Our "participation-test" is based on these ideas.

[18] Resell is not legal in the AB auctions that we study but subcontracts are, although within certain limits.

4 Econometric Tests

In this Section we present our participation and bid tests. We also examine their performance when used with our validation sample data.

4.1 Participation Test

Participation patterns among groups of firms within a suspected set of colluding firms have considerable potential to identify collaborators. Multiple firms within a colluding set obviously have to be present at an auction to influence its outcome. In the case of AB auctions there is a minimum winning coalition size. The descriptive evidence from Turin suggests that coordinated entry is clearly present.

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

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

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

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

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

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

The choice of the set of comparable firms M will be a key decision for implementation of our participation test. For the Italian roadwork procurement auctions 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.

The choice of the number of firms in g is also an important decision. Relatively large and small choices of g may be the most informative. When the group is large, for a fixed set M, power should be good as coincidental attendance of a large group of innocent firms will be unlikely. Using a small group of two should also have good power as the minimum winning coalition must have at least 3 firms. Therefore, size-two groups formed from a set of colluding firms should be less likely than a group of two innocent firms to have both members coincidentally attend an auction.

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

Validation Sample Results: We report the results obtained for the Turin cartels in the right panel of Table 4. In particular, the table reports a one whenever the cartel is identified as such by the participation test at 5% significance level. The first of results are obtained by looking at the cartel altogether. Instead, the last two columns of the table report the fraction of firms classified as cartels if the firm-specific version of the test is used. The results reported in the table are of great relevance because they show that regardless of whether we consider the unconditional version of the test (i.e., the one in which the comparison groups are created by drawing from the set of all bidders) or not, the test suggest that all chosen groups are made of not independent firms. In particular, the result holds both when the comparison groups are chosen by matching the same distribution of capital and distance from Turin of the firms in the cartel (Cond 1) and when they are chosen to match the legal requirements for participation of the firms in the cartel (Cond2). Overall, based on the participation of the largest possible subgroups, independent entry is rejected for all the 8 cartels. Hence, these results are broadly consistent with the prediction that members of groups are more likely to enter together in an auction. The result of the unconditional test for cartel 1 is given a graphical representation in Figure 5. The highest of the red lines is the 95 percent of the distribution of joint participation of random groups of the size reported on the X axis. In this case, our test simply amounts to check that at the highest number of participants (12), the frequency of auctions attended by cartel 1 is greater than that of groups of random firms.

As regards the last two columns of Table 4, they report the fraction of firms in the group that are classified as not independent. To classify a firms as independent we use the firm-specific participation test at 5 percent significance for all firms in all cartels. This amounts to check whether the largest set of firms that have all jointly bid with the chosen firms is that composed by random firms (or, better, the 95 percent of this distribution) or by members of the same cartel. The results are suggestive that our test captures most of the firms belonging to the cartel.

Figure 5: Participation of Cartel 1

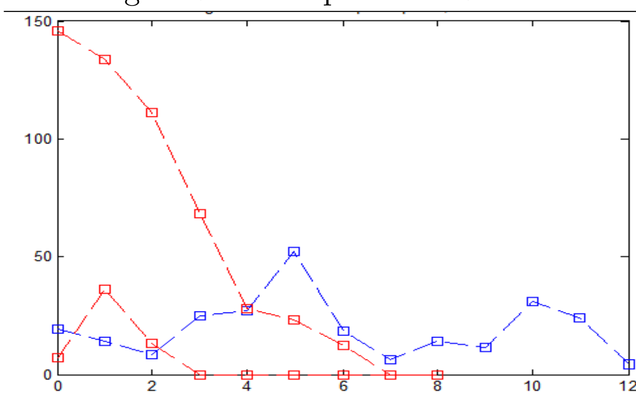

Notes: The blue line indicates the number of auctions attended by subgroups of cartel 1 of size up to 12. The red lines report the same value but for the bottom and top 5 percent of the random groups that participate most often.

4.2 Bid Test

Our bid test is based on a likely basic strategy for a colluding set of firms: a subset of the members clusters their bids in an attempt to pilot the trimmed mean. This strategy certainly seems consistent with patterns in the bidding behavior in our validation sample from Turin.

For ease of exposition, we first present our test without conditioning on any firm or auction characteristics.

We base our test on a measure of how much influence a given set of suspected firms has upon a trimmed mean bid for an auction. First, for a given set of suspected firms define a subset g for a test of whether these g firms are piloting the trimmed mean in an auction. 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 bids from group g and $B^{-g} = \{b_1^{-g}, ..., b_{N-N^g}^{-g}\}$ as the ordered set of remaining bids. The trimmed mean throwing out N' bids[19] 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 bids if the group is trying to pilot the overall trimmed mean up/down. Formally, we test the null hypothesis that the firms in group g are not cooperating to pilot the overall trimmed mean. Our operational definition of 'not cooperating' is that firms are bidding independently.

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

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

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

Our test decides whether a group of firms has unusually coordinated bids by checking whether the realization of $A1^g$ is a tail event relative to the distribution of $A1^s$. A two-sided version of this test at say the 5 percent significance level corresponds to the following decision: reject the null if $A1^g \notin [P^T_{.025}, P^T_{.975}]$ where $P^T_{.025}$ and $P^T_{.025}$ are the $2\frac{1}{2}th$ and $97\frac{1}{2}th$ percentiles of the distribution of the $A1^s$. One-sided tests likewise will reject if $A1^g$ is higher or lower than the corresponding critical values given by the appropriate tail percentile of the $A1^s$ distribution. We approximate the $A1^s$ distribution via simulation, drawing a large number of times from S and calculating for each draw s the corresponding $A1^s$.

[19] N' for the Turin auctions will be the 10th and 90th percentiles, rounding up to solve the integer problem.

A straightforward extension of this test is to account for firms' observable characteristics, Z. In principle, if Z has few possible values we can construct the random groups to match exactly the frequency of Z in the suspect group g.

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

Multiple Auction Testing

Our bid test as stated above applies to a single auction. With data on multiple auctions we may have the capability to conduct multiple tests for a given suspect group g. This multiple testing situation presents a formidable challenge if we allow for firm-level persistent idiosyncrasies in behavior. Our operational definition of non- cooperating needs to be augmented with respect to a firm's actions in multiple auctions. In particular, even if is reasonable to use a benchmark that non-cooperating firms act independently within an auction, we should allow for a given firm's actions to be correlated across auctions in which it participates. The bottom line is that when taking a set of firms s, the set of $A1^s$ outcomes across multiple auctions will not be independent.

First consider a bid test across two auctions. We form a joint test statistic for a suspected-average-pilot group g of firms who participated in both auctions with bid test statistics $A1_1^g$ and $A1_2^g$. For example, using the indicator function $1(\cdot)$, a test statistic J^g could be formed as:

$$J^g = 1(\text{Both } A1_1^g \text{ and } A1_2^g \text{ are such that our one-auction bid tests rejects no cooperation}).$$

This test statistic J^g obviously involves the same set of firms g in both auction one and two, statistics that could be arbitrarily dependent. In order to capture an arbitrary dependence across auctions in $A1_1^g$ and $A1_2^g$, we need to use the corresponding distribution for a randomly selected group s that participates in auctions one and two. Our reference distribution for J^g under the null hypothesis of no cooperation is the implied distribution of

$$J^s = 1(\text{Both } A1_1^s \text{ and } A1_2^s \text{ are such that our one-auction bid test rejects no cooperation}).$$

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

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

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

Validation Sample Results: We first present the results of our one auction bid test for the 8 cartels in Turin. The tests are not conditional on any firm's covariates because, despite the richness of our data, we could not find any attribute that (alone or jointly with others) was robustly associated with firms' bids.[20] To conduct the test, we fix a cartel. Then for each auction we remove the bids of firms in this cartel[21] and compute the trim mean $A1^g$. Then for each auction we repeat 1000 times the calculation of the trim mean but each time we exclude a new set of randomly drawn firms. For each auction we look at these 1000 trimmed means $A1^s$ and we find the percentile of the distribution of simulated $A1^s$ equal to $A1^g$. If firms in g coordinate their bids to push up A1 in a given auction then the corresponding $A1^g$ will tend to be in the left tail percentiles near zero. If instead they push the trimmed mean down, it will tend to be at high percentiles near one. The values taken by the percentiles of $A1^g$ across all

[20] The only exception is the firms' groups affiliation on which, however, we cannot condition.

[21] We consider legal joint bids as cartel's bids if at least a member of the consortium belongs to a cartel.

auctions are reported in the histograms in Figure 4.

Figure 4 shows that with the only exception of the sixth group, all others are remarkably different from a random group of firms. Interestingly, cartels 1, 2, 3 and 7 seem somewhat more prone to push up the average bid. Although this behavior has positive effects for the auctioneer's revenues, it likely has positive returns for these firms since they achieve a higher number of victories (and revenues) relative to the firms in cartels 5 and 8 which bid to push down the trim mean. It appears that focusing on the one tail test, may be enough to identify the most interesting cartels. Finally, the result on cartel 6 is not negative for our methodology. In fact, although we do not find evidence of systematic bids' coordination, this is the only cartel whose members were not charged for "criminal association" because their coordination was sporadic.

Insert Table 4

The results of the histograms are summarized in the left panel of Table 4. A one is reported whenever at least 30 percent of the auction lead to a result for the p-value of the bid test that is either below 0.25 or above 0.975. As discussed above, only cartel 6 is not classified as a cartel. Moreover, the results confirm that it is undesirable to use covariates to construct the random groups in the bid test: draw from small bins of similar firms to construct the random groups implies that we construct random groups containing a high proportion of the original cartel firms. However, as explained above the results in Table 4 cannot be used directly to draw conclusions based on the repetition of the bid test. Instead, if we want to use the multi-auction nature of our data we need to compute the distribution of J^s and assess where J^g falls within this distribution. Table 5 reports the result of this analysis for cartel 1 for two main cases: when the groups used to compute the J^s are composed only of those firms that were never suspected of collusion and when these groups use the mixture sample with both innocent firms and members of cartels other than cartel 1. One major obstacle to this analysis is that it is impossible to find a large number of independent firms that frequently bid together. Therefore, we restrict the attention to groups of size equal to 4. For cartel 1 these 4 firms are the ones that bid most often. We implement the test using as the rejection criterion that enters in the specification of the J^s that of a the one sided single-auction bid test that rejects independence in favor of coordination when $A1^g$ corresponds to a percentile no larger than 5 percent. We perform this analysis for multiple J_T^s corresponding to a different number of auctions T=2, 3,...,10. For each case we repeat 1000 times and we record which percentile of the distribution of J_T^s corresponds to J_T^g. We report the 10th and 90th percentiles obtained through the 1000 repetitions. We want to reject independence in favor of coordination when the percentile reported in the table is below the significance level that we demand. The first row of Table 5 reveals that when we have only two auctions we are unable to reject independence. However, as the number of auctions increases, we become more likely to reject the null that the subgroup if cartel 1 that we are studying behaves like a group of independent firms. In the table it is also clear the trade off between using a sample of truly independent firms but with which we can construct few control groups and using a sample that, despite being mixture by cartel firms, allows to construct a larger number of random groups. Overall, we see that with the mixture sample having at least 10 auctions is sufficient to deliver strong evidence against the fact that cartel 1 is a collection of independent firms. This is reassuring since in the large sample of auctions from the North we will only have mixture samples. Finally, the results for all the cartels and for different rejection rules for the single-unit bid test are reported in the Web Appendix.

Figure 4: Bid Test for the 8 Known Groups

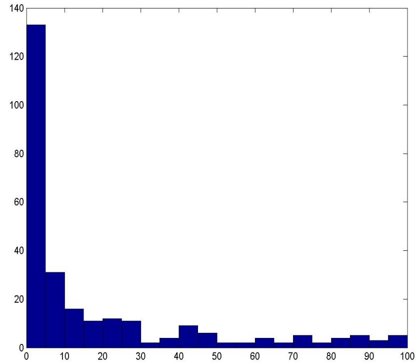

(a) Test Histograms - Cartel 1

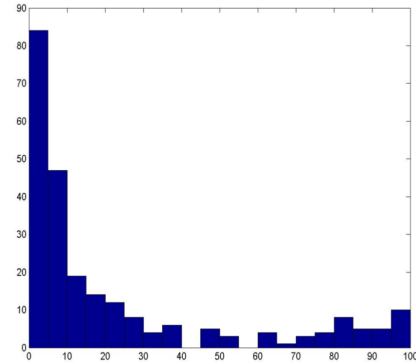

(b) Test Histograms - Cartel 2

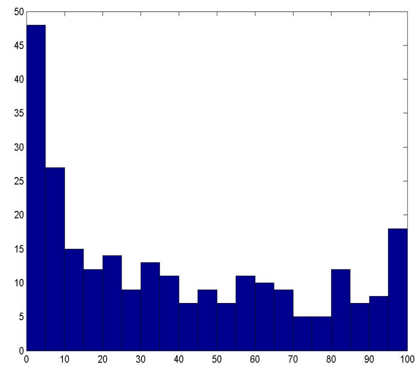

(c) Test Histograms - Cartel 3

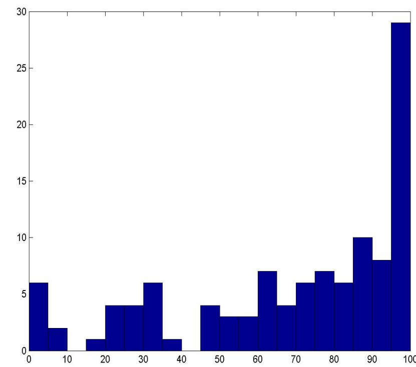

(d) Test Histograms - Cartel 4

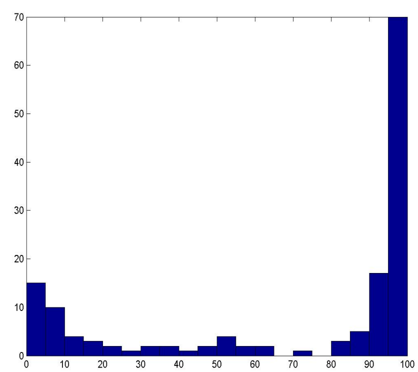

(e) Test Histograms - Cartel 5

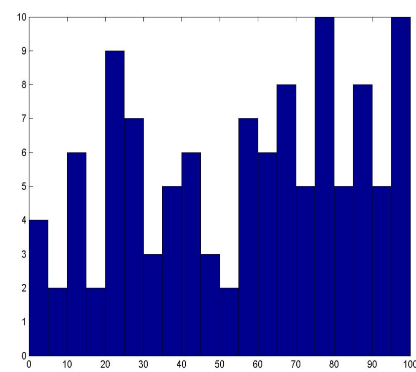

(f) Test Histograms - Cartel 6

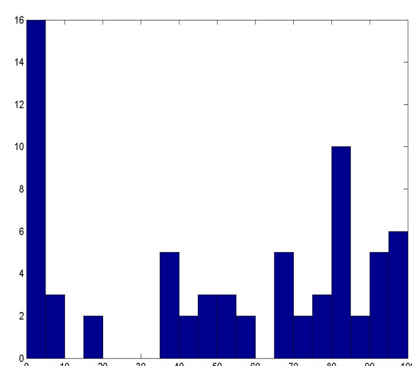

(g) Test Histograms - Cartel 7

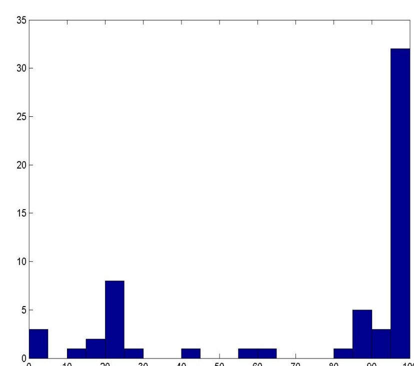

(h) Test Histograms - Cartel 8

5 Testing Coordination with Unknown Groups

This section applies the tests for coordination to auctions where there is no prior knowledge of bidders' groups. In principle, this is a simple task because, given a candidate group, we just need to test it with the bidding and entry test to decide whether its members coordinate their actions. However, the problem consists in appropriately choosing the candidate groups so to avoid checking all the enormous number possible firms' combinations. In this section we describe two methods that we propose to solve the problem. The first method can be used when the researcher observes firms' covariates. These covariates are used to measure the probability that couples of firms are linked together. Then a clustering algorithm uses these probabilities as inputs to generate candidate groups. The second method is less demanding in terms of data because it determines potential groups only on the basis of firms' identity and entry. However, groups determined with the latter method can only be tested with the bid test because they would fail the entry test by construction. Notice also that this part of our analysis is similar in spirit to Bajari and Ye (2003) which shows how to test for collusion in first price auctions without prior knowledge of groups.

Method 1: Observable Firms' Characteristics We use firms' characteristic to construct groups of potentially coordinating firms and then test these groups. We explain the method in steps.

Step 1: In the first step, firms' characteristics are used to construct links between couples of firms. In particular, we identify a link between two firms when: they share some of the owners (managers), or they are geographically close, or they bid together in a consortium, or one did a subcontracting work for the other. First, we quantify the links between the 95 firms in Turin, ending up with a sample of 662 couples of firms connected by at least one link. Since for these firms we know the composition of the cartels, then we are able to tell which of these 662 couples truly belongs to the same group. Hence we can run a probit regression in which the dependent variable is equal to one if the couple is in the same cartel and zero otherwise. For our favorite specification of the model, Table 6 reports the marginal effect of switching from zero to one the various links. The variable Personal is equal to one when the two firms share any owner (top manager). Subcontract, instead, equal one if they ever exchanged a subcontract. The estimated marginal effect is the largest for this variable. The joint bidding variables are, instead, equal to one when the firms formed at least once a bidding consortium in the Turin's data (Joint-Bidding-2) or when they won at least one auction as a consortium in all the auctions held in Piedmont between 2000 and 2003 (Joint-Bidding-1). We exclude all pairs that are linked exclusively by geographical proximity, because, although location helps to identify groups, it exposes these groups to the criticism that any observed failure of independence could be due to the spatial correlation of costs.

For the firms in the dataset where no groups are know, we can use firms' characteristics to create a set of linked couples. Then we can use these linkages together with the estimated coefficients from the probit using Turin's data to forecast the probability that two firms are together in a group. Given N firms, we can then construct an NxN "dissimilarity matrix" which is symmetric and has ones on the diagonal. In the off-diagonal entry (i, j) the matrix has the complement of the predicted probability that the firms i and j are in the same group.

Insert Table 6

Step 2: In this step we use a clustering algorithm to create groups. The algorithm is a standard hierarchical algorithm (Gordon, 1999) which uses as its input the dissimilarity matrix constructed in step 1. This algorithm associates firms (or group of firms) together on the basis of their average dissimilarity. The technical details are provided in the Web Appendix. Instead, to give an idea of how the algorithm works, we report the results obtained by applying it to the set of colluded firms assuming we do not know anymore to which group they belong. Instead of the original 8 groups, the algorithm produces 15 groups.[22] The dendrogram in Figure 6 illustrates the aggregation. The composition of each group is reported in Table 7. Although we did not exactly recover the true groups, the ones produced are almost all subgroups of the original cartels. Only in one case there is a group containing firms belonging to different cartels. In all other cases the groups are either pure subgroups of the original cartel or they contain at most two firms not belonging to any cartel. We do particularly well for cartel 1 for which one of our groups cluster together 10 of its members. In one of the small 2-firm groups composed by one suspect and one independent, the independent firm is owned by the father of the owner of the suspect firm. Therefore, it appears that this method offers a sensible way to use firms' characteristics to obtain candidate groups.

Figure 6: Dendrogram

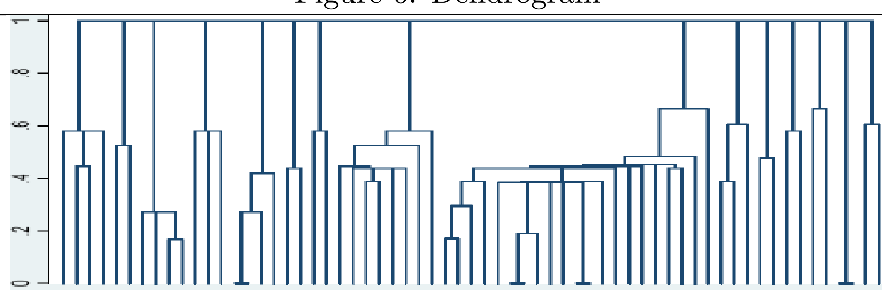

Y axis: dissimilarity (i.e., the complement of the predicted probability of a correct match)

Step 3: The final step consists in applying the entry and bid test to the candidate groups. For instance, Table 7 applies the tests to the 15 groups recovered for the validation data. The details of how the tests are applied are reported at the end of the table.

[22]We test the validity of these clusters using the Monte Carlo approach described in chapter 7 of Gordon (1999). We rejected that the clusters are identical to random groups of firms at 5 percent significance.

Method 2: No Observable Firms' Characteristics Auctions datasets often contain only information on bidders identities and bid. This information is enough to conduct our tests. Nevertheless, when candidate groups need to be found and no firms' characteristics are observable, a possible solution is to preselect firms on the basis of their joint participation. However, this will prevent us from using the participation test, which would reject independence by construction. We know from the case of the validation sample that without having access to firms observable characteristic the usefulness of the participation test is limited. Instead, the bid test works well even without this information. Moreover, using participation patterns allows the formation of candidate groups that exhibit a relevant feature of true groups (coordinated entry) and, hence, that are particularly interesting candidates for the test of coordinated bidding.

Insert Table 7

The main idea of this method is first to identify the most frequent winners and then, for each winner, to construct his candidate group by looking at those firms that participate with him the most. Therefore, cartels are created with a two-step procedure. Step 1: choose the group head by selecting those firm that win suspiciously too much. For instance, we select those firms that win more than an hypotetical independent firm attending the same auctions (absent collusion). Step 2: is a simple iterative procedure. Given the set of winners, construct a group with other N members by looking at the frequency of joint participation. The first firm attached to the group is the one that participates most often with the frequent winner. Then we attach the firm that participates most often with the couple created in the previous step. We continue in this way until we reach the desired group size. To illustrate the performance of this method, we perform it on the firms in Turin and report the results in Table 7. We pretend we do not know the true groups and we impose a group size of 4. Compared to the clustering algorithm, the main error consists not in associating independent firms to cartels but to combine together firms belonging to different cartels. Nevertheless, we obtain several flawless groups and the most informative (one-sided left) bid test rejects independence for all of them.[23]

6 Quantifying the Presence of Groups

The methods illustrated above to test coordinated behavior when groups are unknown can be used to quantify the presence of groups in the AB auctions in our dataset. Since we have more than one test (an several variants of each test), we have multiple ways to define which candidate groups should be considered as groups of coordinating firms. Moreover, we could opt for different levels of significance or require a minimum fraction of auctions in which independence

[23] The exact details of this two methods as well as the results of a different way to compare them are reported in the Web Appendix.

has to fail to decide that some firms formed a group. Having made these choices and obtained a classification of firms between independent and not, it is possible to quantify how many auctions are rigged. However, also in this case we shall decide, for instance, whether every auction participated by at least two firms belonging to the same group has to be considered rigged or whether only auctions in which at least one group fails to pass independence according to the bid test have to be considered rigged.

For illustrative purposes, we use the clustering method described in the previous section with the 164 auctions held in Piedmont. Since firms in groups should be the ones winning more frequently, we focus our attention on the 187 firms that won at least one of the 349 auctions held in Piedmont. For each of these firms, we construct their links to every other bidder in the sample, we obtain the dissimilarity matrix and finally the groups. For the candidate groups, we classify them as true groups only if the (one-sided left, unconditional) bid test at 5 percent level indicates coordination of bids in at least 50 percent of the auctions that they attend. Having identified the groups, we then count in how many auctions at least two firms from the same group bid. Table 8 reports the results obtained at the different significance level of the two tails bid test.

Insert Table 8

In case we classify as group any set of firms that fails bids' independence at least once, then the percentage of rigged auctions would be between 80 and 90 percent. More precise results will be provided in the next version of this paper. In any case the basic picture appears clear: a substantially large fraction of the AB auctions is affected by the presence of groups.

6.1 The Effect on Revenues

Having established that the presence of groups is pervasive in the AB auctions, we would like to evaluate its effects on revenues. As regards the auctioneer's revenues, the interesting insight is that in the AB auction the presence of groups might benefit the auctioneer. We have already discussed that with multiple, sufficiently large groups their competition will prevent a zero-bid equilibrium. Moreover, since the zero-bid equilibrium is unique under full competition of independent firms, in principle the calculation of the benefits for the auctioneer is trivial. It is simply the difference between the reserve price and the true winning price. Given the large volume of AB auctions and the rather high average winning bid (13%), the savings due to the presence of groups are substantial. However, the limitation of this naive calculation is that were all the bids to converge to zero, it is unlikely that the auctioneer would not modify the auction format. Indeed, it is known that the legislators introducing the AB auction were not expecting all the bids to go to zero, therefore they might abandon the mechanism if this would happen. The problem with the AB rule is that the incentives to form groups are so strong that a zero-bid situation is very unlikely to emerge.[24]

In the case of the validation data, the convicted firms are currently facing the prospect of having to pay damages to the city of Turin. In the top panel of Table 9, we report both the actual cost of procurement (i.e., the sum of all winning prices) and some (naive) counterfactual. The first counterfactual looks at what would have been the cost under the zero-bid equilibrium (competition): the saving due to coordination is about €23 million. However, different counterfactual scenarios are currently contemplated in Court. In the first scenario, the cost is simply recomputed by eliminating all the cartel's bids (assuming no changes in the other bids). In the other two scenarios, the bids of a cartel are converted to the average discount (18 percent): all bids are converted in the second scenario, while only 2/3 are converted (and 1/3 eliminated) in the third scenario (as if 1/3 of the cartels were shills). Table 9 reports the cost under these scenarios for cartels 1 and 5 showing that even under these scenarios coordination is not necessarily harmful for the auctioneer.

Insert Table 9

Nevertheless, the activity of the groups is harmful for firms outside the groups. The loss for firms outside the groups is due both to their lower probability of winning and to the fact that, when they win, they do so at a discount greater than zero. The bottom panel of Table 9 illustrates the results of some naive measures of this damage. The first column reports the actual distribution of revenues among independent firms for the validation sample. The following column reports the distribution of revenues that would have resulted if entry was unaltered and the zero-bid equilibrium was played by all firms. The last column does the same but excludes 1/3 of the cartels' bids.

The estimates presented in Table 9 are naive because they assume that the other players do not respond to the change in cartels' behavior. However, following the approach of Asker (2009), it may be possible to estimate a structural model to evaluate the losses of independent firms.[25] Since the zero-bids equilibrium pins down bidding, the contribution of the structural model would be allowing an estimate of entry. This could be done applying to the subsample of first price auctions some recently developed empirical models of entry in FP auctions. Although, we abstain from taking this route in this paper, we acknowledge that this step is needed to account

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

[25] For the auctioneer, as long as the auction format is AB, the zero-bids equilibrium implies that the only relevant action is optimally choosing the reserve price.

for the higher entry in the AB auctions under the zero-bids equilibrium. This is essential to correct for the overestimated losses of independent firms presented in Table 9.[26]

6.2 Drop in Participation

A different aspect of firms' entry that can be usefully studied through our test is the drop in participation that followed the introduction of FP auctions. One of the most striking features of the Italian AB auctions is the phenomenally large number of firms bidding. As shown in Table 1, the AB auctions receive on average 51 bids and auctions with more than 100 bidders are common. Instead, after the switch to FP auctions the number of bids per auction went down to an average of 7 which is line with the turnout at similar auctions in the US. The many studies on FP auctions for road construction contracts in the US report an average bidder turnout that ranges from 3 to 7 firms per auction. When in 2003 the case against the colluding firms started, Turin municipal and county councils imposed the compulsory use of the FP format for all procurement auctions. Similarly to what happened later on in the rest of Italy, the drop in the number of bidders was striking. Figure 7 documents this change by reporting in the left panel the distributions of the number of bidders in Turin both under the AB and the FP auctions (i.e. before and after 2003). The right panel, instead, shows that for all the other local administrations comparable to Turin (in terms of geographic location and size of the population served) that remained with the AB format there was an increase in the number of bidders attending the auctions in the period August 2003 - January 2008 as compared to the period December 2000 - December 2002.[27]

[26] A second problem of the naive estimates, is that it might be inappropriate to use the sample of Turin's auction to evaluate the damages of independent firms because these auctions were selected by the legal office of the city of Turin as the most representative of the cartels' activities. Therefore, the cartels' probability of winning is likely overestimated in this sample.

[27] Also the dynamic over time of the bidders' turnout indicates an interesting correlation with the auction format. In fact, while over time for Turin's auctions the number of bidders kept declining after the introduction of the FP, for the auctions of the other administrations the turnout drastically increased. This could be explained in part by a greater thirst for work caused by the decline in the number of auctions after 2004 (the total value of all public contracts for works decline from about 25 billions of euro in 2004 to about 20 billions in 2006). However it is also possibly evidence of the fact that over time firms understood that the national law transformed the auctions in lotteries and hence payoff maximization could be more helped by a rise in the probability of winning than by a rise of production efficiency.

Figure 7: Distribution of the number of bidders participating at auctions

Treated Group (Provincia & Comune Turin) | Untreated Group (All Other Piedmont Loc. Ad)

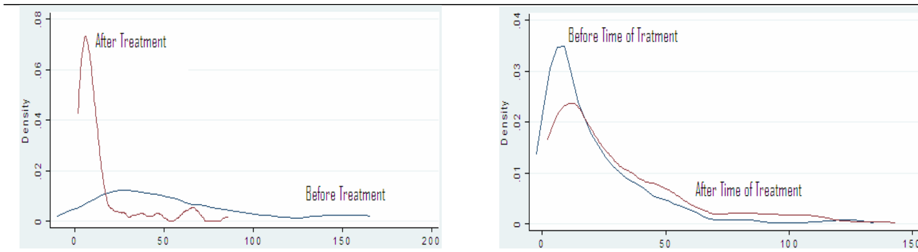

The results of the econometric analysis support what indicated by the raw data densities. The results in Table 10 indicate that a switch from AB to FP is associated with a drop of about 40 bidders. Since the variable measuring the number of bidders is highly not normal (the skewness and kurtosis are respectively much greater than zero and three), the model used is a negative binomial regression with robust standard errors. The negative binomial model is preferred to a Poisson regression because the variance of the number of bidders variable is quite larger than its mean and the estimated coefficient on over dispersion in the negative binomial model is statistically different from zero.

Insert Table 10

These results are not surprising in light of our previous analysis. However, it is hard to decompose this effect between the disappearance of shills and that of true but inefficient firms. Nevertheless, the large size of the market shakeout produced by the change in the auction format makes disentangling the two effects particularly worthy. In this regard, our tests for collusion allow us to make the following considerations. Suppose that we observe a set of firms bidding in some AB and that we classify them between independent and groups' members. Then if an independent firm disappears from the market after the introduction of the FP auction we can claim that this firm exits because it is inefficient. On the other hand, if a firm that we classified as part of a group exits we cannot tell whether it does so because it is a shill of some other firm or because it is a weak member of a group. If we apply this approach to the auctions held in Piedmont and we use the same classification criteria used for Table 8 we obtain that: 288 firms belong to groups of firms coordinating their actions and 966 do not. Of the latter ones, only 264 keep on bidding after the switch to the FP auctions. Therefore, the exit of the remaining 702 firms is likely due to their inefficiency.

7 Conclusions

We constructed two tests that perform well in detecting groups active in AB auctions. Although no statistical test is a final proof, our tests could be useful instruments for the Courts evaluating cases of coordinated bidding. Even if firms were informed about our tests, avoiding detection would require for them renouncing, at least in part, to the benefits of coordination. In this sense our tests have the nice feature of being somewhat "inspector proof". Finally, we believe that the application of our tests to the Italian market have uncovered relevant features of the firms' behavior and of the nature of this market. Our study confirms that firms strategically respond to the incentives generated by the AB rule and that the use of not strategic models by the proponents of the AB rule is incorrect. Therefore, we hope that our results will help to shape the discussion about the various forms of AB rules that are used not only in Italy but in numerous other countries.

8 Bibliography

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

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

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

Asker, J., (2008). "A Study of the Internal Organization of a Bidding Cartel," *American Economic Review*, 100, 3, 724-762.

Autorita' Garante per la Concorrenza e il Mercato, (1992).

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

Burguet, R., J. J. Ganuza and E. Hauk, (2009). "Limited Liability and Mechanism Design in Procurement," mimeo.

Decarolis, F., (2009). "When the Highest Bidder Loses the Auction: Theory and Evidence from Public Procurement", Bank of Italy Working Paper 717.

Decarolis, F., (2010). "When the Highest Bidder Loses the Auction: Theory and Evidence from Public Procurement", mimeo.

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

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

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.

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

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

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

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

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

Li, T., (2005). "Econometrics of First-Price Auctions with Entry and Bidding Reservation Prices," *Journal of Econometrics*, 126, 173-200.

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

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

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

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

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

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

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

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

9 Appendix: Proofs

Proof of Proposition 2: if the coalition is all inclusive, offering a zero discount is the best that can be done. Therefore, if $|N^g| = |N|$ all equilibria have the winning discount equal to zero and at least $|N| + 1$-integer$^+\{(.10)|N|\}$ (where integer$^+\{x\}$ is x rounded off to the highest integer) bids equal to zero. Instead, for coalitions that are not all inclusive, the relevant "minimum winning coalition" in defined as $N^* = 2 + \text{integer}^+\{(.10)(|N|)\}$. Any group that can submit at least N^* bids has profitable deviations when all other discounts are equal to zero. One such deviation is to place $N^g - 1$ identical bids, all equal to ε, for small $\varepsilon > 0$, and the remaining bids equal to $\varepsilon/2$. This strategy gives to the group (approximately) the highest payoff in case of victory and a probability of winning of one (prior to the deviation the probability of winning was N^g/N). However, if the group does not reach a size of at least N^* it cannot profitably deviate from the zero-discount equilibrium because all its bids away from zero would have a zero probability of winning due to the trimming and the requirement that the winner is below A2.

Proof of Proposition 3: We first show that a group always gains from altering the location of the trim mean. Consider the set of all bids except those of the group N^g. Four statistics can be computed with these bids: the trim mean (A1), the trim mean augmented by the positive standard deviation (A2), the lowest bid within the top 10 percent disregarded discounts (Top Bid or TB), and the highest bid within the bottom 10 percent disregarded discounts (Bottom Bid or BB). By definition, the following relationship links the four statistics: $TB \geq A2 \geq A1 \geq BB$. Therefore, there are 8 cases:

Case 1	$TB = A2 = A1 = BB$	Case 5	$TB > A2 > A1 > BB$
Case 2	$TB > A2 = A1 = BB$	Case 6	$TB = A2 > A1 > BB$
Case 3	$TB > A2 = A1 > BB$	Case 7	$TB > A2 > A1 = BB$
Case 4	$TB = A2 = A1 > BB$	Case 8	$TB = A2 > A1 = BB$

First of all, notice that cases 7 and 8 are not possible. Let us define "salient bids" all those bids that are not trimmed. If $A1 = BB$, then all salient bids must be identical and equal to BB. However, this is impossible since $A2 > A1$ implies $A2 > BB$. For all the remaining cases, we will sow that compared to all group's bids that would leave A1 unaltered, we can find at least one strategy which alters A1 and leads to strictly greater gains for the group. Among the many strategies that leave A1 unaltered, several strategies entail placing some bids that are not salient. However, since non salient bids can never win, we disregard these strategies as they are weakly dominated by strategies placing only salient bids.

Case 1: all bids are identical to some b. Hence, not affecting A1 requires that the group's bids are also all equal to b. However, if b=0, theorem 2 shows that a unilateral profitable deviation exists and that its usage would push A1 up. The same logic gives that, if b>0, there is a profitable deviation by placing bids between b and zero. This strategy would push A1 down.

Cases 2, 3 and 4 are almost identical to Case 1. The reason is that A2=A1 implies that all salient bids are identical to the same b. Therefore, not moving A1 requires bidding b but this is a dominated strategy. In Case 2, it is possible that A1=0 and we know that in this case a small deviation above zero is profitable. Apart from this situation, in all three cases a deviation towards a bid lower than A1 leads to higher expected profit. Clearly the payoff in case of victory will be strictly higher. Moreover, the probability of winning can be made equal to one in all three cases. Each of these deviations leads to a change in A1.

Case 5: the winning bid must be in the interval (BB,A2). There are two basic situations in which the group's bids leave A1 unaltered: (a) placing all bids equal to A1 and (b) placing at least some bids on both sides of A1. The former strategy leads to a victory only if there are no bids in (A1,A2). Therefore, a strategy that replicates (a) but that places a bid equal to A1+ε achieves almost the same payoff in case of victory of strategy (a) but has a strictly higher probability of winning. Now consider strategies (b), they can be replicated for all the bids (if any) that have a positive probability of winning, those bids in [A1, A2), and strictly improved by placing the remaining bids strictly within this interval. In particular, these remaining bids are placed symmetrically around A2 so that A2 does not change, this type of strategy leads to a victory every time the strategy (b) was leading to a victory and it strictly increase the probability of those bids that (b) placed below A1. To ensure that these last bids do not lead to a lower payment in case of victory they can all be placed below the highest bid below A2 that the original (b) strategy was placing. Clearly, this type of strategy pushes A1 up. For Case 6 the the argument is identical to the one of Case 5.

The second part of the proposition says that any strategy that achieves an increase of A1 but leaves some bids below the original A1 is dominated by a strategy that clusters all bids above A1. The argument is again based on replication: any strategy with just some bids above A1 can be replicated by a strategy in which all bids above A1 are left unchanged and those that were below are moved between the original A1 and the highest group bid below A2. Through this replication, the group achieves a payoff in case of victory that is at least as large of that of the original strategy and a probability of winning that is strictly larger. The reason why the argument is not symmetric for downward shifts in A1 is that some shifts of A1 might leave A2 unaltered. Therefore, in all these cases, it is not true that all bids between the original A1 and A2 are worthless. Instead, they are the ones most likely to win in these events.

Proof of Proposition 4: Proving the first part of the proposition simply requires noticing that a bid equal to b_h has zero probability of winning unless all bids are equal to b_h. Therefore, since $b_h > 0$ there is always a unilateral profitable deviation by bidding in (b_l, b_h). Therefore, let us define $b_h^I < b_h$ the highest bid that an independent bidder would submit. To prove the second part, notice that regardless of the expected distribution of bids, if the group places a bid equal to b_h this must rise the A1 expected by the group above $\bar{b}$. Hence, any bid $b < \bar{b}$ has now a lower probability of winning. This strategy is therefore dominated by a replication strategy in which all bids are identical with the exception of the lowest one that is now moved to b_h^I.

10 Tables

Table 1: Summary Statistics

Auctions for roadwork contracts below €1 million, Nov 2005 - May 2010

Statistics by Auction							Statistics by Firm						
	Mean	SD	Med	Min	Max	Obs		Mean	SD	Med	Min	Max	Obs
AB Auct.							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
Win-2Bid	.24	.68	.07	0	9.4	802	Reven	170	1081	0	0	$4e^{04}$	4005
With.SD	2.9	1.4	2.7	.14	9.2	802	Age	22.3	13.8	21	1	106	3611
No.Bids	50.7	34.3	43	5	253	802	Capital	447	2411	52	10	$8e^{04}$	2484
Res.Price	312	204	250	11	999	802	Subct	.65	2.9	0	0	53	4005
							Miles	159	234	47.8	0	1102	4005
FP Auct.													
WinBid	28.9	9.9	29	1.2	53.4	232	Frims that ceased activity						3.4%
Win-2Bid	4.5	5.0	3.0	.01	41	232	Location of firms headquarter:						
With.SD	6.9	3.1	6.6	.07	19.1	232	North5						69.6%
No. Bids	7.3	5.5	6	2	48	232	Center and other North						18.4%
Res.Price	342	288	215	30	978	232	South and Islands						12.0%

Notes: Auctions for roadwork contracts procured by municipalities located in 5 regions in the North: Piedmont, Liguria, Lombardia, Veneto, Emilia-Romagna (North5 regions). Top left panel: statistics by auction for the sample of AB auctions. The variables HighBid is the highest discount, while WindBid is the winning discount. Win-2Bid is the difference between the winning bid and the bid immediately below it (sometimes referred to as "money left on the table"). Notice that in the AB auctions Win-2Bid is frequently equal to zero. Mre generally, for the AB sample within auction ties between bids are frequent: in 209 AB auctions at least two bids are identical, for a total of 720 couples and 38 triplets. With.SD is the within-auction standard deviation of bids. No.Bids is the number of bids. Res.Price is the auction reserve price (see Decarolis, 2009 for a discussion of how the reserve price is set). The bottom left panel reports the same statistics for the FP auctions. The values for HighBid are not reported because in 90 percent of the FP auctions the highest discount coincides with the winning discount. In the remaining cases the highest bid is reputed not credible and eliminated.

Right panel: statistics by firm. The variables reported are the number of auctions attended (Entry), the number of victories (No.Win), the probability of winning in the sample (Pr.Win), the total revenues earned (Reven), the age (Age, measured in years) and the capital (Capital, measured in 2005), the number of subcontracts received (Subct, available only for auctions held in the Piedmont between 2000 and 2007), the miles between the firm and the work (Miles), whether the shuts down between 2005 and 2010 (Closed) and whether it is located in the same five regions in the North where also the auctions were held (North5, these regions are), in other northern or central regions (N. & C.) or in the southern regions or the islands (S. & I.). Revenues and capital are in thousands of Euro.

Table 3: Summary Statistics - Turin's Cartels Sample

	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
W-2Bid	.09	.23	.05	0.0	2.9	276	No.Joint	3.0	4.8	1.0	0.0	24	276
	Statistics of Independent Firms							Statistics of Cartels Firms					
Entry	17.2	22.3	9.0	1.0	186	717	Entry	82.9	71.1	54	1.0	263	95
Wins	.13	.42	0.0	0.0	3	717	Wins	1.9	3.1	1.0	0.0	19	95
Reven	51.8	19.6	0.0	0.0	2319	717	Reven	822	1466	327	0.0	$1e^{04}$	95
Miles	237	284	101	0.0	1071	504	Miles	101	207	15	0.0	991	86
Age	27.1	14	25	2.0	106	559	Age	29.6	14.1	30	1.0	72	91
Subct	1.8	5.0	0.0	0.0	53	717	Subct	6.8	8.6	4.0	0.0	44	95

Notes: The variables used to describe the auctions are the same of those in Table 1. The only additional variable is No.Joint which measures the number of (legal) bidding consortia present in the auction. Each consortium places one single bid. The type of jobs and the reserve price of contracts is similar to those in Table 1. The set of 717 independent firms contains 24 firms that share part of their owners and managers with the cartels" firms. Their presence makes the summary statistics of the independent firms slightly closer to those of the cartels. The missing values for miles and age are due to the impossibility of identifying with certainty some firms.

Table 4: Tests for the Known Cartels

		Bid Test (Single Auction)			Participation Test				
Cartel	# Members	Uncond	Cond1	Cond2	N^g	f^g	$Max\ N^s$	$P^T_{.95}$	$Reject$
1	17	(1,0,1)	1	0	16	4	13	9	1
2	13	(1,0,1)	0	0	12	4	10	8	1
3	16	(1,0,1)	0	1	13	11	7	10	1
4	11	(0,1,1)	0	1	8	26	6	17	1
5	11	(1,1,1)	0	0	8	3	7	11	1
6	6	(0,0,0)	0	0	6	5	5	9	1
7	7	(1,0,1)	0	1	7	10	6	18	1
8	14	(0,1,1)	0	1	9	1	6	1	1

Notes: The columns for the bid test report a 1 if at least 30 percent of the p-values of the single-auction bid test are either below .25 or above .975. A zero is reporthed otherwise. The column labeled Uncond reports the results of the single-auction bid test performed without using firms' covariates to construct the random groups. The column Cond1 reports the results obtianed using random groups that match the suspect group in terms of the (quantile of the distribution of) firms' capital and distance from Turin. The column Cond2 reports the results when the matching is done using the legal requirements for participation of the firms in the suspect group.The columns for the participation test report a 1 when we reject indepenedence in favor of coordinated entry at the 5 percent level. The latter two columns indicate the fraction of members of the suspect group for which, using the by-firm version of the bid test, we can reject independence at the 5 percent level.

Table 5: Multi Auction Bid Test - Cartel 1								
	Non Suspects				Mixture Sample			
Auctions	10th	50th	90th	No. Groups	10th	50th	90th	No. Groups
2	1	1	.06	17	1	1	.05	21
3	1	.25	.06	15	1	.61	.06	21
4	1	.19	0	15	1	.14	.05	21
5	1	.19	0	15	1	.14	.05	21
6	1	.19	0	15	1	.14	.05	21
7	1	.19	0	15	1	.14	.07	21
8	1	.16	0	15	1	.12	.08	21
9	1	.13	0	15	.64	.12	.08	20
10	.38	.15	0	12	.16	.09	.05	18
11	.36	.12	0	10	.16	0	0	15
12	.20	.14	0	9	.13	0	0	15
13	.25	.17	0	7	.13	0	0	14
14	.25	.18	0	7	.13	0	0	14
15	.25	.18	0	7	.07	0	0	13

Notes: The table reports the result of the multi auction bid test performed for a subgroup of cartel 1 for T auctions, with T ranging from 2 to 15. The subgroup consists in the 4 members of cartel 1 participating most often jointly. The left panel of the table reports the results using control groups that are composed exclusively of firms that neither were suspected of collusion nor that share any owner in common with any suspect firm. The right panel reports the results in which all firms, except the 4 members of cartel 1 used, are potential members of the control groups. In both cases, the control group firms were selected to match the location and the legal constraints of the 4 members of cartel 1 used. For each auction, the table reports both the number of control groups. The multiunit bid test procedure is repeated 1000 times of each T. We report the result for the one-sided left bid test: the 10th, 50th and 90th percentile of the distribution of the p-value of the test is reported.

Table 6: Marginal Effects Probit Regression

	Fixed at:	
	Zero	One
Personal	.00056	.068
	(.0016)	(.041)
Subcontract	.00056	.95
	(.0016)	(.019)
Joint-Bidding1	.00056	.049
	(.0016)	(.033)
Joint-Bidding2	.00056	.16
	(.0016)	(.037)

Notes: Validation dataset. The marginal effects are reported for each variable holding fixed the other at their means. The estimated coefficients come from a probit model with dependent variable equal to one if the couple of firms is in the same group and zero otherwise. The dependent variables are: Link1 - Firms share at least one owner (regardless of the shares owned) or a manager (regardless of his exact role); Link2 (Personal) - Firms share at least one CEO or other top management, Link3 - Firms share the majority shareholder; Link4 - Firms headquarters are located in the same zip code area; Link5 - Firms headquarters are located in the same municipality; Link6 - Firms headquarters are located in the same county; Link7 (Subcontract) - Firms entered in a subcontracting relationship at least once; Link8 (Joint–Bidding1) - Firms won at least one auction as a legal consortium among all those auctioned in Piedmont (i.e., the auctions in the AVCP dataset for which only the winner is disclosed); Link9 (Joint–Bidding2) - Firms bid at least once as a legal consortium across all the auctions for which all data on bids and identities are available (i.e., the 250 auctions of the Municipality of Turin). We drop all couple that are linked only by the location links (Link4,5,6). The specification also uses interactions between Link1 and Link4, 7, 8, 9; Link4 and Link7, 8, 9; Link7 and Link8, 9; and, finally, Link8 and Link9.

Table 7: Constructing Groups (Validation Sample)

Panel A: Clutering Method

Group	Composition	Part Test	Bid Test	Group	Composition	Part Test	Bid Test
1	(3,0,2)	(27,0)	(0,1,1)	9	(2,0,0)	(71,60)	(0,1,1)
2	(3,0,0)	(21,15)	(0,1,1)	10	(1,0,1)	(2,2)	(0,1,1)
3	(1,0,2)	(48,4)	(1,1,1)	11	(1,0,1)	(10,0)	(0,0,0)
4	(3,0,0)	(17,0)	(1,0,1)	12	(2,0,0)	(33,33)	(1,1,1)
5	(1,0,1)	(42,12)	(1,1,1)	13	(2,0,0)	(1,0)	(1,1,1)
6	(4,0,3)	(44,0)	(1,0,1)	14	(2,0,1)	(23,8)	(1,0,1)
7	(10,0,0)	(4,0)	(1,0,1)	15	(1,0,1)	(11,0)	(1,1,1)
8	(2,5,2)	(35,0)	(1,0,1)				

Panel B: Joint Participation Method

Group	Composition	Bid Test	Group	Composition	Bid Test
1	(4,0,0)	(1,0,0)	9	(4,0,0)	(1,0,0)
2	(4,0,0)	(1,0,0)	10	(2,2,0)	(0,0,0)
3	(4,0,0)	(1,0,0)	11	(2,2,0)	(0,0,0)
4	(3,1,0)	(1,0,0)	12	(4,0,0)	(1,0,0)
5	(2,2,0)	(0,0,0)	13	(3,1,0)	(0,0,0)
6	(3,1,0)	(1,0,0)	14	(3,1,0)	(1,0,0)
7	(4,0,0)	(1,0,0)	15	(4,0,0)	(1,0,0)
8	(2,2,0)	(1,0,0)			

Notes: The groups are constructed applying to the validation data the clustering and joint participation algorithms described in the txt and in the Web Appendix. PANEL A: Pairs of firms having a predicted probability of being together less than 30 percent are discharged, resulting in a sample of 57 firms (43 true cartel members and 14 independents). The clustering algorithm produces 15 groups (using a .995 cutoff for the maximum tolerated dissimilarity). The table reports each group's composition as a triplet: the first value is the number of firms in the group belonging to the same cartel (for the cartel that contributed the most to this group), the second value is the number of firms from other cartels and the third value is the number of innocent firms. The group' size is the sum of the three values. The tiplet for the bid test reports the one-sided left, one-sided right and two sided tests: a 1 is recorded if at least 30 percent of the auctions lead to reject the single-auction (unconditional) test at 5 percent significance. A zero is reporthed otherwise. For the participation test, the fist value reported is the number of auctions in which all the members of the group bid while the second value is the 95th percentile of the distribution of the number of auctions jointly participated by all members of the control groups. Control groups are constructed conditioning on both location and the legal qualifications of firms. PANEL B: We impose a group size of 4 and we report the groups constructed around the first 15 firms that won most auctions. To each firm the algorithm associates the 3 firms that jointly maximize the participation of the group. Analogously to Panel A, the table reports the composition of the group as well as the results of the bid test.

Table 8: Auctions Affected by Groups - PRELIMINARY!

	Piedmont			Lombardy	Liguria	Emilia	Veneto
	Significance level:						
Bid Test	1%	5%	10%				
Affected auctions	21%	48%	52%				
Entry Test							
Affected auctions							

Table 9: Revenues Effects- Naive Estimates for Turin's Data

Panel A: Auctioneer' Cost of Procurement

		True	Counterf.	Difference
Total cost		105,938		
Competition (Zero-bids)			129,346	-23,408
Scenario1	Cartel1		106,580	-642
	Cartel5		105,702	236
Scenario2	Cartel1		105,502	436
	Cartel5		105,859	79
Scenario3	Cartel1		106,164	-226
	Cartel5		105,760	178

Panel B: Independent Firms' Revenues

	True	Competition	Comp. No Shills
Mean	62	109	126
SD	227	147	172
P10	0	9	10
P50	0	50	56
P90	201	314	358

Notes: All values are in thousands of euro. Panel A reports the difference between the realized cost of procurement for the city of Turin for the contracts in our validation sample and the counterfactual costs under different scenarios. The first column reports the true cost of the contracts awarded in the validation dataset (i.e., the sum of all the winning prices). Counterfactual values for the total cost of procurement are reported for four cases: Competition, the cost equals the sum of all the reserve prices; Scenario 1, none of the members of Cartel# bids and all other firms keep their bids unchanged; Scenario 2, all the bids of the members of Cartel# equal 18 percent and none of the other firms' bids is changed; Scenario 3, two thirds of the bids of the firms in Cartel# are equal to 18 percent and the remaing one third are eliminated, no other bid is changed.

Panel B: The first column reports the distribution of the revenues accruing to the independent firms within the validation sample. The second column reports the distribution of revenues under the assumption that for each auction that a firm participates its revenues equal 1/N of the reserve price (where N is the actual number of participants in that auction). The last column is identical to the second one with the only difference that the fraction of revenues accruing to the firm equals 1/N' of the reserve price (where N' is the actual number of bidders minus 1/3 of the number of cartels' bids). This latter scenario could mimick counterfactual revenues under the zero-bids equilibrium if colluded firms could not uses shills and entry into each auction is unchanged.

TABLE 10: Number of Bidders Regressions				
	Turin Area 2000-2007		North Regions 2005-2010	
	NEG.BIN	Pred.Change	NEG.BIN	Pred.Change
First Price	-1.84	-38.32	-1.87	-44.03
	(.15)***		(.18)***	
Observations	2,548		956	
P-Value Chi^2	.000		.000	

Notes: Significance level * is 10%; ** is 5%; *** is 1%. Standard errors are clustered by administration and year. Pred.Change is the predicted discrete change of the number of bidders due to FP switching from 0 to 1. The negative bionmial model estimated includes in addition to the FP auction dummy the following controls: Log(contract value) and dummy variables for type and geographical location of the PA included. Data: the data used for the Turin Area regression is coming from the "Authority sample" of Decarolis (2010). It contains the public procurement auctions for both Turin and similar neighboring PA for both the years beofore and after the 2003 transition of Turin to the FP. The North Regions data consists in the combination of the AB and FP auction samples described in Table 1.