American Economic Journal: Microeconomics 2016, 8(2): 1–38
http://dx.doi.org/10.1257/mic.20130254

Detecting Bidders Groups in Collusive Auctions†

By Timothy G. Conley and Francesco Decarolis*

We study entry and bidding in procurement auctions where contracts are awarded to the bid closest to a trimmed average bid. These auctions, common in public procurement, create incentives to coordinate bids to manipulate the bid distribution. We present statistical tests to detect coordinated entry and bidding choices. The tests perform well in a validation dataset where a court case makes coordination observable. We use the tests to detect coordination in a larger dataset where it is suspected, but not known. The results are used to interpret a major market shakeout following a switch to first price auctions. (*JEL* D44, D47, H57, R42)

> *…At the first meeting they said: 'Why should we kill ourselves and let those coming from the outside laugh at us?' Here [in Turin] firms from the South were coming and getting the jobs, setting the averages, they used to come with 20, 30 or 40 bids, they used to get the jobs and then what was left for us?…*
>
> —Confession of Bruno Bresciani, found guilty of having rigged 94 average bid auctions and other related crimes; sentenced to seven years of jail in 2008

In recent years, economists have helped design new auction markets for activities ranging from electricity supply contracts to the sale of spectrum licenses to mobile operators. The extent to which these auctions can deliver the intended results depends crucially on how bidders respond to strategic incentives. We study the case of a market in which average bid auctions (ABAs) are used for the procurement of public works and show the sophisticated response of bidders to the incentive to use multiple bids to pilot the contract awarding. We introduce two statistical tests that work well to detect this type of behavior and use them to study the drop in firm entry following a switch to first price auctions (FPAs).

* Conley: Department of Economics, Western University, Social Sciences Center, London, Ontario N6A 5C2, Canada (e-mail: tconley3@uwo.ca); Decarolis, Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02115 (e-mail: fdc@bu.edu). The authors would like to thank their colleagues at their institutions as well as Susan Athey, Lanier Benkard, Marianne Bertrand, Peter Cramton, Jeremy Fox, Ken Hendricks, Aldo Sandoval Hernandez, Ali Hortacsu, Jakub Kastl, Jonathan Levin, Greg Lewis, Jozsef Molnar, Nicola Persico, Gustavo Piga, Phil Reny, David Rivers, Marc Rysman, Pierluigi Sabbatini, Jesse Shapiro, Giancarlo Spagnolo, Eric Yao, and Charles Zheng for the useful suggestions. Conley is grateful to the Social Science Research Council of Canada for financial support. We thankfully acknowledge the aid with the data collection of the Bank of Italy, the Italian Authority for Public Contracts, and the Legal Office of the municipality of Turin as well as the Social Science Research Council of Canada grant 410-2011-1740 and 435-2015-1847.
† Go to http://dx.doi.org/10.1257/mic.20130254 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

ABAs are frequently used to award contracts for public works. They are characterized by the fact that, contrary to FPAs where the lowest price wins, the winner in an ABA is decided through an algorithm eliminating all bids that are deemed too good to be true. This often involves eliminating all prices below the average price, or some other mechanically calculated threshold. Their use is not surprising if one considers that, while private buyers typically have the discretion to exclude unreliable bids, corruption concerns usually imply a preference for rigid and transparent public procurement procedures. Institutions like qualification criteria or bid guarantees can mitigate risks associated with awarding contracts at the lowest price, though implementing them may be costly. ABAs are seen as a simple fix to the problem that can be effective even without qualification criteria or bid guarantees and this explains their frequent use in procurement. They are known to be an important procurement mechanism in China, Japan, Italy, and parts of the United States.[1] Despite their wide use, however, little is known about how ABAs affect bidders' behavior.

This paper is an empirical study of bidders behavior in ABAs. We focus on the case of Italy which is emblematic of the relevance of ABAs. For the period from 2000 to 2010 that we study, ABAs were the primary public procurement format. Every year, contracts totaling 10 billion euro were procured via ABAs. Moreover, regulatory reforms trying to replace ABAs with FPAs experienced limited success: an attempted introduction of FPAs in 2008 failed, resulting in the reintroduction of ABAs in 2011. There is limited theoretical knowledge of bidding under ABAs. Thus, our analysis begins with the description of a baseline theoretical framework where all firms can offer at most one bid. We show that, under the specific rules of the Italian ABA, there is a unique equilibrium where all firms offer a price equal to the reserve price. This equilibrium, however, has an obvious weakness. A bidder would prefer to deviate and submit multiple bids in order to pilot the contract awarding, and therefore impact the thresholds that determine the winning bid. Though it will entail winning at a worse price, expected payoffs will typically increase due to the higher probability of winning. We exploit the ABA rules to design a test that is good at detecting threshold-piloting sets of bids and the associated patterns in joint entry among bidders engaged in such a coordinating strategy.

Bid coordination can be achieved either by a single firm controlling multiple bids or by different firms forming a cartel. Since regulations allow each firm to submit at most one bid, offering multiple bids requires firms to game the system by creating shadow subsidiaries, which we will refer to as shills. Creating a shill is complicated by the fact that, to bid in public auctions, these shills must have different majority shareholders than their parent companies and must satisfy certain financial and technical requirements. Instead of bearing the costs of creating shills, a bidder might also try to coordinate with other firms, effectively forming a bidding ring. Both types of behavior are present in our data, and we focus on patterns in entry and bidding that would be viable strategies for either a firm with shills or a bidding ring of actual firms, or a combination of both shills and actual firms.

[1] Specifically the Florida Department of Transportation and the New York State Procurement Agency. ABAs are also used in many other countries including Chile, Colombia, Peru, Malaysia, Switzerland, and Taiwan.

A main contribution of this paper is to operationalize the detection of bidding groups with unusually coordinated actions through two statistical tests of coordination: one for entry and one for bidding. Both tests are based on randomization inference methods. In particular, our entry test compares the frequency of joint entry for members in a suspect group with a set of comparable bidders in terms of various determinants of entry. Fundamental to the effectiveness of this test is the rich set of observed firm covariates and the presence of a complementary dataset of FPAs. This dataset allows us to assess which firm characteristics are relevant determinants of entry (and bidding). Conditioning on these characteristics when constructing the set of comparable bidders groups reduces the chance that a group is labeled unusually coordinating simply due to having unusual costs. Similarly, for our bid test we exploit the exact rules of these ABAs to construct a test statistic tailored to measure the extent to which a given group's bids move the threshold that determines a winner. We then compare this measure of "mean piloting" for the suspect group to its analog for a set of comparable groups. The large number of ABAs in the data allows us to exploit the repeated observation of a group's behavior across multiple auctions to more reliably distinguish coordinated bidding from spurious bid correlation.

When we apply our tests to known coordinating groups, they perform well in detecting them. We use 276 ABAs for roadworks held by the city of Turin between 2000 and 2003. We refer to these auctions as the Validation data. In 2008, the Turin Court of Justice ruled that these auctions had been rigged by 8 cartels made up of 95 firms. From our perspective, this is an ideal scenario because the case was based on very detailed evidence, including confessions of some ring members and phone calls intercepted by the police. Thus, we can reasonably consider that for these ABAs we know the identities of cartel members who engaged in the types of coordinating actions we want to detect. We examine our tests' performance by checking whether they are able to detect coordinated actions of these eight known cartels. The results strongly support the capability of our tests to correctly detect coordinating groups. Of the eight cartels, the only one for which we do not find systematic evidence of coordination received lighter sanctions from the court because its members rarely coordinated bids.

The judicial case also reveals the presence of competition among cartels. Police found hard evidence of a lack of coordination between cartels. Thus, bid coordination in ABAs is not the textbook instance of collusion, but rather a form of competition involving multiple groups and, likely, noncoordinating firms.

We then turn to the problem of detecting coordinating groups in auctions where we have no prior knowledge of their presence. We look at a dataset of 802 ABAs held in the North of Italy between 2005 and 2010. We refer to these auctions as the Main data. Many of the observed features of these ABAs resemble those of the ABAs in the Validation data. Although our tests could be applied to any candidate group, given the large number of firms in these auctions, we suggest various ways to reduce the set of firms to analyze. Our favorite method constructs candidate groups starting from the network of relationships connecting firms along various observable dimensions: overlaps in the identities of owners and managers, the exchange of subcontracts, the formation of temporary bidding consortia and geographical proximity. We then apply our tests to these constructed groups. Based on these tests, we detect numerous groups of firms that appear to be engaged in bid and entry coordination. Our conservative estimates suggest that these groups affect no less than 30 percent of the auctions. We then argue that bid coordination likely produced large savings for the auctioneer relative to the baseline of competitive ABAs with one bid per firm. This is because the counterfactual equilibrium for this case entails all firms offering a discount of zero, while the observed winning discounts are substantially higher: 13 percent of the reserve price on average. However, firms outside the coordinating groups are harmed. They are less likely to win and when they do, they get a worse price than under competition.

In the final part of the analysis, we revisit the role of coordinating groups in the context of a policy change that replaced ABAs with FPAs for certain types of contracts. Our main focus is on the exit of hundreds of firms from the market that followed this switch. We argue that this market shakeout is due to the exit of both inefficient firms unable to effectively compete in FPAs and shill firms that became useless under FPAs. Understanding the relative frequency of these two motivations for exit is key to evaluate policy interventions like the introduction of subsidies for weak bidders in FPAs. We investigate whether a classification into coordinating groups based on our tests can be useful to assess the frequency of shills among exiting firms. Our findings suggest that, among the 774 exiting firms, 159 of them (or 21 percent) belong to detected coordinating groups. We show that exiting firms belonging to detected coordinating groups display characteristics consistent with being shill firms.

Overall, our results have important implications for auction and market design because they describe the degree of bidding sophistication observed in a major procurement auction setting. Moreover, this is the first empirical investigation of entry and bidding into a widely used auction format. Our results highlight the troubling fact that ABAs incentivize the sidestepping of the regulations. From a policy perspective, our results contribute to the debate that followed the market shakeout observed once FPAs replaced ABAs and that represented a major obstacle to the consolidation of FPAs. Our results allow us to understand why the reversal of this reform in 2011 has lead to the resurgence of patterns in the data consistent with the presence of coordinating groups.

I. Literature

This study is related to the vast literature on collusion in auctions.[2] Methodologically, our empirical approach is related to the two major strands in which the empirical literature can be divided: the studies of collusion practices in markets where the presence of cartels has been proved by a court (Asker 2010; Pesendorfer 2000; Porter and Zona 1993; and Porter and Zona 1999) and the studies that try to devise methods to distinguish competition from collusion when collusion is only a possibility (Bajari and Ye 2003).[3] Both approaches have led to the

[2] For a review of the theoretical literature on collusion in auctions see Marshall and Marx (2012).

[3] See also Haberbush (2000) for a review of cases of collusion in US public procurement auctions. Porter and Zona (1993) and Ishii (2009) specifically analyze collusion in auctions for roadwork contracts.

flourishing of a literature on screens for collusion (i.e., statistical tests to detect collusion, see Abrantes-Metz and Bajari 2012). 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 us to assess the likelihood of coordinating groups in markets where its presence has not yet been proved. Thus, our approach exploits the idea of Hendricks and Porter (1989) that collusion is intrinsically tailored to the specific rules of the environment where it takes place. Finally, it is relevant to stress that the type of bid coordination that we study is beneficial for the auctioneer. That certain types of collusion mechanisms can sometimes lead to improvements in the auctioneer's revenues had been found by Asker (2010) in a small subset of the rigged auctions that he studies. More generally, Harrington (2012) analyzes cases wherein tacit price coordination, though possibly qualifying as an infringement of antitrust laws, does not lead to collusive prices. Our results, however, are quite different in that we argue that the presence of groups improves competition relative to a benchmark where all firms are independent, if such firms can only offer a single bid.

This paper is also related to the literature on public procurement surveyed in Dimitri, Piga, and Spagnolo (2006). In particular, our discussion of ABAs and FPAs in Section VI is related to studies of bidding behavior in different procurement auction formats. Recent examples include the analysis of bidding behavior in US forest timber auctions (Athey and Levin 2001) and in Minnesota highway construction contracts (Lewis and Bajari 2011). Our study also complements earlier theoretical (Spagnolo, Albano, and Bianchi 2006; Engel et al. 2006; and Decarolis 2013) and empirical (Decarolis 2014) studies on the role of ABAs in procurement. Relative to all these papers, our study is the first to empirically quantify the presence of coordinated bidder groups in ABAs.

Finally, this paper is related to studies analyzing mechanisms similar to the ABAs. For instance, Abrantes-Metz et al. (2012) study the case of the LIBOR. This rate, to which contracts worth \$300 trillion are linked, is a trimmed mean of bank quotes for interest rates. Evidence that several banks coordinated their quotes to manipulate this trimmed mean emerged in 2012. Morton (1997) and Duggan and Morton (2006) study how drug manufacturers distort prices in response to a regulation setting the mandatory rebate for Medicaid as an average of the drug prices faced by non-Medicaid enrollees. For Medicare Part D, Decarolis (2015) studies how insurers use the multiple plans that they offer to increase the subsidy paid by Medicare which, in turn, is a function of the average of plan premiums.

II. Description of the Market

In this section, we describe both the institutions and our datasets. We study auctions held between 2000 and 2010 by Italian public administrations (PAs) to procure contracts for simple roadworks in Northern Italy. We are motivated to study these auctions because, for the municipality of Turin, we have access to what we call Validation data as a result of legal cases where several firms were convicted for collusion in these auctions. These data are comparable to the remainder of our data, which we refer to as our Main data.

For these contracts, PAs are typically required to select the contractor through sealed bid price-based auctions, either the well-known FPAs or ABAs. In both cases, the PA announces a job description and a reserve price that is the maximum it is willing to pay. Then firms submit sealed bids consisting of discounts on this reserve price. However, while in FPAs the highest discount wins, in ABAs the winner is determined as follows: (i) bids are ranked from the lowest to the highest discount; (ii) a trimmed mean ($A1$) is calculated disregarding the 10 percent of the highest and lowest discounts; (iii) a new mean ($A2$) is calculated as the average of those discounts strictly above $A1$, disregarding those discounts excluded in the calculation of $A1$; (iv) the winning discount is the highest discount strictly lower than $A2$. Ties of winning discounts are broken with a fair lottery.[4]

To better understand the working of the ABA algorithm, consider the following example. Suppose that there are 50 bidders (roughly the sample average in the data). Bidders are numbered $1, 2, \ldots, 50$ and assume $b_1 = 1\%$, $b_2 = 2\%, \ldots$, $b_{50} = 50\%$. The algorithm starts by disregarding bids in the bottom 10 percent (i.e., $b_1, \ldots, b_5$) and top 10 percent (i.e., $b_{46}, \ldots, b_{50}$) of the bid distribution. $A1$ is then equal to 25.5 percent, the simple average of the remaining bids (i.e., $b_6, \ldots, b_{45}$). $A2$ equals 35.5 percent, the simple average of the nondisregarded bids above $A1$ (i.e., $b_{26}, \ldots, b_{45}$). The winner is thus bidder 35 who has the highest bid strictly below $A2$. He will be paid 65 percent of the reserve price to perform the job. If the same bids were submitted in an FPA, bidder 50 would win and would get paid 50 percent of the reserve price to complete the job.

The ABA described above was introduced in 1999 and, until June 2006, it was the compulsory mechanism for the procurement of almost all contracts with a reserve price below €5 million. In this period, 80 percent of all contracts for public works (worth €10 billion a year) were awarded using ABAs. FPAs were the typical format for contracts with a reserve price of €5 million or more. Between July 2006 and May 2011, a series of reforms required by the European Union temporarily limited the use of ABAs and extended the use of the FPAs for contracts below €5 million. Since May 2011, however, ABAs have once again been allowed for all contracts below €5 million and are currently widely used.

A. *Main Data*

Our Main data contain 1,034 auctions held by counties and municipalities between November 2005 and May 2010. Contracts involved the procurement of simple roadwork jobs (mostly paving jobs, worth below €1 million) and were held in five Northern regions (Piedmont, Liguria, Lombardy, Veneto, and Emilia-Romagna).

The data consist of 802 ABAs and 232 FPAs. Table 1 presents summary statistics separately for the two types of auctions. Comparing the statistics for the two sets of

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

Table 1—Summary Statistics—Main Data

	Mean	SD	Median	Minimum	Maximum	Observations
Panel A. Statistics by auction–ABAs						
HighBid	17.4	5.4	17.4	1.6	37.4	802
WinBid	13.4	5.2	13.5	0.51	36.8	802
$\Delta W2_{nd}$	0.24	0.68	0.07	0	9.4	802
With.SD	2.9	1.4	2.7	0.14	9.2	802
No.Bids	50.7	34.3	43	5	253	802
Res.Price	312	204	250	11	999	802
Panel B. Statistics by auction–FPAs						
WinBid	28.9	9.9	29	1.2	53.4	232
$\Delta W2_{nd}$	5.1	5.5	3.3	0.01	41	232
With.SD	6.6	3.3	6.1	0.07	19.1	232
No.Bids	7.3	5.5	6	2	48	232
Res.Price	342	288	215	30	978	232
Panel C. Statistics by firm						
Entry	13.1	22.1	4	1	205	4,005
Wins	0.31	0.87	0	0	18	4,005
Pr.Win	0.03	0.12	0	0	1	4,005
Reven	170	1,081	0	0	$4e^{04}$	4,005
Miles	159	234	47.8	0	1,102	4,005
Age	22.3	13.8	21	1	106	3,611
Capital	447	2,411	52	10	$8e^{04}$	2,484
Subct	0.65	2.9	0	0	53	4,005

Notes: Panels A and B: Statistics for ABAs and FPAs for roadwork contracts procured by the municipalities of five Northern regions: Piedmont, Liguria, Lombardia, Veneto, Emilia-Romagna. Panel A: statistics by auction for the sample of ABAs. HighBid is the highest discount. WinBid is the winning discount. $\Delta W2_{nd}$ is the difference between the winning bid and the bid immediately below it. With.SD is the within-auction standard deviation of bids. No.Bids is the number of bids. Res.Price is the auction reserve price in thousands of euro. Panel B reports the same statistics for FPAs. The HighBid is (almost) always WinBid and so is not reported. Panel C: Statistics by firm. The variables reported are the number of auctions attended (Entry), the number of victories (No.Win), the probability of winning in the sample (Pr.Win), the total revenues earned (Reven), the age (Age, measured in years in 2010) and the capital (Capital, measured in 2005), the number of subcontracts received (Subct), the miles between the firm and the work (Miles). Revenues and capital are in thousands of euro.

auctions reveals several differences in terms of firms' entry and bidding. Regarding entry, the number of bidders is several times larger in ABAs than in FPAs: on average there are 7 bidders in an FPA and 51 in an ABA. Regarding bidding, the winning discount is on average 13 percent in an ABA, while it is 29 percent in an FPA. Moreover, in ABAs there is substantially less within-auction variation in the bids than in the FPAs: this is shown by both the lower within-auction standard deviation of bids and the lower difference between the winning discount and the next highest discount in the ABAs relative to the FPAs. This latter variable, sometimes defined as "money left on the table" is on average 5.1 percent of the reserve price in FPAs, but only 0.2 percent in ABAs. Finally, the bottom panel of Table 1 reports summary statistics for the bidders. There are approximately 4,000 firms that bid at least once. They exhibit strong asymmetries both in their characteristics and in their performance in the auctions. Although we do not report the data broken down by the format in which the firms participate, on average the firms bidding in FPAs have higher capital and are located closer to the work area.

Descriptive Analysis of the Main Data.—We analyze separately correlations for entry and bidding under the two formats. For entry, we estimate probit regressions where the dependent variable is one if the firm bids in the auction and zero if the firm does not bid but is a potential participant.[5] For bidding, we use all the bids in each auction and estimate OLS models where the dependent variable is the discount offered. For all regressions, the set of independent variables includes both firm and auction characteristics. The former are particularly relevant as we will control for firm characteristics in our tests. The firm variables are the log distance between the firm and the work site, the log capital, backlog,[6] the workforce size, and dummy variables for limited liability status and for the firm region. We also report estimates that include additional variables recording for each firm how many of the other bidders are connected to it through a series of linkages that we will later show to matter for predicting coordinating groups.[7] The auction variables are six dummies for value categories of the reserve price and dummies for each year and PA region. We also consider specifications with auction fixed effects.

Table 2 reports the entry and bidding regression estimates. ABA and FPA entry estimates are similar: The coefficients on all firm covariates (with the exception of the number of workers) are identical in terms of sign and significance. Magnitudes are also nearly the same for two particularly relevant cost proxies: distance and capitalization. ABA and FPA bid estimates are, instead, not similar to each other: Most coefficients differ in terms of sign and/or significance. Consistent with the findings in the literature, we obtain that for FPAs a greater distance to the work site is associated with smaller discounts. For ABAs, however, distance is disconnected from discounts (or associated with them with the "wrong sign"). Furthermore, absent auction fixed effects, the R^2 is only 13 percent in the case of ABAs, while it is 21 percent in the case of FPAs. The inclusion of auction fixed effects reverses this order with the R^2 increasing to 66 percent in the case of ABAs and 55 percent in the case of FPAs. Both the disconnection between bids and costs and the relevance of auction fixed effects are well explained by what we describe below about firm behavior in the ABAs of the Validation data.

Altogether, the evidence from the two datasets will make clear that ABA bids are of little help to infer underlying firm costs. However, given that for our method it is crucial to disentangle coordinated behavior from the presence of common cost shocks, we will proceed under the assumption that costs relevant for ABA and FPA bids are the same. Hence, those cost proxies that appear as significantly associated with FPA bids will be our cost-controls in the bid tests for ABAs. Furthermore, since the estimates show that the same firm covariates associated with FPA bids are also associated with ABA entry, this justifies a step of our method employing ABA entry choices as an approximate sufficient statistic for bid-relevant costs.

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

[6] The backlog variable is calculated as in Jofre-Bonet and Pesendorfer (2003) and measures the amount of unfinished work across the stock of contracts won at the time of the entry/bid decision.

[7] These links are common personnel, common owner, common manager, common zip code, common municipality, common county, subcontracts, winning consortium, and bidding consortium.

Table 2—Probability of Entry and the Discount Offered—Main Data

	Probability of entry				Discount offered			
	FPA		ABA		FPA		ABA	
	Probit (1)	Probit (2)	Probit (3)	Probit (4)	OLS (5)	OLS (6)	OLS (7)	OLS (8)
log(miles firm-work)	−0.84***	−0.85***	−0.86***	−0.86***	−0.66**	−0.40**	0.23*	0.04
	(0.02)	(0.02)	(0.01)	(0.01)	(0.23)	(0.15)	(0.13)	(0.07)
log(firm capital)	0.06***	0.06***	0.05***	0.05***	−0.46**	−0.09	−0.10**	−0.03*
	(0.01)	(0.01)	(0.01)	(0.01)	(0.21)	(0.17)	(0.04)	(0.01)
Backlog	0.12	0.13	0.07	0.07	−2.79	−0.34	0.09	0.19
	(0.21)	(0.21)	(0.05)	(0.05)	(3.94)	(3.64)	(0.22)	(0.15)
Unlimited liability	0.45**	0.46**	0.04**	0.05**	−5.30**	−1.81	−0.39***	−0.17*
	(0.19)	(0.15)	(0.02)	(0.02)	(2.01)	(2.43)	(0.12)	(0.09)
Number of workers	0.04	−0.02	−0.66***	−0.67***	14.73	10.10	0.30	1.10*
	(0.55)	(0.55)	(0.17)	(0.17)	(13.21)	(11.93)	(1.23)	(0.54)
Firm links	No	Yes	No	Yes	No	No	No	No
Auction FE	No	No	No	No	No	Yes	No	Yes
Prob. χ^2	0.00	0.00	0.00	0.00	—	—	—	—
R^2	—	—	—	—	0.21	0.55	0.13	0.66
Observations	11,806	11,806	80,274	80,274	1,886	1,886	37,699	37,699

Notes: Sample: Main data. Columns 1–4 report probit regression where the dependent variable is one if the firm bids in the auction and zero if the firm does not bid but is a potential participant. All probit regressions include a constant, six dummies for the categories of value of the reserve price, and dummies for each year, and the PA region and the firm region. Relative to models 1 and 3, models 2 and 4 include "firm link" variables (how many other bidders in the auction are linked to the firm along each one of the links described in Table 6). Columns 5–8 report OLS regressions for the discount offered. Standard errors are clustered by PA and year. Relative to the sample of the probit regressions in columns 1–4, only submitted bids are part of the sample of OLS regressions in columns 5–8. All regressions include a constant, six dummies for the categories of value of the reserve price, and dummies for each year and region of the auction. Relative to models 5 and 7, models 6 and 8 also include auction fixed effects.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Finally, we operationalize the choice of firm characteristics used as cost controls by looking at model fit. Conditioning on too many dimensions would make impossible to find an adequate number of firms to be used as controls in the tests. Thus, we select the subset of the most relevant variables by looking at the marginal change in the R^2 when one variable is excluded from the regression. For both the FPA bid regressions and all the entry regressions, we find that distance and capitalization are the two most relevant control variables.[8]

B. *Validation Data*

The ABAs in the Validation data were collected by the legal office of the municipality of Turin as part of a legal case against several firms accused of committing auction rigging. This dataset consists of 276 ABAs held by the municipality of Turin between 2000 and 2003 to procure roadwork jobs. There is a substantial overlap of

[8] For the entry regressions, we consider the linear probability analogue of the probit model in Table 2.

Table 3—The Eight Sanctioned Cartels in Turin

Cartel name and ID	Firms	Victories	Auctions
1–Torinisti (B)	19	91	270
2–San Mauro (C)	14	39	257
3–Coop (G)	19	22	263
4–Pinerolesi (A)	12	3	125
5–Canavesani (E)	12	7	178
6–Settimo (D)	6	10	245
7–Provvisiero (F)	7	10	74
8–Tartara-Ritonnaro (H)	13	2	109

Notes: This table shows the eight cartels of the Validation data. The first column reports the name of the cartel and, in parentheses, the capital letter that we use to identify the group. The last three columns of the table report the size (i.e., the number of firms) of the cartel, the total number of auctions its members won, and the total number of auctions attended by at least one member of the cartel (out of the 276 auctions of the Validation data).

bidders among the Main and Validation data. In April 2008, the Court of Justice of Turin convicted the owners and managers of numerous construction firms. The court documents identify 95 firms that operated in 8 cartels.[9] We use the term cartels for these eight groups to better distinguish them from the groups of potential coordinators detected by our tests. These cartels were very successful in their activity. Despite representing no more than 10 percent of the firms in the market, they won about 80 percent of all the auctions held in the Piedmont region between 2000 and 2003. Cartels were formed mostly by firms geographically close to each other. This is unsurprising both because some firms are shills and because proximity lowers coordination costs among different cartel members.[10]

In Table 3, we use capital letters, from A to H, to indicate each cartel. The table shows that the eight cartels are quite heterogenous in their size, entry, and victories. Moreover, six of these cartels have all their firms located close to Turin. For the remaining two cartels, G and H, we know from the court case that their strategy typically entailed winning the contract to resell it via subcontracts to other firms closer to the work site.

In addition to the asymmetries across cartels, Table 4 shows that there are also significant asymmetries within cartels, panel B, and between cartel and non-cartel firms, panel C. Given that this sample was assembled for the court case, it is not surprising to see that all variables measuring outcomes of the auctions (entry, victories, subcontracts, etc.) take larger values for the members of the cartels. Regarding the auctions themselves, however, panel A of Table 4 suggests that these auctions are similar to those in the Main data described in Table 1 on the basis of entry and

[9] Turin Court of Justice, first Criminal Section, April 28, 2008, sentence N. 2549/06 R.G. Of the 95 suspect firms, 29 were sentenced. Proscription lead to the acquittal for two firms. The judgment of the other firms was decided in different court cases. In our study we consider the full network of 95 firms.

[10] See Ortner and Chassang (2014) for a recent theoretical contribution on the role of information frictions for the sustainability of collusion. For the type of market that we study, Porter and Zona (1993) suggest various specific reasons, mostly related to information frictions, for why cartels frequently emerge: (i) bids are evaluated only along the price dimension and so product differentiation is absent; (ii) firms are relatively homogeneous because of the similar technology and inputs; (iii) every year there are many auctions and they take place quite regularly; (iv) there are legal forms of joint bidding; (v) the same firms repeatedly interact, (vi) ex post the auctioneer discloses the identities and bids of all bidders.

TABLE 4—SUMMARY STATISTICS–VALIDATION DATA

	Mean	SD	Median	Minimum	Maximum	Observations
Panel A. Statistics by auction–ABAs						
HighBid	22.8	5.6	22.1	12.5	47.5	276
WinBid	17.4	5.0	17.3	6.7	37.7	276
$\Delta W2_{nd}$	0.09	0.23	0.05	0	2.9	276
With.SD	3.6	3.9	1.7	0.34	10	276
No.Bids	73.3	37.1	70	6.0	199	276
Res.Price	510	400	460	50	3,710	276
Panel B. Statistics by firm: Firms in the eight cartels						
Entry	82.9	71.1	54	1.0	263	95
Wins	1.9	3.1	1.0	0	19	95
Reven	822	1,466	327	0	$1e^{04}$	95
Miles	101	207	15	0	991	86
Age	29.6	14.1	30	1.0	72	91
Subct	6.8	8.6	4.0	0	44	95
Panel C. Statistics by firm: Other firms						
Entry	17.2	22.3	9.0	1.0	186	717
Wins	0.13	0.42	0	0	3	717
Reven	51.8	19.6	0	0	2,319	717
Miles	237	284	101	0	1,071	504
Age	27.1	14	25	2.0	106	559
Subct	1.8	5.0	0	0	53	717

Notes: Panel A: Summary statistics by auction. All auctions in the Validation data are ABAs. The definition of the variables is that given in Table 1. The reserve price is expressed in millions of euros. Panels B and C: Summary statistics by firm, distinguishing between the firms in the eight cartels and all the remaining firms. The definition of the variables is again the one given in Table 1.

dispersion of the bids. Interestingly, the average winning discount is higher in these "colluded" auctions than in those of Table 1, 17.4 percent compared to 13.7 percent. This is not totally surprising, however, considering that the court investigation reveals a fierce competition between the 8 cartels to win auctions. Understanding this paradoxical situation of competition between cartels requires a careful analysis of bidding in ABAs.

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

The first feature of the bid distributions is that a basic range for winning discounts is predictable across auctions within a PA. The winning bids are almost always near the approximate mode of the bid distribution, which in the Validation data is around 17 to 18 percent. Court documents report the cases of various defendants claiming that it was known to all players in this market that most of the discounts would be

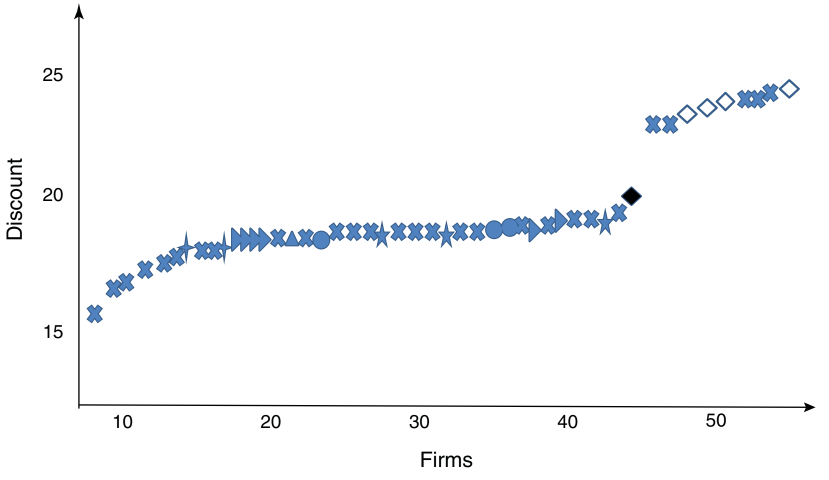

Figure 1. Example of an ABA in the Validation Data

Notes: Discounts offered in one of ABAs in the Validation data. The horizontal axis marks all the 56 bidders that participated in this auction, while the vertical axis reports the discount they offered. The firms are sorted to be in increasing order of the discount offered. Almost all bidders are offered a discount close to 18 percent. The different symbols mark different cartels, but the symbol **x** indicates noncoordinating firms. The diamond symbol, ◇, is used for the bids of the winning cartel: the dark, full diamond represents the winner's bid, while the hollow diamonds represent the nonwinning members of this cartel. The nine highest discounts comply with the description of "supporting bids" presented in the text.

near this range. Figure 1 illustrates this for one Validation data auction. Individual bids are plotted in increasing order with discounts on the vertical axis. There is a clear mode in the distribution around 18 percent with the winning bid highlighted by the black diamond on the edge of this mode. Auctions for this PA within a year of this auction have very similar modes and winning bids. This basic pattern occurs in the ABAs in the Main data as well. For example, both the difference between the winning discount and the next discount and the within-auction standard deviation are similar in Main and Validation data (see Table 1 and Table 4). This evidence about predictability of modes and range of winning bids is confirmed by accounts given by market participants and is consistent with the large amount of public information about past auctions.[11]

The second feature about bidding is that, despite the fact that most bids are typically in a range near the winning discount, there are often some extremely high and/or low discounts. The explanation offered in the court documents is that sometimes bids are not placed to win but to pilot the average. The bidders themselves refer to these very high/low bids as "supporting bids" because they are too extreme to have any chance of winning the auction, but can help a connected firm to win. In Figure 1, the nine highest discounts illustrate the idea of supporting bids. Recall that the vertical axis is the discount offered while the horizontal axis lists the bidders in

[11] The sources of information are both public and private. Regulations require the publication of auction outcomes on the PAs notice board. Moreover, an active market exists for firms reselling information on auctions. Coviello and Mariniello (2014) study the effects of these sources of information on auction outcomes.

increasing order of their discounts. Different symbols indicate different cartels with the thick **x** representing firms not in cartels. The majority of discounts are near the 18 percent approximate mode. However, several members of the cartel, represented by a diamond, submitted discounts that are "discontinuously" greater than those of all other bidders. In this case, their strategy was successful in making a member of their coalition win the auction (the black diamond). In this example, the winning cartel exercises a competitive pressure on the winning discount. However, the association between coordinated bids and competition is not trivial because support bids are also used to push the average discount downward. Interestingly, the Validation data reveals that since multiple cartels were simultaneously trying to manipulate the average, non-cartel firms were sometimes winning by bidding near the 17 percent to 18 percent mode. Many cases of evident average manipulation are present in the Validation data. Correspondingly, numerous extreme discounts suggest piloting of the awarding threshold in the Main data. It is routine for there to be clusters of bids in the tails of the distribution separated by a substantial distance from the bulk of the bids.

The third relevant behavioral feature regards the joint entry of firms. As mentioned in the introduction, it is illegal for two firms sharing the same majority shareholder to submit bids in the same auction. However, the Validation data reveal that entry by closely connected firms is common. Several of the firms composing the eight sanctioned cartels had shareholders in common. Moreover, some of them also shared managers, ownership by members of the same family, registration at the same street address, or systematically exchanged subcontracts. Since we observe all these characteristics for the firms in the Main data, we know that in both datasets it is extremely common to find several closely connected firms entering the same auction. Sometimes the connections between firms in the Validation data were so strong that the sentence mentions the possibility that some firms could have been considered shills of some other firm in the same cartel. However, not even the court could convincingly identify shills because that would require demonstrating that in the absence of ABAs the firm would not exist. One aspect that the sentence clearly sorts out is that no cartel was made only of a parent company and its shills. This is evident from the multiple accounts of conflicts within cartels on revenue sharing that, in a few cases, culminated with firms exiting cartels. In Section VI, we explore this issue in greater detail, but for most of our analysis it will be convenient to think of a coordinating group as a collection of firms acting jointly as if they were all subsidiaries of a mother company.

III. ABA Bidding and Incentives for Coordinating Bids

This section presents a simple framework to discuss both why bidders have an incentive to coordinate bids and what is a likely method to manipulate an ABA through coordinated bids. Use of this method creates observable patterns in firms' bids and entry choices that are the basis of our tests to detect coordinating groups.

We begin by considering a standard characterization of a procurement auction as an independent private value (IPV) game: there are N firms; each firm j has a cost $c_j \in [c^l, c^h]$ that is is privately and independently drawn from the same absolutely continuous distribution $F_C(\cdot)$; each firm simultaneously submits one single bid that must be between zero and one and represents a discount over a publicly announced reserve price R;[12] the expected profit for firm j offering b_j is: $[(1 - b_j)R - c_j] \Pr(b_j \, wins)$. The winner is determined according to the ABA rules described in Section II: b_j wins if it is the highest discount strictly below $A2$.

We can now turn to the ABA equilibrium analysis. One equilibrium entails all firms offering a discount of zero. Provided that $N - 1$ discounts equal zero, an individual firm cannot profitably deviate by offering any other discount. Negative discounts are not allowed, while any positive discount would lead to the certainty of the bid being eliminated. In contrast, this individual firm would have a probability of $1/N$ of winning at a price equal to the reserve price if it keeps at zero the discount offered. This result does not generalize to other strategy profiles where all bidders offer the same nonzero discount. This depends on the details of the ABA rules. The lowest 10 percent of the discounts are disregarded in the calculation of $A1$, but not excluded from the auction (the logic of the rule is to exclude excessively *high* discounts). Hence, if bidder j deviates to $b_j = 0$ when all other bidders bid $b > 0$, then j wins because its bid is the *highest bid strictly below* $A2$.[13]

This logic can be extended to conclude that the zero bid equilibrium is unique. Although this also relies on the fine details of the complex ABA rule, the repeated participation in ABAs together with the high stakes associated with winning, make the zero bid a reasonable outcome in environments where each firm submits one bid: Absent frictions, firms should rapidly learn that high bids lead to both low profits and high chances of being eliminated.[14]

Nevertheless, a major source of frictions in ABAs is the likely presence of coordinated bids. A basic modification of the game consisting of allowing bidders to submit multiple bids readily reveals the fragility of the zero-bid equilibrium characterization. To see this, consider again the prototypical ABA with 50 bidders. Since the average reserve price in the data is about €300 thousand, if the zero-discount equilibrium is played, each firm has an expected revenue of €6,000. Next, suppose that out of these 50 bids 7 are all submitted by the same firm. This firm would clearly find profitable to deviate and submit, for instance, one bid equal to some small discount $\varepsilon > 0$ and the remaining six bids strictly higher than ε. This would ensure that the multi-bid firm wins the auction at a discount of ε. Its expected revenues would thus increase from €42,000 to close to €300 thousand. This example illustrates that whenever there is at least one player who coordinates a sufficiently high number of bids in the auction (and unless all bidders collude), the zero-discount bid profile is not an equilibrium. The size of this "minimum breaking coalition" that makes a deviation from the zero bid strategy profile profitable equals 2 plus 10 percent of N.

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

[13] Recalling the rules about bid ties from Section II, if all bids equal a constant $b > 0$, then $A1 = A2 = b > 0$ and the winner is randomly selected. When j deviates to $b_j = 0$, its bid is disregarded in the calculation of $A1$ and so $A1 = A2 = b > 0$. Thus, j wins with certainty because it has the highest bid strictly below $A2$.

[14] Decarolis (2013) presents a formal proof for various instances of the exact ABA rule. Decarolis and Giorgiantonio (forthcoming) describe instances of variations if the fine details of the ABA (like the tail trimming procedure) at local PA level and show how this resulted in marked shifts in observed bidding patterns.

Bid coordination can result from a single firm having access to shill bidders or from multiple firms agreeing on their bids. The latter case introduces questions of stability of the agreement. However, even neglecting this, the equilibrium characterization of a game where multiple players have access to multiple shills is difficult due to the complexity of the strategy space. Therefore, we focus on one specific type of behavior that is particularly relevant for our environment. Continuing from the previous example, suppose again that 7 bids are all from one firm and that the remaining 43 are all from noncoordinating firms. Suppose that these latter 43 bids are realizations of independent draws from a uniform distribution between 17 percent and 18 percent (this is roughly the typical winning range described in Section II). In this scenario, if bids are in discrete increments of 0.1 percent, then the multi-bid firm can ensure victory at a discount of 18.1 percent if it submits 1 bid equal to 18.1 percent and 6 bids of at least 56 percent. With only 7 bids out of 50, the group cannot manipulate the average downward and win for sure. However, this will typically be possible for a large enough group.[15] In general, a coordinating group could employ either upward or downward manipulations.

Since influencing $A1$ is key to manipulating the awarding rule, we conjecture that a distinguishing feature of coordinated bids is to have a higher influence on the trimmed mean than that of a set of noncoordinated bids. An equilibrium characterization is presented in the online Appendix for the special case of one group and many noncoordinating firms. Although, when multiple groups compete to win the auction a formal equilibrium analysis is hard, intuitively the lack of intergroup coordination can advantage noncoordinating firms. With some groups trying to increase the average relative to the typical winning range, while others try to decrease it, a noncoordinating firm bidding within the typical winning range might end up winning. We conjecture that noncoordinating firms randomize their bids in a small interval around the typical winning range and that this is a best response to the average-piloting strategies of the competing groups of coordinators. These conjectures on group and noncoordinating behaviors capture the features of the Validation data described in Section II. The *support bids* are a clear example of average-piloting strategy, while randomization by noncoordinating bidders is consistent with the large number of bids by noncoordinating firms falling in the narrow interval between 17 and 18 percent.

Our conjectures on bidding behavior motivate a test for whether coordinated bids have an unusual influence on the trimmed mean $A1$. However, at least two caveats limit the applicability of this idea. First, correlation in costs might induce correlation among bids of noncoordinating firms. This makes it important for our testing procedures to always control for cost determinants. Second, information on how a cartel intends to rig a given auction might leak to some noncoordinating firms and induce them to adjust their bid up/down to match those of the coordinating group. To the extent that such events are idiosyncratic across auctions, we might still be

[15] The multi-bid firm cannot guarantee itself the certainty of winning at a discount below 18.1 percent for two reasons. First, the winning bid has to be *above* $A1$ and, so, bids that are too low cannot win. Second, given a range for competing bids of 17 percent to 18 percent, the multi-bid firm can submit extremely high discounts (up to 100 percent) with a greater impact on the average than the lowest bid it can submit, 0 percent. A successful downward manipulation would thus require for the multi-bid firm to further increase the number of its bids.

able to distinguish coordinated bids from noncoordinated bids by looking at bids across multiple auctions.

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

IV. Econometric Tests

A. *Participation Test*

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

We look at whether a statistic reflecting g participation patterns is a tail event relative to a reference distribution that is uniform with points of support given by the analogous statistics for all the groups in H. Define T as the total number of auctions and use the indicator $d_{it} = 1$ to indicate that firm i participates in auction t. Then, for group g having size N^g, the fraction of auctions participated in by $K \leq N^g$ members of g is

$$f_K^g = \sum_{t=1}^{T} \mathbf{1}\left\{K = \sum_{i \in g} d_{it}\right\}.$$

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

$$f_K^h = \sum_{t=1}^{T} \mathbf{1}\left\{K = \sum_{i \in h} d_{it}\right\}.$$

Formally, we test the hypothesis that firms in g do not have unusually coordinated entry by testing that f_K^g is a draw from the distribution f_K^h, induced by a uniform draw from H. This is commonly referred to as randomization or permutation inference (see Rosenbaum 2002). A two-sided 10 percent level test of our null that g is not unusual relative to the set of groups in H corresponds to the following decision: reject if f_K^g is outside the range between the fifth and ninety-fifth percentiles of the f_K^h distribution. This distribution can be exactly calculated or approximated via simulation.

The choice of comparable groups H is the key decision for implementing our participation test. H must be chosen so that its groups have comparable cost structures to g. We have the data to do this in our application. We combine two strategies. First, we carefully assess all the formal legal restrictions on entry and consider for each possible firm-auction pair if the firm has the right certifications to bid in the auction. Second, using a regression approach, we evaluate what factors are most associated with entry in the ABA and FPA datasets and then condition choice of the groups H to such determinants. Thus, the usefulness of the participation test is crucially linked to the availability of a rich set of observable covariates. As discussed later, this limit will be somewhat less stringent for the bid test.

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

It is important to note that in typical (nonvalidation) datasets, H is very likely to contain both noncoordinating firms and undetected coordinating firms. We are testing the participation patterns of a group g compared to the groups in H. We are not testing g compared to a representation of the conduct of noncoordinating firms. Implementing this ideal comparison of g to known noncoordinating firms would require either typically unavailable data on known noncoordinators or an estimable model of their entry decisions, which we leave for future research. This composition issue for H can be an important consideration when choosing conditioning information used to construct H. We first present our main results for the Validation data and, then, we return to the issue of how they are affected by this composition issue at the end of this section.

Validation Data Results.—The first step in applying the test to the eight known cartels in our Validation data is to choose H. We construct H as the set of all groups of firms whose composition of distance, capitalization, and legal qualification match the given cartel. Matching is determined by categorizing subscribed capital and distance. We divide each characteristic into small, medium, and large categories (a third of all firms in each) and match firms based on the joint distribution of these distance and capital categories. For example, consider a cartel with eight members who are small distance and small capital and two members with medium distance and medium capital. This cartel will have an H that contains all groups of ten firms in our dataset that have the cartel's distance/capital configuration of eight small/small and two medium/medium firms and that have a configuration of the legal qualifications to bid equivalent to that of the cartel.[16]

We report the results obtained for the Turin cartels in Figure 2. For each of the eight cartels, the figure shows the frequency of participation of subgroups of all sizes. The dotted lines are the fifth and ninety-fifth percentiles of the reference distribution. For example, focus on panel A, we observe the largest subset of cartel B

[16] In addition, we are implicitly conditioning on factors that disciplined our dataset construction. All auctions involve roadwork jobs, a rather standardized type of contract, and were procured by the same PA.

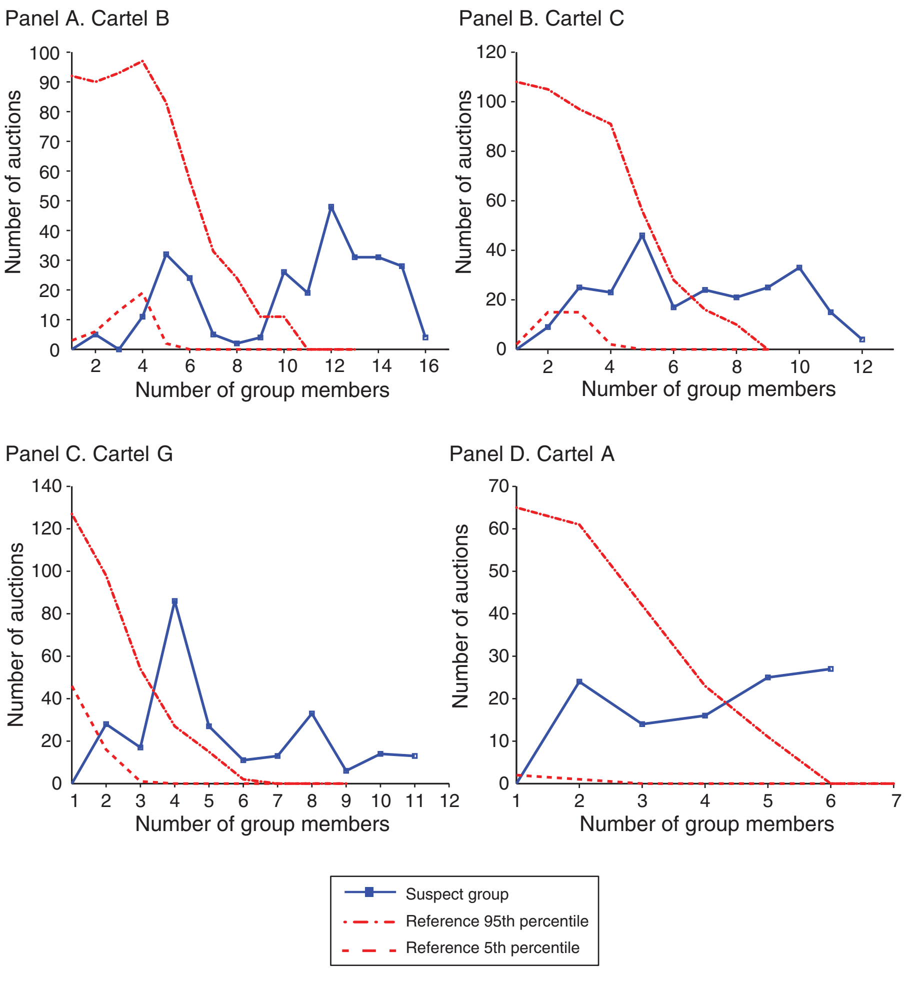

Figure 2. Participation Test—Validation Data (*continued*)

that jointly enters has size 16. However, the ninety-fifth percentile of the reference distribution for such a large group is approximately zero. Indeed, the ninety-fifth percentile of the reference distribution is estimated to be positive only for subgroups no larger than ten. Across cartels, the frequency of joint entry for larger sized suspect groups is much higher than that of the ninety-fifth percentile of the reference distribution. Larger-sized groups provide clear rejections of the noncoordination in entry. Therefore, the evidence presented in the remaining seven panels of Figure 2 also shows an entry behavior compatible with coordination between cartel members. A second relevant aspect for cartel B is that small subsets, of size 2, 3, and 4, have joint participation frequencies that are *lower* than the fifth percentile of the reference distribution. The same is true for cartel C for the subset of size 2. Thus, firms in B and C exhibit behavior consistent with a cartel considering minimum breaking coalition size when coordinating entry.

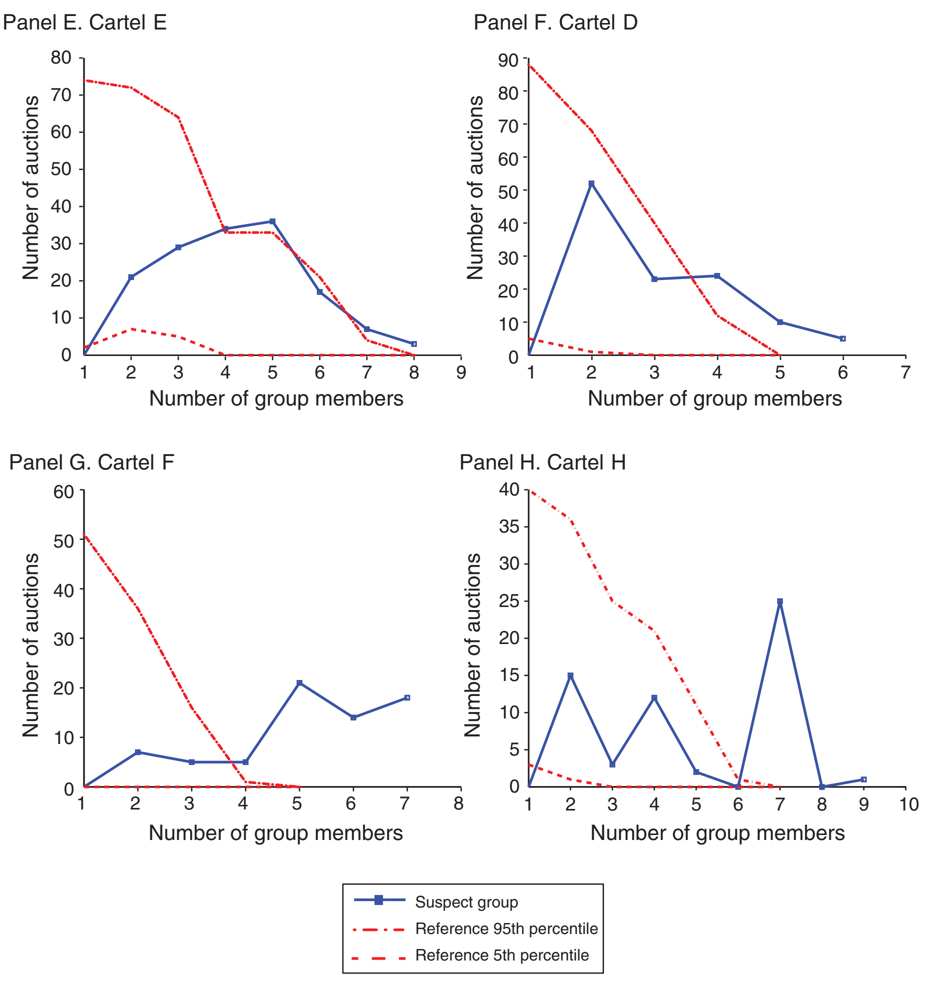

Figure 2. Participation Test—Validation Data (*continued*)

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

B. *Bid Test*

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

We base our test on a measure of how much influence a given set of suspected firms has on a trimmed mean discount $(A1)$ for an auction. Consider a group g suspected of piloting averages. We begin by considering a single auction with N total firms with N^g firms in group g and N^{-g} firms not in this group. We define $B^g = \{b_1^g, \ldots, b_{N^g}^g\}$ as the ordered (from small to large) set of discounts from group g and $B^{-g} = \{b_1^{-g}, \ldots, b_{N-N^g}^{-g}\}$ as the ordered set of remaining discounts. The trimmed mean throwing out N' discounts[17] on either end is

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

This statistic $A1^g$ will be systematically lower/higher than the trimmed mean of all the discounts if the group is trying to pilot the overall trimmed mean up/down. We compare $A1^g$ to its analogs for a set H of comparable groups. The trimmed mean without $h \in H$ is

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

We consider how $A1^g$ compares to the distribution of $A1^h$, induced by a uniform draw from H. Specifically, we compute the percentile of this distribution that corresponds to $A1^g$ and call it p^g. If the number of combinations makes H very large, one can approximate the percentile via simulation instead of calculating it exactly.

Our bid test combines together the p^g statistic from two or more auctions. First, consider a bid test for an environment with two auctions available. Let's indicate by g a suspect group that participates in both auctions. Use the notation p_1^g for the percentile of $A1^g$ in the first auction. For the second auction, use the notation p_2^g for the analogous statistic. Our joint test statistic J^g describes the extent to which these percentiles are extreme, either small or large, across the two auctions. Since manipulations of the mean can be either upward or downward, we construct a test that can detect both types of manipulations.[18] For below median percentiles we use the percentile itself and for percentiles above the median we use 100 minus the percentile as a measure of how far it is in the tail. To aggregate across auctions we add the individual "tail percentile" measures forming our statistic as

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

where $1\{\cdot\}$ is the indicator function. This test statistic will take on small values if both p_1^g and p_2^g are tail events and larger values otherwise. J^g clearly involves the same set of firms g in both auction one and two, and many other firms may also bid in both auctions, so the p_1^g and p_2^g statistics could have substantial dependence. In order to capture dependence across auctions in our bid test statistics we condition on participation by constructing a reference set M using only groups m that participate

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

[18] However, since upward manipulations are more likely because they are easier (in the sense described in Section III), results reported in the online Appendix repeat the analysis by looking at upward manipulations only. The results are qualitatively similar and are commented at the end of this section.

in auctions one and two. Our reference distribution for J^g under the null hypothesis of no coordination is the distribution of

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

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

Conditioning on costs is an important consideration for our bid test, clearly one reason for g to have unusual bidding structure even in the absence of coordination is that its firms' costs are unusual. In our bid test we have several options for how to condition on costs. We could construct H in each individual auction in a manner analogous to our participation test, so that its groups have the same cost determinant composition as g. We could also impose conditioning at the second stage where we define the set M across auctions. In addition to requiring that firms in M participate in auctions one and two, we can require them to have common cost determinant composition to firms in g. There is also an implicit conditioning on entry patterns since we are by construction looking at sets of firms that bid in the same auctions. It is important to examine the effectiveness of this implicit conditioning on participation by itself compared to explicit conditioning on cost determinants plus participation patterns. If these two approaches yield similar performance it would expand the scope of our testing procedure to contexts where detailed cost data are not available.

Conditioning on participation is not a perfect control for costs, but there is reason to believe it can work well. As shown in Section II, cost measures of distance-to-job and capital are strong predictors of FPA bids, determined in equilibrium by costs and markups. These cost measures are also strongly correlated with entry in both FPAs and ABAs. This correlation implies that entry patterns have the potential to be a useful summary statistic capturing variation in costs. We do not claim that participation patterns completely reveal cost structures, so conditioning on participation patterns does not guarantee firms in M have identical cost structures to those in g. We claim only that this conditioning on participation may be a useful control for costs. We utilize our validation data to explore the performance of conditioning on participation alone versus the best cost conditioning our data allow via legal qualification to bid, distance-to-job, and capital information.[19]

It is important to note that the number of auctions used in our bid test will impact its properties due to the same undetected coordinating versus noncoordinating firms

[19] Conditioning on a choice variable like entry implies that, in addition to conditioning on costs, we are conditioning on elements of firm strategies. This endogeneity of our participating set of firms does not affect our tests' ability to compare g versus the groups in M. However, it could have a big impact upon the composition of M, e.g., in terms of the proportions of noncoordinators and undetected coordinating firms. M could be dominated by undetected coordinating firms rather than noncoordinating firms. In this case our test will tend to correctly indicate that a coordinating g is not unusual relative to M, just because M itself has many coordinating groups in it. We are reassured that this is not too large a problem by the good detection performance of the bid test in the Validation data.

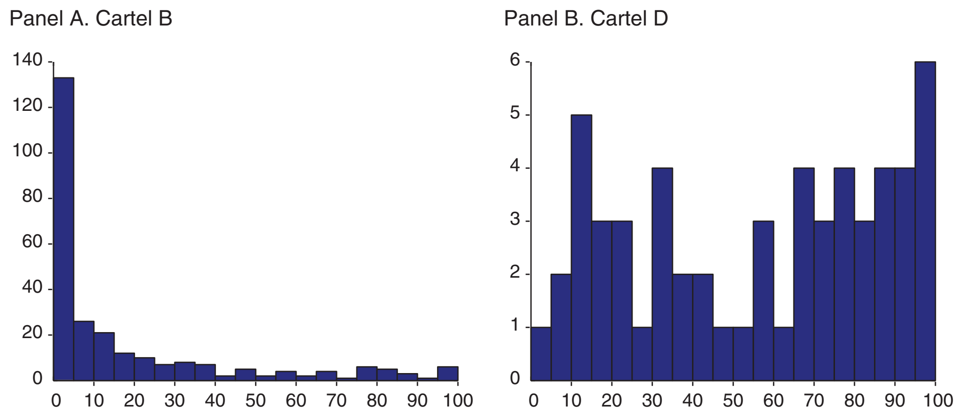

Figure 3. Cartel Bids Impact on the Trimmed Mean—Validation Data

Notes: Histograms of the statistic p^g for cartels B and D across all auctions in the Validation Data. A large spike near zero, like that observed for cartel B, means that in most of the auctions the trimmed mean $A1$ calculated excluding cartel B bids is substantially lower than that calculated for the comparison groups. A uniform-looking histogram, like that observed for cartel D, indicates that the impact of cartel D bids on the trimmed mean $A1$ is typically no different from that of the comparison groups.

composition issue we mentioned with our participation test. The set M will likely contain coordinating and noncoordinating firms. As we increase the number of auctions jointly attended, there will be a change in the composition of the groups in M and hence the distribution of J^m. The proportion of undetected coordinating firms relative to noncoordinators will likely grow as attendance is required at an increasing number of auctions. For example, if we conditioned upon all the members of a group attending dozens of auctions, we anticipate that very many of the firms satisfying these criteria would be those in undetected coordinating groups. It is implausible for uncoordinated firms to coincidentally attend auctions in such numbers. Our statistical test would often (correctly) indicate that a collusive group g was not unusual relative to groups in M, but this would not be an indicator of a lack of coordination. Thus, as the number of auctions jointly considered increases, there is a cost in terms of coordination detection performance eventually decreasing due to this composition effect.

Validation Data Results.—To illustrate the usefulness of our $A1$ statistic to detect unusual bidding behavior, the distributions of p^g values for groups whose members are in cartels B and D are illustrated in Figure 3. Consider first cartel B. The histogram describes the percentile of the reference distribution of $A1^h$ to which $A1^g$ corresponds for all the auctions in the Validation data where at least 3 members of cartel B were present. The test group g for each auction is comprised of all the cartel B firms in attendance. In each auction, H consists of all groups with the same size as g. A small percentile for $A1^g$ is consistent with a group trying to pilot the winning discount up and a large percentile is consistent with the group trying to pilot the winning discount down. Thus, for cartel B it appears that, in the large majority of the ABAs in which members participated, their behavior is consistent with upward average piloting. In contrast, the histogram for cartel D suggests that this cartel bids in a way that is not dissimilar from that of the groups in H. Thus, we do not have indications that cartel D manipulates $A1$. As previously discussed, cartel D was sanctioned less heavily due to less frequent bid rigging. We now discuss our bid test, which combines the p^g of the different auctions and accounts for common cost determinants.

The first choice needed to implement this test is the group g. We use a group that is a strict subset of the cartel since too large a group will result in too few noncoordinating firms jointly attending auctions. We choose the size of g to be four or the most frequent size in which cartel members participate in an auction (see Figure 2), whichever is greater. When there are multiple groups of this size, we choose g to be the one with the highest frequency of joint entry. As described in Table 5, we typically end up using groups of five firms.

Table 5 reports the results of our bid test for two sets of conditioning information applied to sets of two, four, six, and eight auctions. The columns labeled "Firm Controls" report results using tests that condition upon legal qualifications, firm distance to job, and capital, while the columns labelled "No Controls" report results where the only conditioning occurs through participation in the same auctions. For each cartel, there are two rows of entries. The first reports the median p-value over sets of two to eight auctions and the number below reports a count of these sets. We require that the auctions in these sets have at least 30 participants in common. The sets of auctions reported are randomly chosen from among all the potential combinations of auctions. The number of selected sets was chosen by imposing a time limit of one month for the Matlab routine searching for the elements to be included in the set, or 1,000 elements, whichever was reached first. For example, in the first column, the entries of 0.13 and 739 indicate that among 739 sets of pairs of auctions attended by the group from cartel B, the median p-value of our test was 0.13.

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

Since the p-values reported in Table 5 are not from independent sets of auctions, their distribution needs to be considered along with prior information/assumptions about the strength of dependence. Our strong prior beliefs are that this dependence is weak enough for substantial fractions of small p-values to be taken as evidence against the no-coordination null. Therefore, our conclusion from the findings reported in Table 5 is that our bid test is successful at detecting six of the eight cartels. In fact, when considering sets of up to eight auctions, the median p-value reaches a value below 0.10 (for the "Firm Controls" case) and below 0.05 (for the "No Controls" case) for all cartels, with the only exceptions of D and G. It is interesting to explore these latter two cartels in more detail.

As regards cartel G, failing to detect it is a bit surprising as the distribution of the p^g calculated separately for each auction and without conditioning provides clear evidence consistent with coordination in this cartel. The explanation lies in the structure and behavior of cartel G: this is a relatively large group of 19 firms, but only 5 of them win auctions. The nonwinning partners always place supporting

Table 5—Bid Test: Median p-Values

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

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

bids, generally consisting of very high discounts, while the few designated winners always bid closer to the center of the distribution. This implies that if we look at an auction in isolation and use all the firms in the cartel to construct p^g and all the firms bidding to construct the reference distribution, we often detect G as a coordinating group. What allows this cartel to evade detection in the bid test is that the designated winners are the only groups that frequently participate together, while individual supporting bidders participate sporadically. Therefore, our group selection method, selecting a subset of four firms within cartel G that jointly participate the most, results in a subgroup of four firms who are frequent winners and do not bid in an unusual manner. This highlights an important caveat of our bid test: its performance can be sensitive to the choice of group g.

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

V. Testing Coordination with Unknown Groups

Our testing methods can in principle be applied to any candidate group. In applications with a small number of firms, all possible groups could be examined. However, this is computationally infeasible for situations like that in our Main data with hundreds of bidders. Feasible strategies for selecting groups of firms will of course depend on the available information. Given the richness of our data, in this section we describe our favorite method that exploits both the presence of a Validation dataset and various firm covariates observed in both the Main and Validation data. Our Validation data allow estimation of the links between firms to predict their probability of being in the same cartel. The fact that our Main data are comparable to the Validation data allows us to use this estimated model to predict groups in the Main data. We examine both the "in-sample" performance of this method using the Validation data itself as the target, as well as its performance using our Main data. We make no claim that this group selection method is optimal, leaving the question of optimal group selection for future research. Our favorite group selection method has three steps.

Step 1: In both our Validation and Main data, we observe measures of firms' association along three dimensions: common ownership and management, formation of temporary bidding consortia and exchange of subcontracts.[20] Using the Validation data, we construct all pairs of firms that can be formed by linking each one of the 95 convicted firms to any of the other bidders, through any of the three association measures described above. This results in 775 pairs. Since in this dataset we know the composition of the eight cartels, we can estimate a model predicting which of these pairs are in the same cartel given their characteristics. We estimate a probit model where the dependent variable is one if the pair is in the same cartel and zero otherwise. Table 6 shows that the characteristics that we are analyzing help in predicting group membership. We also include measures of the geographical proximity between firms. Specification (1) in Table 6 indicates a positive association between

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

Table 6—Probit Regression–Validation Data

Probability that for a pair of firms both firms belong to the same cartel	(1)	(2)
Common personnel	0.94	1.67
	(0.26)***	(0.40)***
Common owner	0.07	−0.04
	(0.45)	(0.56)
Common manager	−0.67	−0.48
	(0.47)	(0.54)
Common zipcode	0.18	0.12
	(0.22)	(0.55)
Common municipality	−0.06	−0.03
	(0.16)	(0.17)
Common county	0.33	0.35
	(0.14)**	(0.14)*
Subcontract	0.88	1.89
	(0.19)***	(0.46)***
Winning consortium (all Piedmont contracts)	0.46	1.66
	(0.22)**	(1.29)
Bidding consortium (Validation data)	1.01	−2.15
	(0.18)***	(1.02)**
(1 − common personnel) × common zipcode		0.01
		(0.59)
(1 − common personnel) × W.Consortium		−0.59
		(1.14)
(1 − common personnel) × B.Consortium		1.42
		(0.66)**
(1 − common zipcode) × W.Consortium		−0.48
		(0.63)
(1 − common zipcode) × B.Consortium		0.07
		(0.47)
(1 − subcontract) × W.Consortium		0.94
		(0.58)
(1 − subcontract) × B.Consortium		0.97
		(0.51)*
(1 − W.Consortium) × B.Consortium		1.85
		(0.69)***
Constant	−2.23	−3.29
	(0.21)***	(0.47)***
Prob. χ^2	0.000	0.000
Observations	775	775

Notes: The dataset consists of all pairs of firms (from the Validation data) that share at least one owner (manager) or exchanged subcontracts or bid at least once as a legal temporary bidding consortium. The table presents probit coefficients and, in parentheses, their standard errors corrected following Conley (1999) for the correlation across any pairs that share firms. The dependent variable equals one if the pair belongs to the same cartel and zero otherwise. All independent variables are all dummy variables. The first three variables listed in Table 6 are equal to one if the couple shares, respectively, any white collar worker, any owner (regardless of the shares owned), or any top manager (regardless of his exact role). The following three variables equal one if the firms' headquarters are located, respectively, at the same zip code, in the same municipality, or in the same county. Subcontract equals one if the couple ever exchanged a subcontract. Winning Consortium equals one if the couple has won as a legal temporary bidding consortium at least one contract for public works held in Piedmont between 2000 and 2003. Bidding Consortium, instead, equals one if the pair of firms ever bid in the Validation data as a legal temporary bidding consortium. Model 2 differs only in that it includes interactions.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

the probability of being in the same cartel and exchanging subcontracts, sharing personnel, being located in the same county, and having bid jointly in a consortium. In our favorite specification, model (2), we also use interactions between the links to improve the model's predictive capacity.

Step 2: We use our estimates from the cartel membership probit model (Step 1) to generate predicted cartel membership probabilities for pairs of firms from the Main data. We will refer to these predictions as predicted group membership probabilities. To form a set of firm pairs, we begin by selecting the top 10 percent of firms in terms of participation, a set of likely suspects for group leaders. Each one of these firms is paired with the other firms in the Main sample with which they have at least one linkage due to common ownership and management, formation of temporary bidding consortia, or exchange of subcontracts. For each of these pairs, we construct a predicted probability of group membership using the estimates of model (2) of Table 6. The complements of these predicted probabilities are interpreted as a dissimilarity array.

Step 3: We use the constructed dissimilarity array from Step 2 with a standard hierarchical clustering algorithm (Gordon 1999) to partition the firms into clusters. In the first round of the algorithm, all firms are singleton clusters. In the next rounds, firms (or groups of firms) are associated together on the basis of their average dissimilarity. The process stops when a maximum tolerance for dissimilarity is reached. The clustering algorithm has a tendency to yield some very large and small clusters that we trim away to arrive at a set of candidate groups. Since this procedure entails arbitrarily chosen tolerance parameters, we provide its exact details in the online Appendix in the note to Table A.5.

The "in-sample" performance of this group selection method with our Validation data is reported in Table 7. Our method should work well in this case as it was in a sense tailored to this dataset. The first column is an integer enumerating each of the 14 clusters created by our three-step procedure. The second column reports a letter from A to H that identifies the cartel most often represented in the cluster. The following column reports the size of this set of cartel members. The following two columns report the number of members from different cartels and the number of noncoordinating firms. The last two columns report, respectively, the total number of victories of the members of the cluster and which, if any, of our two tests leads to a detection of unusual coordination. Our participation tests use the largest jointly participating set of firms within each cluster for g and detection coincides with the frequency of joint participation for g being above the ninety-fifth percentile of the reference distribution. For our bid test, we treat each cluster exactly as we treated cartels in the Validation data. We test for combinations of 2, 4, 6, and 8 auctions in which members of group g participated.

Overall this group selection method appears to perform reasonably well. The only cartel with no members in any assigned cluster is cartel D, which, as aforementioned, coordinated bids only sporadically. Although several noncoordinating firms are assigned to clusters, clusters 1, 4, 5, 7, and 9 contained a substantial fraction of members of the same cartel in addition to some noncoordinating firms. When clusters do not contain firms from cartels, our tests correctly do not indicate coordination. The

Table 7—Clusters in the Validation Data

Three-step method						
Assigned group	Known cartel	Members cartel	Members other cartels	Non-suspects	Auctions won	Detection
1	B	13	5	11	106	Both
2	B	1	0	3	6	No
3	B	1	1	2	5	No
4	C	4	0	3	15	Both
5	G	3	0	1	12	No
6	A	3	1	7	7	Part
7	E	10	0	7	6	Bid
8	F	2	0	2	4	No
9	H	3	0	2	0	Both
10	—	0	0	4	3	No
11	—	0	0	3	2	No
12	—	0	0	2	1	No
13	—	0	0	2	1	No
14	—	0	0	4	0	No

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

same lack of coordination evidence occurs when there are two or fewer members of the same cartel in a cluster. In five of the six clusters with three or more firms from a cartel, one or both of our tests rejects noncoordination. Table 7 also shows the limits of the procedure: our tests do not detect coordination for cluster 5, despite three of its four members coming from cartel G. However, in this case the reason is specific to the bidding strategies of cartel G. As discussed in the previous section, this is a large cartel with many fringe firms making piloting bids, but with a very small core of designated winners placing less extreme discounts. The three members of cartel G in cluster 5 belong to this subset of designated winners and this is why detection fails.

Poor Data Scenario.—Since many auction datasets often contain information only on bidder identities and bids, we also examine the performance of a method that forms groups based on participation patterns and then applies only our bid test to analyze coordination. When applied "in-sample" to Validation data, this method performs poorly compared to our favorite method. This is a further indication of the importance of firm covariates for our analysis. The online Appendix contains a more detailed discussion of this alternative approach while, in the following sections, we proceed with our method using firm covariates.

VI. Search for Coordinating Groups in Main Data

This section applies our tests to the Main data. We use the test results to identify a set of unusually coordinating firms. Using these firms as a benchmark, we then investigate the potential effect of coordination on revenues and on a regime switch from ABAs to FPAs.

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

We apply our tests to these 49 clusters producing the outcomes reported in the top panel of Table 8. The table provides details about those clusters for which at least one of our tests suggests coordination. We replicate the exercise detailed in Section IV and illustrated in Figure 2, and we find that the typical patterns are similar to those in that figure. We label a cluster as being unusually coordinated in entry if the majority of pointwise joint participation test statistic for groups of size 6 to 12 are greater than the ninety-fifth percentile of their reference distribution. This results in 42 clusters being classified as unusually coordinating with an average size of 10 firms each. This is indicated in the first row of Table 8. In total, there are 408 firms in these 42 clusters and their average number of bids, victories, and revenues are reported in the final columns of the table. For comparison these values can be related to those in the whole sample of firms reported in Table 1. Along all these dimensions, the average firm in the 42 clusters appears orders of magnitude larger that the average firm in the whole sample.

The second row of Table 8 reports results for bid tests. For each of our 49 clusters, we conduct a bid test by treating the cluster in the same manner as we treated cartels in the Validation data with the group g selected based on joint participation as detailed above. We conducted one and two-sided bid tests for all sets of 2, 4, 6, and 8 auctions in which the group g participated. Table A.4 in the online Appendix reports the median, tenth, and ninetieth percentile of the resulting distributions of p-values. Four clusters show clear indications of coordination having a median two-sided p-value less than 0.05 for at least one auction-set size. A fifth cluster shows some evidence of coordination having a median p-value for the two-sided test of 0.11 and one sided test of 0.05. We label these 5 clusters as being detected to have unusual coordinators according to our bid test. They are a subset of the 42 clusters detected as unusual by our participation test.

Given these definitions of coordinating groups, we can quantify the number of auctions potentially impacted by firms engaging in coordinated behavior. A basic measure of the volume of auctions impacted by coordination is the share of auctions receiving bids from at least three members from at least one of the clusters of coordinating firms. When using the 42 clusters detected by the participation test, this definition implies that 79 percent of the 802 ABAs in the Main data are affected. This share is 43 percent when using the 5 clusters detected by the bid test (and also

Table 8—Detection Results in the Main Data

Clusters detected as groups of coordinating firms					
Rejected test	Number of clusters	Cluster size	Entry	Number of victories	Revenues
Participation test	42	10	45.2	0.82	350,231
Bid test	5	16	59.0	1.08	462,914

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

by the participation test).[21] Using a more conservative measure that requires bids from at least 5 members instead of 3, the share of affected auctions becomes 64 percent and 34 percent using the clusters detected, respectively, by the participation and bid tests.[22]

An alternative approach to quantify how many auctions are affected by bid coordination would be to look for departures from the zero bid equilibrium. Using this approach implies that essentially all auctions involve coordination. This, however, likely overestimates coordination for at least two reasons. First, bidders might not know whether coordinated bids will be submitted or not. Hence, they might bid close to the typical winning range even when no coordinating groups are present. Second, although the zero-discount equilibrium is a likely outcome of a learning process where firms repeatedly bid in ABAs without coordination, the occasional presence of coordinating groups can prevent learning. Although the zero bid equilibrium is of little use to classify coordinated ABAs in the data, it is nevertheless the most appropriate benchmark to quantify the effects of the coordination detected in the data relative to what would happen were all firms to always compete with a single bid.

A. *Potential Effect of Coordination on Revenues*

The set of unusually coordinating clusters detected by our tests captures a significant share of the revenues in this market. For instance, considering the 5 clusters detected by the bid test, their members win 333 out of 802 ABAs, corresponding to a cumulative reserve price of €143 million out of a total of €370 million. However, contrary to typical cases of collusion in auctions, this is not necessarily an indication

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

[22] A formal test of whether an auction has suspect behavior from one of a set of groups is an alternative approach here and straightforward to implement. Testing a null that more than one group of specified sizes have the same distribution as a comparably sized comparison set of groups can be done via randomization inference in the same fashion as our tests. Test statistics determined by the set of groups outcomes can be compared to a reference distribution determined by randomly choosing sets of groups.

that the PAs could have paid a lower procurement price were these firms not engaged in bid coordination. In the counterfactual zero-discount equilibrium, the auctioneer pays the reserve price. Regulations mandate that this reserve price cannot be set based on the PAs' expectations about bidder behavior. It must be calculated by applying an official menu of prices to the estimated input quantities required by the work. This makes the observed reserve prices reasonable values for their counterparts in a counterfactual thought experiment without coordinated bids. This gives us a clear benchmark for this counterfactual scenario: all PAs would have paid an amount equal to the observed reserve price. Thus, in the Main data, at an average reserve price of €312,000, the average winning bid of 13.4 percent implies that the PA savings due to firm coordination is €42,000 per auction.

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

Consider an example scenario in which we can assess the relative importance of win probability reduction versus win payoff reduction in expected revenues for noncoordinators. A typical auction in our Main data has about 51 bidders, 17 of whom are members of our detected coordinating groups. Consider a hypothetical auction with 34 noncoordinating firms and 17 colluders. In the zero-bid equilibrium, each of the 51 bidders has a 1.96 percent chance of winning. Suppose that with coordination the 17 colluders can increase the probability that one of them wins to our sample group win frequency of $333/802$ and noncoordinating firms all have the same probability of winning. Thus, the win probability of the 34 noncoordinating firms drops to $(1 - 333/802)/34 = 1.70$ percent. Hence, there is a 13.2 percent decrease in the win probability for noncoordinators due to coordination among their competitors with a corresponding 13.2 percent decline in expected revenues. As above, we take our sample's 13.4 percent winning discount as representing the effect of coordination upon winning discounts. Insofar as this example is a reasonable benchmark for firms in our Main data, the effect of coordination upon win probabilities of noncoordinating firms appears to be as important as its effect on discounting a winning payoff in impacting expected revenues.

B. *The Policy Switch to FPAs*

Between 2006 and 2011, the introduction of new regulations from the European Union forced Italian PAs to replace ABAs with FPAs. In this section, we exploit this switch to study two separate issues: the performance of cartel firms across different sets of auctions and how bid coordination in ABAs helps our understanding of the drop in entry observed under FPAs.

Out of the 95 firms involved in the Turin case, 60 firms also bid in the Main data. Most of these firms had to change managers and owners as a consequence of their

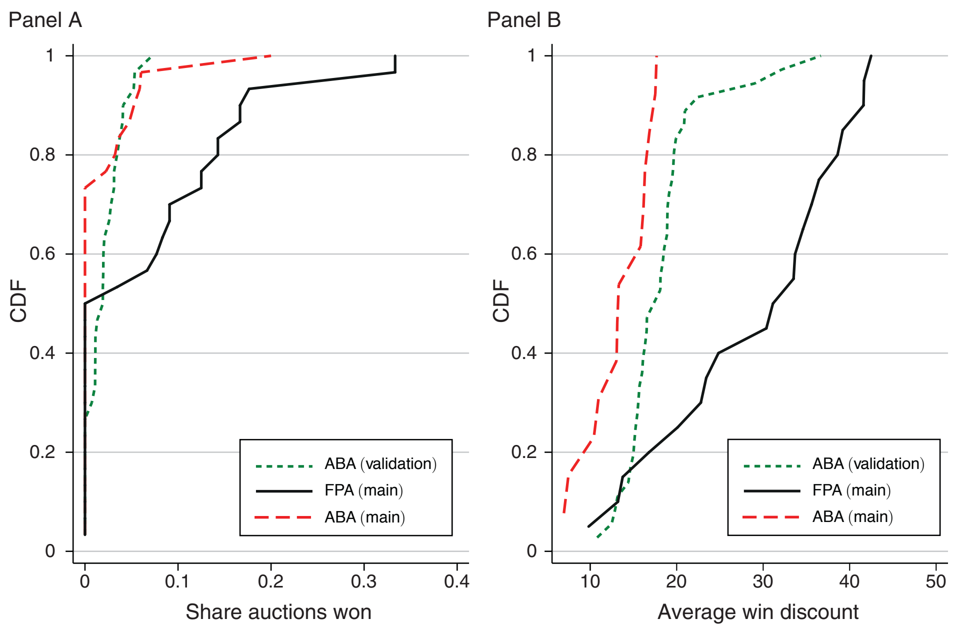

Figure 4. Performance of Cartel Firms across Formats

Notes: Empirical CDFs of the share of auctions won (panel A) and the average winning discount (panel B) for the 60 cartel firms of the Turin case observed in the Main data. CDFs are reported separately for three samples: ABAs in the Validation data (dotted line), ABAs in the Main data (dashed line), and FPAs in the Main data (solid line).

conviction.[23] Nevertheless, the new owners and managers are often close relatives of the former ones. Given this tight connection, it is therefore interesting to compare the performance of these 60 cartel firms in the ABAs of the Validation data to their performance in the ABAs and FPAs of the Main data. Figure 4 shows the empirical analogue of the cumulative density function of the share of auctions won over those entered (left) and the average winning discount (right) for these 60 firms across the three sets of auctions. Together the two plots show that none of the three environments is strictly better for these firms. Although the share of victories in FPAs is typically the highest, so is the discount at which they win in FPAs. The rather high winning discounts observed in the ABAs of the Validation data are consistent with the presence of competition driven by the rivalry between the eight cartels. Finally, the lack of a clear ranking in terms of win shares in the two sets of ABAs is consistent with the possibility of the formerly convicted firms continuing to coordinate bids. Indeed, some former cartel members belong to the suspect groups of Table 8.

A clearer ranking is offered in Figure 5, where we compare the 60 cartel firms to non-cartel firms participating in the same auctions. The comparison is within each of the three sets of auctions and is again in terms of probability of winning (top three

[23] For the remaining 35 firms, their exit is indeed in part connected to their legal prohibition to bid in public auctions without replacing all convicted owners and managers. In part, however, it is also the result of sample attrition as the Main data begins in 2005, about two years after the end of the Validation data.

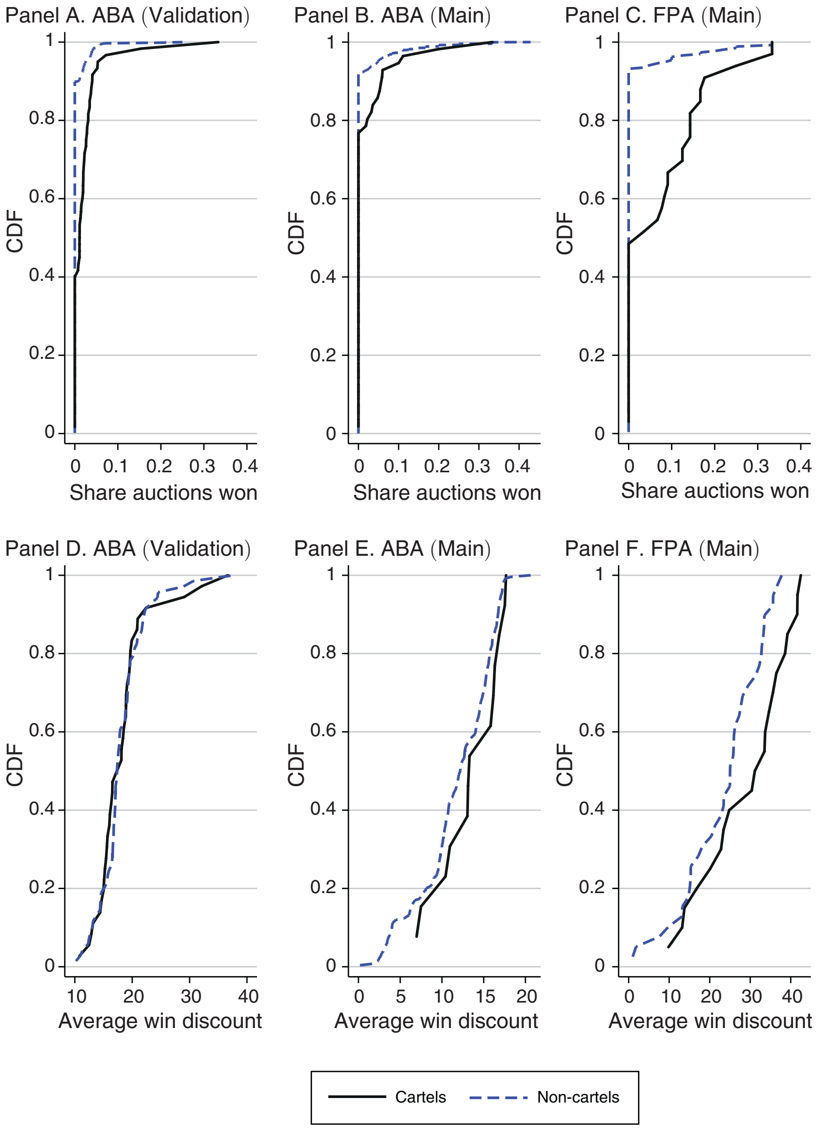

FIGURE 5. RELATIVE PERFORMANCE OF CARTEL AND NON-CARTEL FIRMS

Notes: Empirical CDFs of the share of auctions won (top plots) and the average winning discount (bottom plots) for two sets of firms: The 60 cartel firms used in Figure 3 (solid line) and the non-cartel firms selected as described in the text (dashed line).

plots) and winning discount (bottom three plots). In Validation data ABAs, cartel firms (solid line) outperform non-cartel firms (dashed line) due to the higher probability of winning despite the identical winning discounts. In Main data ABAs, cartel firms tend to have a higher probability of winning, but also to win at a worse price. A similar and even more pronounced pattern is present in FPAs. Therefore, while in the ABAs of the Validation data the cartel firms do better than the other firms, no clear ranking can be established in the other two cases.

We study the drop in entry associated with FPAs by focusing on all the firms active in the Main data. Figure 6 offers a clear image of this drop: the black triangles mark the ABAs and the hollow grey circles mark the FPA. The top panel reports the number of bidders in ABAs and FPAs held by four PAs in the Main data that switched to FPAs. The systematically lower values of the circles (FPAs) relative to the triangles (ABAs) is evident. The bottom panel of Figure 6 documents that a similar drop in participation occurred for the PAs in Turin.

There are two main causes for the drop in participation with the introduction of FPAs. The first reason is the exit of less efficient firms that have too little chance of winning FPAs. As shown in Table 1, winning discounts are indeed substantially higher in FPAs than in ABAs. The second reason is the disappearance of shills in FPAs. Since known coordinating groups are composed of both separate firms and shills, the fact that shill firms are not useful in FPAs can explain a large share of the market shakeout. From a policy perspective, distinguishing between the exit of shills and inefficient firms matters because the regulator might want to foster the participation of some less efficient firms, but most likely not of shills. In the Main data, about 4,000 firms bid at least once in ABAs and only about 1,000 bid once or more in FPAs. Not all of the ABA bidders are, however, potential FPA participants. Focusing on firms that were qualified and near to prospective FPAs, we examine 1,482 firms who attended at least three ABAs in counties where subsequently at least three FPAs for which they were legally qualified to bid were held. The 1,482 firms contain 298 members of our 42 coordinating groups (see Table 8) and 1,184 noncoordinating firms.[24] Of the 298 coordinators about half do not participate in an FPA and likewise about half of the 1,184 noncoordinators also do not participate in an FPA. Referring to those not entering in FPAs as exiters, the frequency of exiters does not depend on coordinating status.

Characteristics for these four sets of firms are reported in Table 9. We anticipate that shill firms will be predominately located in our detected coordinating clusters (with a perfect measure of coordination, shills would *only* be present among coordinators). Thus, the composition of exiters in terms of shills versus inefficient firms should vary according to whether the firms are coordinators and should show up in the firm characteristics. We find clear differences in the characteristics of exiters according to whether they are labeled coordinators or not. Among noncoordinators, exiters have smaller capital and labor force relative to those who participate in FPAs despite being slightly older firms, possibly signaling their relative inefficiency. Exiters among coordinators also have less capital and workers than FPA participants, but these gaps are smaller than for noncoordinators.

An important caveat to the interpretation of the ownership and management characteristics reported in Table 9 is that there are serious missing response issues. We do not have the data to address this issue and necessarily proceed to interpret these statistics as though nonresponse was random.[25] With this caveat in mind,

[24] The 42 groups in the top row of Table 8 are used to classify group firms. Qualitatively, the results do not change if one of the two more stringent classifications is used.

[25] With more complete data, the shock given by the switch to FPAs could have been exploited to more rigorously trace out the connections between firms, in the spirit of Bertrand, Metha, and Mullaliatan (2002).

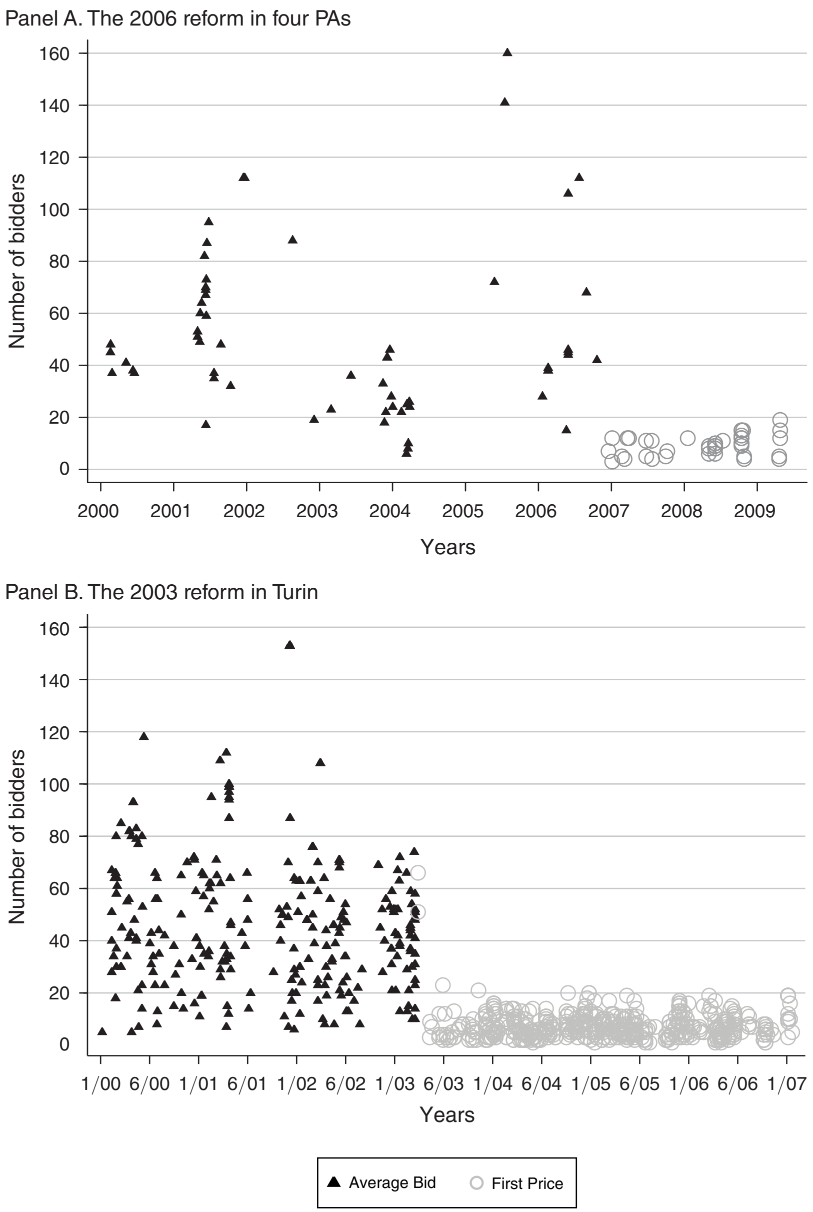

Figure 6. Number of Bids in ABAs and FPAs

Notes: Panel A: Scatter plot of the number of bidders in ABAs and FPAs held by four PAs in the Main data: Padova, Varese, Sondrio, and Cremona, which all switched to FPAs. Panel B: Scatter plot of the number of bidders in the ABAs and FPAs held by the municipality of Turin.

Source: Italian Authority for Public Contracts

there do appear to be female ownership and management differences according to coordination status. For noncoordinators, exiters have lower or nearly the same frequency of female ownership and management presence. In contrast for coordinating firms there is modest evidence of exiting firms having more female owners and

Table 9—Firms' Size and Gender Composition

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

Notes: The table reports statistics for four sets of firms: (i) non-coordinating firms that never bid in FPAs (top left), (ii) non-coordinating firms that bid in FPAs (top right), (iii) group members that never bid in FPAs (bottom left), and (iv) group members that bid in FPAs (bottom right). Firms are classified as group members if they belong to any one of the 42 clusters described in the top row of Table 8. A firm is in the entering-FPA group if it bids in at least one FPA. A firm is in the not-entering-FPA group if: (i) it never bids in any FPAs and (ii) it bids in at least three ABAs held in counties where at least three FPAs (for which the firm was qualified to bid) were held. For each of the four sets, the columns Mean and SD report the average and standard deviation taken across all firms in the set. The column N reports the number of firms considered. The firm characteristics considered are: the number of years between the beginning of activity and 2010 (Firm age) and the average value between 2006–2010 of the number of all dependent workers (Number of workers), the fraction of female white collar workers over all white collar workers (Proportion of women), the number of female owners (managers) (Number female owners (managers)), the ratio of the number of female owners (managers) to that of the total number of owners (managers) (Proportion female owners (managers)) and (expressed in €1,000) capital, revenues, and profits.

managers versus those that stayed and participated in FPAs. This is in line with the legal case in Turin where shill firms were often formally owned and managed by the mothers, sisters, or wives of the men convicted for collusion. The presence of shills is also suggested by some ad hoc comparisons of the firms in the five groups detected by our bid test. For example, we have a few instances of pairs of firms registered at the exact same street address that bid together in almost all the ABAs in which they participate, but that have only a single member of the pair bidding in FPAs.

VII. Conclusions

In this paper, we document that ABAs give strong incentives for bidders to coordinate their entry and bidding choices. We propose two statistical tests to investigate bidder coordination and show that they work well in Validation data where eight cartels have been identified by a court. These are tests for whether groups of firms participate or bid differently than other comparable groups of firms. Our metrics for describing participation and bidding patterns are motivated by how coordinating firms can coordinate their bids to pilot the thresholds that determine the awarding of the contract. Finally, we apply these tests to a different dataset of ABAs in which the presence of coordinating groups has not been previously known and show that the tests suggest the presence of several coordinating groups influencing numerous auctions. Thus, although no statistical test is a final proof, a natural application of our tests could be of help to courts evaluating cases of coordinated bidding. In this respect, a good feature of our tests is that they are somewhat "inspector proof" in that even if firms knew of them, avoiding detection would require foregoing, at least in part, the benefits of coordination.

We are optimistic that our tests could be adapted to detect coordination in other environments where similar incentives to manipulate thresholds exist. Similar types of manipulable mechanisms are fairly common in numerous relevant markets ranging from the procurement of public works to financial markets (the LIBOR being the most striking case), health care markets (like the subsidies awarded to insurers in Medicare D), and even labor markets.

Importantly, our results also indicate that it is not obvious that bidder coordination should always be sanctioned. We present the case of a market in which bidder coordination reduces the procurement cost for the auctioneer relative to an environment where firms compete in ABAs submitting one bid per firm. Thus, our results argue against any automatism in antitrust activity. Instead, we see a role for the use of an accurate economic analysis of bidder behavior as a guide to the quantification of the effects of coordination. Our results are thus coherent with the current view of the US antitrust policy where, at least since *Broadcast Music v. Columbia Broadcasting System*, 441 *US* 1, the Supreme Court has favored a careful analysis of price fixing practices rather than considering them per se violations.

REFERENCES

Abrantes-Metz, Rosa, and Patrick Bajari. 2012. "Screens for Conspiracies and Their Multiple Applications (reprint)." *Competition Policy International Journal* 8 (1) 177–83.

Abrantes-Metz, Rosa M., Michael Kraten, Albert D. Metz, and Gim S. Seow. 2012. "Libor manipulation?" *Journal of Banking and Finance* 36 (1): 136–50.

Asker, John. 2010. "A Study of the Internal Organization of a Bidding Cartel." *American Economic Review* 100 (3): 724–62.

Athey, Susan, and Jonathan Levin. 2001. "Information and Competition in U.S. Forest Service Timber Auctions." *Journal of Political Economy* 109 (2): 375–417.

Bajari, Patrick, and Lixin Ye. 2003. "Deciding between Competition and Collusion." *Review of Economics and Statistics* 85 (4): 971–89.

Bertrand, Marianne, Paras Mehta, and Sendhil Mullainathan. 2002. "Ferreting out Tunneling: An Application to Indian Business Groups." *Quarterly Journal of Economics* 117 (1): 121–48.

Conley, Timothy G. 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics* 92 (1): 1–45.

Conley, Timothy G., and Francesco Decarolis. 2016. "Detecting Bidders Groups in Collusive Auctions: Dataset." *American Economic Journal: Microeconomics*. http://dx.doi.org/10.1257/mic.20130254.

Coviello, Decio, and Mario Mariniello. 2014. "Publicity requirements in public procurement: Evidence from a regression discontinuity design." *Journal of Public Economics* 109 (1): 76–100.

Decarolis, Francesco. 2013. "Comparing Public Procurement Auctions." http://people.bu.edu/fdc/Decarolis_ProcAucCompare.pdf.

Decarolis, Francesco. 2014. "Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions." *American Economic Journal: Applied Economics* 6 (1): 108–32.

Decarolis, Francesco. 2015. "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" *American Economic Review* 105 (4): 1547–80.

Decarolis, Francesco, and Cristina Giorgiantonio. 2015. "Local Procurement Regulations: The Case of Italy." *International Review of Law and Economics* 43: 203–26.

Dimitri, Nicola, Gustavo Piga, and Giancarlo Spagnolo. 2006. *Handbook of Procurement.* Cambridge: Cambridge University Press.

Duggan, Mark, and Fiona Scott Morton. 2006. "The Distortionary Effects of Government Procurement: Evidence from Medicaid Prescription Drug Purchasing." *Quarterly Journal of Economics* 121 (1): 1–30.

Engel, Andreas R., Juan-José Ganuza, Esther Hauk, and Achim Wambach. 2006. "Managing Risky Bids." In *Handbook of Procurement*, edited by Nicola Dimitri, Gustavo Piga, and Giancarlo Spagnolo, 322–44. Cambridge: Cambridge University Press.

Gordon, Aaron D. 1999. *Classification: Methods for the Exploratory Analysis of Multivariate Data.* Boca Raton, FL: Chapman and Hall.

Haberbush, Kara L. 2000. "Limiting the Government's Exposure to Bid Rigging Schemes: A Critical Look at the Sealed Bidding Regime." *Public Contract Law Journal* 30 (1): 97–122.

Harrington, Joseph E., Jr. 2012. "Exploring the Boundaries of Unlawful Collusion: Price Coordination when Firms Lack Full Mutual Understanding." http://www.tilburguniversity.edu/upload/35c32b08-bee4-47c3-bdb2-1749f008e859_5december.pdf.

Hendricks, Kenneth, and Robert H. Porter. 1989. "Collusion in Auctions." *Annales D'Economie et de Statistique* 15–16: 217–30.

Ishii, Rieko. 2009. "Favor exchange in collusion: Empirical study of repeated procurement auctions in Japan." *International Journal of Industrial Organization* 27 (2): 137–44.

Jofre-Bonet, Mireia, and Martin Pesendorfer. 2003. "Estimation of a Dynamic Auction Game." *Econometrica* 71 (5): 1443–89.

Lewis, Gregory, and Patrick Bajari. 2011. "Procurement Contracting with Time Incentives: Theory and Evidence." *Quarterly Journal of Economics* 126 (3): 1173–1211.

Marshall, Robert C., and Leslie M. Marx. 2012. *The Economics of Collusion: Cartels and Bidding Rings.* Cambridge: MIT Press.

Ortner, Juan, and Sylvain Chassang. 2014. "Making Collusion Hard: Asymmetric Information as a Counter-Corruption Measure." Princeton University William S. Dietrich II Economic Theory Center Research Paper 064-2014.

Pesendorfer, Martin. 2000. "A Study of Collusion in First-Price Auctions." *Review of Economic Studies* 67 (3): 381–411.

Porter, Robert H., and J. Douglas Zona. 1993. "Detection of Bid Rigging in Procurement Auctions." *Journal of Political Economy* 101 (3): 518–38.

Porter, Robert H., and J. Douglas Zona. 1999. "Ohio school milk markets: An analysis of bidding." *RAND Journal of Economics* 30 (2): 263–88.

Rosenbaum, Paul. 2002. *Observational Studies.* Springer Series in Statistics. New York: Springer Science and Business Media.

Scott Morton, Fiona. 1997. "The strategic response by pharmaceutical firms to the Medicaid most-favored-customer rules." *RAND Journal of Economics* 28 (2): 269–90.

Spagnolo, Giancarlo, Gian Luigi Albano, and Milo Bianchi. 2006. "Bid Average Methods in Procurement." *Rivista di Politica Economica* 96 (1): 41–62.