Comparing Public Procurement Auctions

Francesco Decarolis*

May 17, 2013

Abstract

This paper contrasts two auction formats often used in public procurement: first price auctions with ex-post screening of bid responsiveness and average bid auctions, in which the bidder closest to the average bid wins. In equilibrium, their ranking is ambiguous in terms of revenues, but the average bid auction is typically less efficient. Using a dataset of Italian public procurement auctions run alternately under the two formats, a structural model of bidding is estimated for the subsample of first price auctions. Lower bound estimates of the efficiency loss under the counterfactual average bid auctions range between 11 and 41 percent.

JEL: L22, L74, D44, D82, H57.

*Boston University, Department of Economics. E-mail address: *fdc@bu.edu*.
Acknowledgements: This is a revised version of a chapter of my Ph.D. thesis, supported by the Henry Morgenthau dissertation fellowship of the University of Chicago and the Economic Fellowship of the Bank of Italy. I am grateful to Ali Hortacsu, Roger Myerson, Philip Reny and Timothy Conley for their guidance.

When procuring a contract to execute a public work, auctioning it off at the lowest price does not ensure paying the lowest procurement cost. Because of cost uncertainty at the time of bidding, a low price in the auction stage might come at the cost of poor ex post contract performance. In the context of public procurement, where transparency considerations have fostered the use of sealed bid auctions as the main allocation mechanisms, this has lead to the proliferation of auction formats that deviate from the well know first price auction.

This study contrasts from both a theoretical and an empirical perspective two such auction formats frequently used in public procurement. The first format consists of supplementing a conventional first price auction (FPA) with an additional stage in which the bids received are screened for their reliability. Hence, the winner is not necessarily the firm offering the lowest price, but the firm offering the lowest price among those deemed reasonable by the auctioneer. Instances of this modified first price auction are common. For example, in the context of the public procurement of roadwork contracts by the California DoT, Bajari, Houghton and Tadelis (2007) report that in 4 percent of the FPAs in their study the lowest price is disregarded because this price is considered unreasonably low by the DoT engineers.

The second auction format that I consider consists of awarding the contract to the firm offering the price closest to the average price (or to a more complicated function of the average, like a trim mean). The winner is then paid his own price to complete the contract. This format is typically known as an average bid auction (ABA). Although not common in the US, where it appears to have been used only by the Florida DoT and the New York State Procurement Agency, the ABA is present in the public procurement regulations of many countries including Chile, China, Colombia, Italy, Japan, Peru, Malaysia, Switzerland and Taiwan. Moreover, its usage has been suggested by both the civil engineering literature, Ioannou and Leu (1993), and international institutions, European Commission (2002).

In the first part of this paper, I present a stylized model of public procurement where firms face production cost uncertainty and asymmetric costs of defaulting on their bid. This model exhibits the well known perverse property of first price auctions: Those firms that have lower costs of defaulting anticipate this benefit, offering low prices that make them highly likely both to win and to default after the cost uncertainty is realized. Then, I turn to analyze equilibrium bidding under ABAs and FPAs with bid screening and show that both mechanisms are effective at limiting default risk. For the latter format, this is directly due to bid screening. For the ABA, instead, this occurs because in equilibrium this auction resembles a random lottery that awards the contract at a high price. Both the fact that the allocation is random and that the price is high limit the scope for strategic bidder defaults.

Although both formats limit the risk of a winners' default, they are not equivalent. In particular, I show that their ranking in terms of the revenues generated for the auctioneer is ambiguous: The winning price is lower in the FPA with screening, but since screening is costly the overall auctioneer cost under the ABA might be lower. Nevertheless, I show that their ranking is essentially unambiguous in terms of allocative efficiency. Since the ABA in equilibrium resembles a lottery, this format will typically be less efficient.

The size of the inefficiency produced by ABAs, however, crucially depends on the dispersion of firm production costs. To simplify, if the production costs were essentially the same across all firms, the inefficiency produced by the random allocation of the contract would be negligible. Thus, the relative inefficiency of ABAs is ultimately an empirical question.

In the second part of the paper, I address this question by analyzing a dataset of Italian public procurement auctions held alternately under the ABA or the FPA with screening. This dataset, collected for this study, covers several thousand auctions for road construction and maintenance held between 2000 and 2013 by counties and municipalities in the North of Italy. The descriptive analysis of the data confirms various theoretical predictions and, in particular, that the allocation produced by ABAs is substantially different from that of FPAs and that it resembles a random lottery at a high price. This motivates me to conduct a structural estimation procedure to more thoroughly explore the relative efficiency of the two mechanisms. Since the bids offered in ABAs do not bear any clear connection to firm costs, the structural estimation relies exclusively on the subsample of FPAs. The estimation method used extends that of Krasnokutskaya (2011) to permit identification and estimation with auction datasets where the econometrician does not observe all the bids, but observes at least the winning bid, the reserve price and the number of bidders.

The main estimation outcomes are the estimates of two separate distributions, one for the private, idiosyncratic production cost of each bidder and one for their common production cost. These estimates very clearly suggest that the inefficiency of the ABA is potentially quite large. In particular, to explore the performance of the ABA, I use these estimates to simulate three counterfactual scenarios mimicking the allocation produced by ABAs under different cases that appear to be relevant in the data. These scenarios account for both the higher bidder turnout in ABAs relative to FPAs and for the presence of groups of cooperating bidders in ABAs. Across these different scenarios, the simulation reveals that the cost of the winner under an ABA is between 11 and 41 percent higher than what would be in the corresponding FPA. Moreover, the share of auctions in which the winner of the ABA has a cost strictly above that of the winner of the corresponding FPA ranges between 50 and 86 percent. Importantly, these estimates are obtained under assumptions that make them best interpreted as a lower bound on the inefficiency of ABAs.

This paper has three main contributions. The first contribution is to bridge the vast theoretical literature on the perverse effect of FPAs when bidders can default and two alternative formats that are frequently encountered in real world public procurement. An incomplete list of the main studies in this literature includes: Spulber (1990), Waehrer (1995), Zheng (2001), Rhodes-Kropf and Viswanathan (2005), Board (2007), Chillemi and Mezzetti (2009), Burguet, Ganuza and Hauk (2009) and Che and Kim (2010). More in detail, the characterization of equilibrium bidding in the ABA is an important result that sheds light on this format that has received limited attention in the existing theoretical literature. The two previous studies that analyzed this format, Spagnolo, Albano and Bianchi (2006) and Engel et al. (2006), characterized its properties under assumptions on the number of bidders and their cost and information structures that were more restrictive than those used in this paper.

The second contribution is to quantitatively compare two different auction formats used in the same market. Only a few other studies have accomplished this goal because auctions are typically very persistent institutions so that format changes are rarely observed. All these other studies involve public auctions in the US: Athey, Levin and Seira (2011) compare open vs. sealed bid auctions used for the sale of timber harvesting contracts, Lewis and Bajari

(2011) compare first price vs. scoring rule with time incentive auctions for the procurement of roadwork contracts and Marion (2007) compares first price vs. first price with small-business bid subsidy for roadwork contracts.[1] Methodologically, in order to estimate the bidder cost distributions that allow me to compare the two auction formats I follow the method of Krasnokutskaya (2011). This method is also used in Asker (2010) and originates from the pioneering study of Li, Perrigne and Vuong (2000) on unobserved auction-level heterogeneity. Similarly to Athey, Levin and Seira (2011) and Athey, Coey and Levin (2013), I only use one of the two auction formats observed in the data to estimate bidder costs because only in this format does the theory provide a mapping between bids and costs. In a complementary paper, Decarolis (2013), I study the effects of the format switch on observable auction outcomes.

The third contribution concerns the policy implications stemming from the paper findings. The major inefficiency estimated for the ABA suggests that its continued use in public roadwork contracts procurement is particularly wasteful. Nevertheless, an effective solution is unlikely to be either a naive adoption of first price auctions, because of the risk of costly defaults, or of first price auctions with screening, because of the high screening cost. Adequate solutions, instead, should involve the simultaneous adoption of an efficient auction format, like a first price auction, *and* of effective methods to reduce the default risk, which combine elements of a centralized bid screening system, stricter qualification criteria for bidders, insurance policies (in the form of performance bonds) and higher penalties in case of default.[2]

I Theoretical Analysis

This section presents a stylized bidding model. Similar to Zheng (2001), a main feature of the model is that bidders face different default costs. However, contrary to Zheng the default

[1] Krasnokutskaya and Seim (2011) study the same bid preference system studied by Marion (2007). However, all the data in Krasnokutskaya and Seim (2011) come exclusively from the first price auctions with bid subsidy. Athey, Coey and Levin (2013) study a closely related question analyzing timber auction run alternately with or without set asides for small business.

[2] Spulber (1990) and Calveras, Ganuza and Hauk (2004) analyze the relative merits of these methods.

cost is observable to all bidders and the source of asymmetric information is represented by a privately observed project completion cost. In addition to this privately observed cost, bidders also share a commonly observed completion cost as in Krasnokutskaya (2011), but I assume this cost to be uncertain ex ante. The first subsection shows that FPAs pose a severe default risk in this environment, while the second one illustrates how ABAs and FPAs with screening alleviate this problem.

A. Baseline First Price Auction

Consider a first price procurement auction in which N risk-neutral bidders compete to win one project. When bids are submitted, the cost to complete the project is uncertain. For any bidder i, with probability $(1-\theta)$ the cost of the project is $c^i = y + x^i$, while with probability θ the cost is $c^i + \varepsilon = y + x^i + \varepsilon$, where $0 < \varepsilon < y$ and $0 < \theta < 1$. One part of the cost, x^i, is only privately observed by bidder i, while the other part, y, is commonly observed by all bidders. Likewise, ε and θ are constants known to all bidders.

After being awarded the contract, the winner observes the full cost of the project. At this stage, the winner has two options: either he completes the project at the promised bid, or he defaults. In the latter case, his payoff is equal to $-p \leq 0$, the penalty that he pays. To capture in a simplified manner features of the application that I will discuss later, I assume that there are two types of bidders, L and H, who face different penalties for defaulting: There are $n_H > 2$ bidders of type H, who pay a large penalty (p_H), and $n_L = N - n_H$ bidders of type L, who pay a low penalty (p_L), $p_H > p_L \geq 0$.

Both the type and the number of bidders are observable to all bidders. Moreover, bidders know that each type of bidder independently draws his privately observed cost x from a type-symmetric distribution F_{X_j}, $j = \{H, L\}$, that is assumed to be absolutely continuous and have support on $[\underline{x}_j, \overline{x}_j]$, where $0 \leq \underline{x}_j < \overline{x}_j < \infty$.

This model fits squarely into the commonly used *independent private value* paradigm, with the sole complications coming from common uncertainty regarding the shared cost component (y) and the possibility of costly default. However, the possibility of default affects the game only when ex post the project turns out to be costly to complete because defaulting on a cheap contract is a dominated strategy.[3] Thus, disregarding dominated strategies, the expected payoff for a bidder of type $j = \{L, H\}$ bidding b_j can be written as:

$$[(1-\theta)(b_j - (y + x_j)) + \theta \max\{-p_j, b_j - (y + x_j + \varepsilon)\}] \Pr(win|b_j).$$

To simplify the analysis, I make the following restriction on the game parameters:

Assumption (i): $\frac{\overline{x}_L - \underline{x}_L}{1-\theta} < \varepsilon < y$, *and the two bidder types have* $p_H > p_H^*$ *and* $p_L < p_L^*$,

where p_H^* and p_L^* are two constants characterized in the appendix. Their role is to ensure that for type H bidders the penalty is high enough that it is never optimal to default, while for type L bidders the penalty is low enough that they always optimally default if the cost is high when the format is a first price auction. This greatly simplifies the game by allowing me to write bidder expected payoffs in the FPA conditional on bidding b_j as:

$$\begin{cases} [b_H - x_H - a_H] \Pr(win|b_j) & \text{if bidder type H,} \\ [b_L - x_L - a_L](1-\theta) \Pr(win|b_j) & \text{if bidder type L,} \end{cases} \tag{1}$$

where a_H and a_L are constants such that $a_H \equiv (y + \theta\varepsilon)$ and $a_L \equiv (y + \frac{\theta}{1-\theta} p_L)$. Finally, I assume that there is a commonly known reserve price, R, which represents the maximum price that the auctioneer is willing to pay. This reserve price is assumed to be non binding in the sense that even the least efficient bidder can earn a profit if he wins at the reserve price.

The equilibrium analysis focuses on type-symmetric Bayes-Nash equilibria (BNE), which consist for every bidder i of type $j = \{L, H\}$ of a continuous function $b_j : [\underline{x}_j, \overline{x}_j] \to R_+$ and a decision of whether to default if the cost of the project is high; these two elements together maximize i's payoff conditional on the other bidders bidding b_j.

I begin by showing the perverse features of first price auctions in this environment. The game described above is isomorphic to a FPA with asymmetric bidders. Thus, under

[3] Under the stated assumptions, if a bidder optimally chooses to default when the cost is low, then he must do so also when the cost is high. Thus, the payoff of this strategy in case of victory is $-p \leq 0$. However, this strategy is strictly dominated by bidding $c + \varepsilon$, which guarantees a payoff in case of victory of $(1-\theta)\varepsilon > 0$.

Assumption (ii) below, Lemma 1 follows from results in de Castro and de Frutos (2010):

Assumption (ii): *Type H hazard rate dominates type L:* $\frac{f_{X_H}}{1-F_{X_H}} < \frac{f_{X_L}}{1-F_{X_L}}$.

Lemma 1. *An equilibrium exists. In equilibrium, if* $(\underline{x}_L - \overline{x}_H) < (a_H - a_L) < (\overline{x}_L - \underline{x}_H)$, *despite type H shading their cost less than type L for the same cost draw, the bid distribution of type H bidders first order dominates that of type L bidders. (Proof in appendix)*

The restriction that $(\underline{x}_L - \overline{x}_H) < (a_H - a_L) < (\overline{x}_L - \underline{x}_H)$ ensures that the supports of type L and H's cost distributions overlap. I will maintain this restriction throughout the analysis since without it the game would only have equilibria where one type always wins. Lemma 1 is an example of the well known result that weakness leads to aggression: Type H bidders shade their cost less to try to compensate for the cost advantage that the possibility of default gives to type L bidders. The auctioneer benefits from the need of type H bidders to bid aggressively. Nevertheless, the downside for the auctioneer is that a default is likely to happen whenever the contract is costly to complete: The FPA favors allocating the project to the less reliable type L and does so at such a low price that a default is likely. Since a default can entail monetary, welfare and even political cost for the government, it is evident why alternative mechanisms are often preferred to the FPA for public procurement.

B. Alternative Auction Format I: Average Bid Auction

The two alternative mechanisms that I analyze are an average bid auction and a first price auction with bid screening. I start from the average bid auction. Since this format was not characterized earlier, I initially analyze equilibria under a simplified awarding rule and under the hypotheses of the classical independent private value paradigm (Theorem 1). Then, I extend the result to the more complicated average bid rule used in Italy (Lemma 2).

The simplified awarding rule, which I will refer to as the *Florida average bid auction*, states that *(i)* the bid closest to the average of all bids wins, *(ii)* ties of winning bids are broken with a fair lottery and *(iii)* the winner is paid his own bid to complete the project. The model is identical to the classical independent private value paradigm by imposing the following parameter restrictions: *(i)* $p_H = p_L = \infty$ *(no defaults), (ii)* $\varepsilon = y = 0$ *(no uncertainty and no common cost element)* and *(iii)* $F_{X_H} = F_{X_L} = F_X$ *(symmetric bidders).*

When $N = 2$, for any pair of bids both bidders are equally distant to the average. Thus, for both bidders to bid the reserve price R is the unique equilibrium. Theorem 1 deals with the more interesting case of $N > 2$.

Theorem 1: *For any $N > 2$, the strategy profile in which all players bid according to the common constant bid $\xi \in [\overline{x}, R]$ is a symmetric BNE. Moreover, four properties characterize any other symmetric BNE that might exist. The continuous bidding function $b(x)$: (i) is weakly increasing, (ii) is flat at the bottom, (iii) has all types lower than the highest cost one bidding strictly more than their own cost and (iv) the probability of a bidder not bidding $\xi \in [\overline{x}, R]$ is arbitrarily small for N large enough. (Proof in appendix)*

To understand this theorem, consider first the special case where $\overline{x} = R$. Clearly, a flat bid function equal to R is an equilibrium: By unilaterally deviating a single bidder certainly loses. Instead, by bidding R this bidder has one out of N chances of winning and making a profit. I cannot prove that when $N > 2$ this equilibrium is unique. However, the four properties described in the second part of Theorem 1 indicate that any other equilibrium that might exist is approximately a flat bidding function. Moreover, simulation results indicate that the lower bound of this bidding function (property *(iv)*) rapidly converges to R as the number of bidders increases. The intuitive explanation is that as N grows large the chance of a bidder drawing a high cost and offering a high bid increases enough to induce the other bidders to revise their bids upward.

When the reserve price is not binding, $\overline{x} < R$, a multiplicity of equilibria exists: Every constant bid function taking a value in $[\overline{x},\, R]$ is an equilibrium. Similarly, the main model with potential defaults also admits equilibria where all bidders offer a common bid: Such equilibria exist for any bid comprised between R and the expected cost of the (ex ante) least efficient bidder. Thus, the *Florida average bid auction* has equilibria that entail both a random allocation across all bidders and have high winning prices. Both motives make a default less likely than in the FPA.

These properties also characterize the more complex awarding rule used in Italy which I now turn to explain. The Italian average bid auction, which I refer to simply as average bid auction (or ABA), determines the winner as follows: Disregard the top and bottom 10 percent of the bids; calculate the average of the remaining bids (call it A1); then calculate the average of all the bids strictly above the disregarded bottom 10 percent and strictly below A1 (call this average A2); the first price above A2 wins. Ties of winning bids are broken with a fair lottery and the winner is paid his own price to complete the work.[4]

Lemma 2. *In the unique equilibrium: Bids equal R; type H bidders never default; type L default only if the contract cost exceeds R by more than their penalty p_L. (Proof in appendix)*

When all bidders offer R, no individual bidder can deviate without being excluded with certainty by the 10 percent trimming of the lower bids. Moreover, this bidding function is the only one compatible with an equilibrium because of nuances in how tails trimming works: Even when all bids are identical but less than R, an individual bidder who deviates to R wins with probability one and earns the highest possible payoff. The reason being that a bid equal to R will be disregarded in the calculation of A1 and A2, but will then be the closest bid strictly above A2. The more technical discussion is left for the appendix.

The relevance of Lemma 2 is in showing how the ABA can limit defaults by both inducing a random lottery across bidders and inducing a high winning price that makes defaulting less likely. Indeed, an appropriately high R prevents defaults altogether. However, both the high price and the inefficient allocation might be a source of concern for the public authority awarding the contract. The second mechanism that I consider addresses these two problems.

C. Alternative Auction Format II: First Price Auction with Screening

The second mechanism is an FPA augmented by bid screening. By this I mean that after having received the bids, but before awarding the contract, the auctioneer can learn bidder types. To maintain model simplicity, I assume that screening entails: *(i)* perfectly learning the type and *(ii)* disqualifying all bidders that have a positive default probability. Hence, by Assumption (i) this auction essentially becomes a FP auction involving only the n_H type H bidders who never default. Thus, the unique equilibrium bidding function is:[5]

[4] Details on how the rules deal with other types of bid ties and special cases are presented in the appendix.

[5] Equation (2) is standard. To complete the equilibrium, the behavior of the type L bidders needs to be specified. Since type L cannot win, different strategies are compatible with equilibrium. I consider the case in which all type L bidders bid a_H and default when the contract is costly.

$$b(x) = a_H + x + \frac{\int_x^{\overline{x}} (1 - F_X(u))^{n_H - 1} du}{1 - F_X(x)^{n_H - 1}} \tag{2}$$

Bidders offer a price that equals the sum of their commonly shared expected cost, their private cost and a markup term that becomes smaller as the number of competitors increases. Under this equilibrium, the FPA with screening prevents defaults without eliminating competition. However, when comparing this format to the ABA, it is necessary to factor in the auctioneer's bid screening cost. In applications, this cost entails at least the cost of the administration's engineers analyzing bid justifications and of lawyers defending the decision to eliminate a firm. Depending on the amount of the screening cost, the auctioneer's expected revenues under the FPA with screening may or may not exceed those under the ABA. Thus, a revenue comparison between these two formats leads to an ambiguous result.[6]

In terms of allocative efficiency, however, their comparison is more conclusive. In equilibrium, the ABA is equivalent to a random lottery. Therefore, if the same set of bidders were to bid in the two formats, the ABA would be more inefficient. The exact size of this inefficiency, however, crucially depends on the firm cost structure: If the cost that firms face is mostly driven by their commonly observed cost, y, then the inefficiency will be limited. In contrast, strong variations in the private cost component, x, imply that ABAs are particularly wasteful. In the structural analysis that follows, I separately estimate the commonly and privately observed cost distributions and, hence, quantify the potential inefficiency associated with the widespread use of ABAs in the Italian public procurement.

II The Market

The market that I analyze is that for the execution of public work contracts awarded by counties and municipalities (which I will refer to as public administrations or PAs) in the North of Italy. In particular, I focus on road construction and maintenance contracts, which

[6] A formal proof of this statement is presented in a working paper version of this study. That version also shows two other revenue comparisons. First, the FPA without screening dominates the FPA with screening when the screening cost is high and the cost of a default is low. Second, augmenting the ABA with bid screening produces a mechanism whose revenues are strictly inferior to those of the FPA with screening.

represent about a quarter of all public works procured (both in terms of the value of the contracts and in terms of the number of auctions held). This market exhibits at least four of the key elements characterizing the stylized model of section 2. First, firms face cost uncertainty when bidding because contracts are fixed price and their total cost will be fully observed only about 10 months after bidding.[7] Second, the cost of default is known to differ across bidders. This is due to at least two forms of firm heterogeneity: *(i)* their distance to the PA holding the auction, which matters because the standard punishment for defaulting entails the exclusion of the firm for one year from the auctions of the specific PA with which the default occurred; *(ii)* their subscribed capital, which is a proxy for the maximum amount that a PA can obtain as a compensation for the damages incurred because of a default.[8]

Third, there is a reserve price that, although formally binding, is non-binding in practice. Indeed, this reserve price is set using formulas that substantially overestimate the cost of the contract. The result is that the discounts offered often exceed 50 percent of the reserve price and, on average, equal 31 percent. Moreover, an aspect that will be of particular importance is that the administrations must use the same set of formulas to compute the reserve pice. Thus, for a given work the reserve price has to be the same regardless of whether the ABA or the FPA is used.

The fourth element linking the market to the model is the usage of both ABAs and FPAs with screening. The procurement of public works in Italy is almost entirely conducted through sealed bid auctions.[9] A few differences exist between different PAs and over time, due to changes in the regulations. However, in essence the steps needed to award a contract are as follows: First, the administration releases a call for tenders that illustrates the contract characteristics, including the reserve price and the awarding rule (for instance ABA or FPA

[7]Bids are submitted about 4 months before the work begins, and then the work lasts for about 6 months.

[8]The distance between the bidder and the PA (measured at the zip code level) exhibits a strong variation across the firms bidding in these auctions. Its average is 78 miles, while the standard deviation is 134 miles. Similarly, subscribed capital has a mean of €538,000 and a standard deviation of €4 million.

[9]The system is extensively described in Decarolis, Giorgiantonio and Giovanniello (2010) and Decarolis (2013).

with screening).[10] Then every firm qualified to bid[11] for public contracts can submit its sealed bid consisting of a discount over the reserve price. Finally, bids are all opened at the same time. If the awarding rule is the ABA rule, then the winner is selected in the way described in the previous section.

When the awarding rule is the FPA with screening, the PA's engineers first assess the reasonableness of the bids received. The process proceeds sequentially: If the lowest bid is considered reliable, then the contract is awarded to this bidder and no additional bids are screened. If instead, the lowest bidder is judged unreliable, an administrative procedure commences, during which the firm is requested to present justifications for its low price. The process entails a series of steps at the end of which the PA can either eliminate the lowest bid and move on to screen the next bid, or can accept the explanations received and award the contract to this firm. The highest discount is excluded in about 10 percent of the FPAs within my sample because it fails this screening. Since FPAs must be conducted with screening, in the rest of the paper I refer to these FPAs with screening only as FPAs.

In Italy, these two auction formats are especially important in limiting default risk because the letter of credit that is used as bid guarantees typically only covers around 20 percent of the contract value. In contrast, the Miller Act in the US mandates that the winning bidder post a 100 percent performance bond that guarantees the execution of the contract by a third party, the *surer*, in case of a default. Nevertheless, the relative importance of ABAs and FPAs has shifted through time: ABAs have been the most frequently used format since their introduction in 1998. Indeed, between January 1998 and June 2006 the ABA was the mandatory format to award contracts with a reserve price below (approximately) €5 million. Contracts totaling in worth about €10 billion per year were auctioned off through ABAs in this period. After June 2006, however, a reform mandated by the European Union reduced the relevance of ABAs: First, the usage of ABAs was made voluntary.

[10] In addition to ABAs and FPAs, negotiated procedures and scoring rule auctions can be used. In this study, I will disregard these latter two procurement methods. Thus, my results do not necessarily extend to contracts of small economic value (below €300,000), for which negotiations are typically used, and to contracts involving projects of high technical complexity, for which scoring rule auctions are typically used.

[11] Pre-qualification criteria are based on very mild quantitative requirements about the financial viability of the firm and on lack of mafia charges for any person connected with the firm ownership and management structure. They are assessed every three years.

Then, between November 2008 and May 2011 the ABA was forbidden for contracts above €1 million. After that, ABAs were once again allowed for contract worth up to €5 million.

The main reason for the alternation between ABAs and FPAs is that both system have problems that the regulations have been unable to fix. In particular, the main complaints about ABAs regarded the emergence of collusion. Indeed, the fact that bidder payoffs were linked to an easily manipulable trim mean induced firms to form groups coordinating their bids to pilot the contract allocation.[12] In Turin in 2003 a major collusion episode involving 95 firms triggered a local reform mandating a switch from ABAs to FPAs for all contracts awarded by both the county and municipality of Turin. The central government opposed this local reform as a violation of the national law. However, by 2006 both the emergence of other similar collusion episodes in other cities and the victory of Turin against the opposition to its reform before the European Court of Justice lead to the national reform described above.

Nevertheless, the process of switching toward FPAs encountered strong opposition within the PAs. All the reforms failed to account for the severe cost this switch imposed on PAs given the highly decentralized nature of the procurement process (which takes place at the level of single municipalities) and the mandate for in-house bid screening. This cost was lower for the largest PAs that had both engineers to conduct the screening and legal teams to face the appeals of excluded bidders in court. The cost was instead substantial for the smallest PAs which opposed the ban on the usage of ABAs for contracts above €1 million introduced in 2008 and obtained the ban lift in May 2011. Subsequently, even some large PAs, including the county of Turin, returned to ABAs to speed up the procurement process and avoid the delays from the rigid bid screening protocol.

III Data

The data consist of ABAs and FPAs held between 2000 and 2013 by counties and municipalities in five Northern regions (Piedmont, Lombardy, Veneto, Emilia and Liguria). All contracts involve road construction and maintenance and have a reserve price below €5 mil-

[12] This manipulation is similar to that of the trim mean determining LIBOR emerged in the 2012 scandal.

lion. In the rest of this section, I describe two different datasets. The first one is a short panel of 892 ABAs and 338 FPAs for which I observe all bids submitted. I use these data to describe a few key features of bidding under the two formats. The second dataset, which I refer to as the Main data, covers 2,063 FPAs and is the dataset used for the structural estimation.

A. Panel Data

The auctions in this dataset are a subset of those in the Main data for which I was able to obtain the entire set of bids submitted in each auction.[13] Their importance is in identifying certain marked differences between ABAs and FPAs. As the summary statistics in Table 1 illustrate, the two auction formats clearly differ in terms of winning bids. The discount that the winning bidder offers relative to the reserve price is on average 36.67 percent in FPAs, while it is only 13.75 in ABAs. In a paper complementary to this one, Decarolis (2013) establishes that for the county and municipality of Turin the causal effect of the switch from ABAs to FPAs in 2003 is a statistically significant increase of the winning discount that ranges between 6 percent and 12 percent. That study, however, also shows that the switch to FPAs worsens contract performance (in terms of both increased delays in job completion and cost overruns) unless the PA intensively screens bids (by devoting more days to the evaluation of firms' ability to honor their low offered price). This evidence confirms the ambiguous ranking in terms of revenues of the two formats, which crucially hinges on the cost of bid screening, which is unobservable and particularly hard to measure.

As regards the efficiency of the two formats, the data strongly suggest that ABAs resemble random lotteries. The allocation, however, typically differs from the perfectly fair and random lottery implied by Lemma 2. To illustrate this point, I report in Figure 1 the entire set of bids offered in two ABAs from the dataset. The discount (over the reserve price) that was offered is reported on the vertical axis, while the horizontal axis lists all bidders in increasing order of their discount. One auction has 25 bidders and has bids represented by a circle, while the other has 26 bidders and has bids represented by a diamond. For both auc-

[13] There is no centralized system collecting this type of information. The data were manually extracted into a spreadsheet from scanned pdf copies of the auction outcomes released by the single PAs. The pdf were purchased from Telemat spa a company that sells them to firms interested in bidding for public contracts.

tions, I indicate with a square the winning bid. Despite strong similarities in the contracts auctioned off, the bid patterns look remarkably different. Bidding in the 25-bidder auction resembles the case described in Lemma 1 with all bids extremely close to a zero discount. The 26-bidder auction, instead, shows two plateaus, one around a discount of 3 percent and one around a discount of 6 percent, plus six bids markedly higher than all others.

A second study complementary to this one, Conley and Decarolis (2012), shows that the pattern of this latter auction is representative of what is found in a large share of ABAs and is due to the presence of groups of cooperating firms that coordinate their bids to pilot the average that determines the winner. This type of behavior is considered collusion by the Italian criminal law and its discovery by the judiciary has led, for instance, to convictions for 95 firms involved in the Turin case studied in Conley and Decarolis (2012). The evidence from various known collusion cases indicates that a bidding pattern like that observed for the 26-bidder auction in Figure 1 is likely the result of two competing cartels, one trying to manipulate the average discount upward and one downward, in an environment with a few non-colluded bidders offering the intermediate discount of 6 percent. Indeed, independent bidders typically all offer very similar discounts which are PA-specific and are approximately equal to the historical modal winning discount in the auctions of that PA. Market participants sometimes refer to these modal bids as *focal bids*. Hence, the allocation resembles an unfair lottery at a price close to the focal bid.

The statistics in Table 1 confirm that bids within an ABA are typically very concentrated. For instance, in ABAs the average difference between the winning discount and the next highest discount, a quantity often referred to as money left on the table, is only 0.46, while it is 4.89 in FPAs. Similarly, both the average bid range and the average within auction bid standard deviation are almost twice as large in FPAs relative to ABAs. Finally, bid rigging in ABAs partially accounts for their high bidder turnout. Although the lower winning discount in ABAs is one likely reason for why this format attracts a higher number of bidders than FPAs (on average 55.67 vs. 8.56), a second reason is that firms often illegally create shill firms. A shill is a firm that, despite being from a legal perspective like any other firm, exists only because the original firm wanted to be able to submit multiple bids at the auction to enhance its chances of guessing the average (or to rig this average).[14]

The above discussion has two main implications for the analysis that follows. The first is that there is no clear mapping between the bids observed in ABAs and firm costs. Hence, the structural estimation of firm costs will be performed exclusively using FPAs. The high winning discounts in these auctions, which are on average 31 percent of the reserve price and often exceed 40 percent, suggest that collusion is not a concern in these auctions. The second implication, however, is that collusion is an important element for ABAs. Thus, the efficiency comparison that I will conduct will consider not only the theoretical benchmark of a fair lottery, but also the case of the unfair lottery induced by collusion.

B. Main Data

The Main data consist of 2,063 FPAs. In contrast to the panel data, for each auction I only observe the winning bid, together with other auction characteristics like the reserve price. Table 2 reports summary statistics for these data dividing the auctions by the number of bidders. The most frequently occurring number of bidder participating in an auction is six (194 auctions), followed by the case of eight participants (184 auctions). All other cases of bidders comprised between 2 and 10 appear with similar frequencies ranging from 120 to 147 auctions. The table shows that the distribution of both the winning bid and the reserve price (both expressed in €100,000) does not appear to vary systematically as the number of bidders grows. This suggests that we can ignore considerations related to endogenous entry driven by differences in the reserve price. Moreover, an additional indication that these data are qualitatively consistent with the implications of the FPA bidding model comes from the values reported in the column labeled MLT (i.e., money left on the table). I calculate this quantity using the subset of auctions for which I also observe a second classified bid. The values reported are the difference (in €100,000) between the winning bid and the next lowest bid. The declining pattern of this quantity, which monotonically declines as the number of bidders increases from 2 to 10, is what we should expect to observe if bidders were competing more aggressively with more bidders participating.

[14]The law forbids the use of shills asking that every bid emanates from a single "decisional center." In practice, it is often hard for the PAs to establish which firms are de facto sharing the same decisional center.

In the structural analysis that follows, I conduct the estimation using the subsample of the Main data consisting of the 194 auctions with six bidders participating. In the final section, however, I evaluate the robustness of the results using the subsample of auctions with eight bidders. The estimation procedure described next requires that for each auction both the winning bid and the reserve price are observed.

IV Structural Analysis

This section illustrates how to separately estimate the commonly observed and idiosyncratic components of firm costs using the Main data. The method proposed is a variation of Krasnokutskaya (2011) that works for datasets of first price auctions where the only bid observed is the winning bid, but the researcher also observes the reserve price.

A. Empirical Model of FPA Bidding

To map the theoretical model to the Main data, I assume that the observed bids originate from bidders behaving according to the bid function described in Equation (2). I drop from the data any bid excluded through bid screening and consider all remaining bids as coming from type H bidders. Moreover, since the only bid that I observe is the winning bid, $\bar{b}_w$, it is useful to rewrite Equation (2) specifically for the winning bid as follows:

$$\bar{b}_w = a + x_{(n:n)} + \frac{[1 - F_B(b_w)]}{(n-1) f_B(b_w)}, \tag{3}$$

where a is the commonly observed expected cost (defined as $y + \theta\varepsilon$ in section 2), $x_{(n:n)}$ is the lowest private cost draw among the n bidders, b_w is the bid that this bidder would have made if the commonly observed cost a had been equal to zero and F_B and f_B are, respectively, the cumulative and probability density functions of the equilibrium bid conditional on $a = 0$. This formulation follows from the well known inversion approach of Guerre, Perrigne and Vuong (2000) and is convenient because it expresses costs only in terms of bids and bid distributions.

To further link the theoretical model to the data, I make two statistical assumptions (the notation follows the convention of denoting random variables with capital letters and their realizations with lower case letters):

Assumption (i): *The reserve price, R, is a random variable equal to the sum of the commonly observed component of firm costs and an idiosyncratic shock Z, $R = A + Z$.*

Assumption (ii): *The cost and reserve price components are independently distributed according to the joint probability distribution function: $Pr(Z < z_0, A < a_0, X_1 < x_{10}, ..., X_1 < x_{N0}) = F(z_0, a_0, x_{10}, , x_{N0}) = F_Z(z_0)F_A(a_0)\Pi_{i=1}^{N} F_X(x_{i0})$, where F_Z, F_A and F_X are the marginal distributions of the shock Z, the commonly observed cost, A, and privately observed cost, X. The supports of these three marginal distributions are, respectively, $[\underline{z}, \overline{z}]$, $[\underline{a}, \overline{a}]$ and $[\underline{x}, \overline{x}]$ with $0 < \underline{z} < \overline{z} < \infty$, $0 < \underline{a} < \overline{a} < \infty$, $0 < \underline{x} < \overline{x} < \infty$. The distributions of Z and X are continuously differentiable and strictly positive on the interior of their supports.*

Assumption *(i)* serves to link the reserve price to one of the quantities that are the object of the estimation, the commonly observed cost. Since I observe only the winning bid, it would be impossible to distinguish whether a high winning bid is due to a high A or to a high X unless for the same auction another variable conveying information about A is observable. Assumption *(ii)* states the independence of Z, A and X which, together with the additive separability structure of both the reserve price and firm costs and the differentiability of the distributions ensures the applicability of the following identification argument.

B. Identification

In essence, the Main data allow us to separately identify the two firm cost components because the variation of the winning bid and reserve price across auctions identifies the distribution of the common cost, while their within-auction variation identifies the private cost. A formal proof is presented in Krasnokutskaya (2011) and is built upon the idea of treating the common cost component A as auction-specific unobserved heterogeneity.[15]

[15]This idea builds on the work of Li and Vuong (1998) on measurement error. Li, Perrigne and Vuong (2000) were the first to introduce it into auctions, but Krasnokutskaya (2011) extended their method making it suitable for more general cases of bidder asymmetry. Contrary to this paper, all these studies consider environments where multiple bids are observable in each auction. Roberts (2009), instead, considers a similar environment where only the winning bid and the reserve price are observable to the econometrician.

The explanation of how identification is achieved is as follows. First note that, as shown by Equation 2 and 3, the separability of firm costs is preserved in equilibrium. Thus, the winning bid, $\overline{B}_w$, can be written as $\overline{B}_w = A + B_w$, where by B_w I indicate the winning bid conditional on the common cost A being equal to zero. The pair $(\overline{B}_w; R)$ can therefore be thought of as a pair of convolutions $(A + B_w; A + Z)$. Since by Assumption *(b)* the idiosyncratic cost X is independent of A and Z, then B_w, which is a nonlinear transformation of X, is independent of A and Z. Independence and additive separability permit the application of a deconvolution result due to Kotlarski (1966), which leads to the separate identification of the characteristic functions of A, B_w and Z subject to a location normalization.[16] Then, Fourier transformations permit identifying the three marginal probability density functions of A, Z and B_w from their characteristic functions. Finally, once the pdf of B_w is recovered it can be used to simulate a sample of pseudo-winning-bids which, in turn, identify the distribution of the private cost X as the well known result of Guerre, Perrigne and Vuong (2000) shows.

It is important to stress that the choice of the most appropriate method to deal with unobserved auction heterogeneity crucially hinges on both the data and the institutions governing the market. In this application, the availability of data is such that the Krasnokutskaya (2011) approach is infeasible. However, her method might be preferable when all bids are observed since it does not require making assumptions on the nature of the reserve price. As regards the market institutions, section 3 explained that the reserve price from the sample auctions is not set in an attempt to maximize the auctioneer revenues by strategically excluding some bids. Indeed, despite the estimation not imposing a non-binding reserve price, the estimates reveal that it is non-binding in more than 95 percent of the the simulated FPAs.[17] In different applications, the reserve price and the bids might be linked in ways that do not allow the implementation of this approach. However, other variables

His study shows how to separately estimate the common and private cost components in English auctions through a control function method which requires different assumption than the ones used in this study.

[16] The normalization that I use is $E(B_w) = 0$, but other normalizations would be possible.

[17] I estimate all distributions from the subsample of 6-bidder auctions and then use the estimate to simulate 1,000 FPAs. I consider a successful aspect of the model that in more than 95 percent of these simulated FPAs the reserve price is non-binding. In the analysis that follows, I drop all those auctions where the reserve price is binding.

could be used for that. For instance, in the US the auction datasets released by the DoT of many states report the engineers' project cost estimate. This quantity might work well with the proposed method because it is both linked to firm costs and is non binding for bidders.

C. Estimation

The Main data consist of m auctions for which $(n_i, \bar{b}_{wi}, r_i)_{i=1}^m$ are recorded: n_i is the number of bidders, b_{wi} is the winning bid and R_i is the reserve price. Consider fixing a subset of auctions with $n_i = n_0$. The estimation method, which closely follows that of Krasnokutskaya (2011) and Asker (2010), mimics the logic of the identification argument and consists of the following two-step procedure:

Step 1: Estimation of the probability density functions of A and B_w. The first task is estimating the joint characteristic function of a winning bid and the reserve price. This is done non parametrically using the empirical analogue of the joint characteristic function:

$$\hat{\psi}(t_1, t_2) = \frac{1}{m} \sum_{j=1}^m \exp(it_1 \bar{b}_{wj} + it_2 r_j),$$

where i denotes the imaginary number. Then, the deconvolution result of Kotlarski (1966) is exploited together with the normalization and independence assumptions to estimate the characteristic functions of A, Z and B_w:

$$\hat{\phi}_A(g) = \exp \int_0^g \frac{\partial \hat{\psi}(0,t_2)/\partial t_1}{\hat{\psi}(0,t_2)} dt_2$$

$$\hat{\phi}_{B_w}(g) = \frac{\hat{\psi}(g,0)}{\hat{\phi}_A(g)} \quad \text{and} \quad \hat{\phi}_Z(g) = \frac{\hat{\psi}(0,g)}{\hat{\phi}_A(g)}$$

Finally, the estimated probability density functions of A, B_w and Z are obtained through an inverse Fourier transformation.[18]

Step 2: Estimation of the probability density function of X. This step involves simulating a sample of size M of pseudo-winning-bids, B_w^s, from the estimated density of B_w. A rejection

[18]More specifically, these densities are estimated as: $\hat{g}_u(q) = (2\pi)^{-1} \int_{-T_u}^{T_u} d_{T_u}(t) \exp(-itq) \hat{\phi}_u(t) dt$ where $u \in \{A, B_w, Z\}$, where d_{T_u} is a dumping factor that reduces the problem of fluctuating tails. This factor is constructed as in Krasnokutskaya (2011) so that $d_{T_u}(t) = 1 - (|t|/T_u)$ if $|t| < T_u$ and zero otherwise. The smoothing factor T_u should diverge slowly as m goes to infinity to ensure uniform consistency of the estimators. The choice of T_u employs a grid search with a starting point found as in Diggle and Hall (1993) and a termination value that minimizes of the integrated absolute error, $\int |f(x) - \hat{f}(x)| dx$, where the densities in the integral are those of the bid data and the simulated bid data. I end up with $T_A = 2$ and $T_{B_w} = .4$.

method is used for this task.[19] These simulated winning bids are distributed according to the same distribution that would govern equilibrium winning bids in an environment with no unobserved heterogeneity and costs distributed according to the F_X that we seek to estimate.

Therefore, the standard procedure of Guerre, Perrigne and Vuong (2000) can be applied to this sample of simulated bids. This entails first nonparametrically estimating the cdf and pdf of B_w^s.[20] Then, these distributions of the lowest bid are converted into the parent distributions of all bids, $\hat{F}_B$ and $\hat{f}_B$, using properties of order statistics. Finally, substituting these latter two estimates for the cdf and pdf of all bids into Equation (3), implies that for every simulated winning bid we can use equation Equation (3), with A set to zero, to calculate the corresponding simulated winner's cost x_w^s. Finally, with this sample of simulated lowest costs it is possible to proceed as done for B_w^s to non parametrically estimate the relative cdf and pdf, and then to convert them into the corresponding parent distributions $\hat{F}_X$ and $\hat{f}_X$.

V Results

A. Baseline Estimates

For the baseline estimates, I use a subsample of FPAs in the Main data where six firms bid. The estimates of the distributions of the commonly observed and idiosyncratic cost components are shown in Figure 2. Since the location of the two distributions is indetermined, I plot the distributions fixing the lowest bound of both of their supports at zero. This figure reveals that the distribution of the common cost is characterized by more variation than the one of the private cost. Under the additively separable structure of total cost and

[19] In practice, this step requires knowing the support of the distribution because the deconvolution estimator is imprecise at the distribution tails. I estimate these bounds using the following procedure: First, to estimate the length of the support of B_w I use the maximum difference between the winning and the least qualified bid, across all auctions in the sample used for the estimation. The least qualified bid is observable for most of the auctions as the AVCP collects this datum. The length of the support of A is the difference between the support of the bids and that of B_w. For the estimation, the support of B_w is initially centered at zero. If f_{B_w} turns out to be perfectly symmetric around zero, no further adjustments are needed. Since in my estimates f_{B_w} is not symmetric, I shift its support until the mean of the recovered distribution is zero.

[20] This is accomplished using the empirical analogue for the cumulative density function of B_w^s: $\hat{F}_{B_w^s}(b_w^s) = \frac{1}{M}\sum_{j=1}^{M} 1(B_{wj}^s \leq b_w^s)$. The kernel estimator: $\hat{f}_{B_w^s}(b_w^s) = \frac{1}{M}\sum_{j=1}^{M} \frac{1}{h_g}[\frac{35}{32}(1-(\frac{B_{wj}^s - b_w^s}{h_g})^2)^3 1(|\frac{B_{wj}^s - b_w^s}{h_g}| < 1)]$ with bandwidth $h_g = (M)^{\frac{-1}{6}}(2.978)(1.06)(St.Dev.(B_w^s))$ is used to estimate the probability density of B_w^s.

the cost components independence, the total cost variance is the sum of the variances of the two cost components. Thus, the estimates imply that the variation in the common cost component alone explains 94 percent of the total cost variation. This estimate is rather close to that of 86 percent found for the same industry in the US by Krasnokutskaya (2011).

Although this finding indicates that most of the variation is due to common costs, the variation in the private cost is non negligible, as was already suggested by the summary statistics on the within auction bid dispersion. The implication is hence that inefficiencies can arise if contracts are allocated through the ABA. For a more in depth exploration of this inefficiency, I use the cost estimates to simulate allocations under three different scenarios capturing the main features of the observed ABAs.

B. Efficiency Comparison

Starting from the estimated cost distributions, I first create a 1,000 simulated set of FPAs. Since the average number of bidders across all FPAs in the Main data is 7, each simulated FPA consists of 7 draws from the distribution of X. The seventh lowest draw is taken as the cost of the winner. The average cost of the winner across the 1,000 simulations is the FPA efficiency benchmark against which I compare the performance of the ABA.

I consider three scenarios for the ABA and present the associated findings in Table 3. In the first scenario, for each of the 1,000 simulations I use the same seven draws used for the FPAs, but select at random one of the seven draws taking it to be the cost of the winner. The average winner's cost across the simulations is the average cost of a counterfactual in which ABAs replace FPAs and bidders behave according to the equilibrium in Lemma 2. As a first measure of inefficiency, I consider the percentage difference between the costs of the winner in an ABA and the winner in the corresponding FPA. As shown in Table 3, on average the cost of the winner in the ABA is 38.3 percent higher than the cost of the winner in the corresponding FPA. The second inefficiency measure that I consider is the share of auctions in which the ABA selected a winner with a cost strictly above that of the winner in the corresponding FPA. Since in this first scenario there is one out of seven chances that the ABA allocates to the lowest bidder, the share of inefficiently allocated auctions is 86.3 percent. This is mechanically true because in this counterfactual every bidder has one out of seven chances of winning, however, this second measure becomes more interesting for the following counterfactuals.

The second scenario that I consider acknowledges the fact that ABAs exhibit higher participation than FPAs. Therefore, for each of the 1,000 simulations, I add to the original 7 draws plus 66 new draws for a total of 73 bidders. Then, I calculate the cost of the winner in each auction by drawing at random among these 73 costs. As Table 3 shows, the first inefficiency measure remains essentially the same, 40.6 percent. However, the second performance measure improves with 74.4 percent of the ABAs selecting a winner whose cost is above that of the corresponding FPA. This happens because the set of bidders out of which the ABA selects the winner is a superset of that of the FPA bidders. Thus, the ABA can randomly select a bidder whose cost is lower than that of each FPA bidder. Nevertheless, these estimates are best interpreted as a lower bound on the inefficiency of the ABA because for all draws I use the same distribution estimated from the FPAs. The bidders that select into these highly competitive FPAs are likely the most efficient firms and the idiosyncratic cost distribution in this subgroup gives a very conservative estimate of the potential cost dispersion among the less homogeneous and more inefficient ABA bidders.[21]

The third scenario captures how bidders' cooperation in ABAs interacts with the efficiency of the allocation. Since I cannot rely on an equilibrium characterization for this scenario, I calibrate the simulation using parameters from the ABAs involved in the large collusion case that lead to the conviction of 95 firms in Turin that are studied in Conley and Decarolis (2012). In the 276 ABAs that were presented in the court case, out of the 73 bidders participating on average, 43 were non-cooperating firms, while the remaining 30 belong to groups of cooperators. The six groups into which these 30 firms are divided have size: 11, 6, 6, 3, 2 and 2. I evaluate two sub-scenarios: In the first one, which I refer to as the "fair lottery" case, all bidders have an ex ante probability of $1/73$ of winning. In the second one, the "unfair lottery" case, the probability of winning in each group equals the relative winning frequency of this group in the data. That is, a winning probability of 36 percent,

[21]The data shows that two proxies for firm costs, the firm capitalization and its distance from the job location, are have average values such that the firms bidding in FPAs seem significantly more efficient than those bidding in ABAs.

13 percent, 10 percent, 2 percent, 4 percent and 1 percent, where the order goes from the largest to the smallest cartel.

I seek to characterize a lower bound for the inefficiency of the ABAs in this environment. Therefore, I assume that the groups are efficient in the sense that all group bidders have the cost of the lowest cost bidder in their group. Moreover, within each group all firms are assumed to have the same probability of winning. For each of these two cases, I use the cost drawn for the second counterfactual assigning at random the participation of firms into groups. Table 3 reports the results of this analysis. As expected, group bidding alleviates the ABA inefficiency. In particular, in the fair lottery case the amount of extra cost of the ABA relative to the FPA reduces to 14.4 percent. Nevertheless, the share of auctions that select a winner with a higher cost than in the FPA remains high, 72.2 percent. The reduction of the inefficiency is even stronger in the unfair lottery case: The extra cost of the ABA declines to 10.5 percent, and the share of inefficiently allocated auctions declines to 50.4 percent. More in detail, the share of auctions in which the winner of the ABA has a cost strictly below that of the FPA rises from 25.5 percent in the case of groups and fair lottery to 44.6 percent in the case of groups and unfair lottery. This happens because in this latter scenario the largest groups are highly likely to win and, conditional on winning, is assumed to give the contract to its most efficient member. However, both because of the selection argument discussed above and because the allocation might not be perfectly efficient within groups, these figures are best interpreted as a lower bound on the inefficiency of ABAs with collusion.

C. Robustness

To assess the reliability of the above results, I conducted a series of robustness checks which are summarized in Table 4. The first robustness check consisted of adopting the original method of Krasnokutskaya (2011). For the baseline estimates, I preferred to use the alternative method based on the reserve price because the sample size that I can use to apply Krasnokutskaya (2011) is very small: There are only 34 auctions that have exactly two bidders bidding and for which I observe both bids. Under this caveat, I report these estimates in the second row of Table 4. Among the different counterfactual scenarios, the average winner cost in the ABA is substantially above that in the corresponding FPA: This inefficiency ranges from a maximum of 81 percent for the scenario where the set of bidders in the two formats is identical, to 23 percent in the case with cooperating firms and uneven winning probabilities. These estimates are larger than the baseline ones but, as shown by the following row, this seems mostly driven by the use of a different sample and from the different estimation method. Indeed, the third row shows that the estimates are similar when I use the method based on the reserve price with the sample of 143 2-bidder auctions for which both the winning bid and the reserve price are observed. [22]

A second set of robustness checks consisted of restricting the auctions to those held by the county and municipality of Turin. As mentioned earlier (and described at greater length in Decarolis (2013), these two administrations were the first to switch from ABAs to FPAs. Indeed, they hold the majority of the FPAs in the sample. Focusing exclusively on their auctions increases the homogeneity of both the auctions and the set of bidders, thus making the empirical model closer to the simple theoretical model which does not account for heterogenous bidding behavior in auctions held by different administrations. Estimating firm costs using again the subsample of 6-bidder auctions, but restricting the attention to the 133 auctions held by the Turin administrations, I obtain estimates close to the baseline ones: As shown in the fourth row of Table 4, the inefficiency ranges from 59 percent to 13 percent. However, in the following row of the table, I report that the inefficiency is larger when I account for an important form of heterogeneity across administrations: Their average generosity in renegotiating the contract. In particular, I assume that all bidders know ex ante that the administration will concede a cost renegotiation equal to the average renegotiation observed in the sample.[23] Under this perfect foresight assumption, the sum of the observed winning bids and the expected renegotiation can be used to replace $\overline{B}_w$ in the procedure

[22] An additional robustness check on the estimation method entailed replacing the Li and Vuong (1998) deconvolution estimator used in step 1 of the estimation procedure with Bonhomme and Robin (2010) "generalized deconvolution" estimator. I do not report these additional estimates because they are very similar to the ones presented above. Moreover, the generalized deconvolution estimate should not have advantages in my setup since the assumption of finite supports of the distributions (not required for this estimator, but required by Li and Vuong (1998)) does not seem restrictive when estimating distribution of firm costs. Furthermore, the small sample size may be problematic for this estimator given its slow asymptotic rates of convergence.

[23] Using the values reported in Decarolis (2013), the average renegotiation is 13.94 percent of the reserve price for the municipality of Turin and 6.66 percent for the county of Turin. I assume that all firms expect the same renegotiation because renegotiations are not significantly associated with any firm observable characteristic.

described above. Under this formulation the lowest bound of all the inefficiency estimates across the different scenarios increases to 18 percent relative to 11 percent in the baseline case. Finally, as an additional check to control for the potential bias due to observed auction heterogeneity, I follow the convention of "homogenizing" the auctions by first running an OLS regression of the winning bids (and reserve price) on observable auction characteristics (i.e., dummy variables for the year and type of job and administration) and then using the regression residuals to perform the analysis described earlier. The last row of Table 4 reports the estimates obtained after the homogenization. They are in the same ballpark of the other estimates in the table and, in particular, the lowest bound on the inefficiency is 13 percent across the counterfactual scenarios.

VI Conclusions

This paper analyzed two auction formats often used to award public work contracts. Their theoretical comparison revealed that both formats might help an administration to reduce the risk of a winner's default relative to a conventional first price auction. The two mechanisms have an ambiguous ranking in terms of revenues, but the FPA with screening dominates the ABA in terms of efficiency. Using a dataset of Italian procurement auctions for public works, I estimate that the most conservative estimates on the amount of this inefficiency range from 11 to 41 percent. The wide range of estimates were derived under different counterfactual scenarios. Given that ABAs are used to award about €6 billion per year, even the most conservative estimate suggests that they induce a major efficiency loss.

The conclusion is therefore that the usage of ABAs in Italy should be reduced. Moreover, this study suggests that in the numerous other countries where ABAs are used, this procurement method should undergo a careful assessment of its costs and benefits. On the other hand, this study very strongly points out that the limits of the alternative solution represented by the FPA with bid screening. A policy suggestion for the Italian case would be to centralize the screening process to make it cost effective even for the small administrations procuring few contracts per year. More generally, this study stresses the usefulness of future empirical research on which are the most effective methods to procure public contracts.

VII Appendix: Proofs

A. Proof of Lemma 1

To prove Lemma 1, I first introduce the following lemma:

Lemma A.1. *Assuming* $(\frac{\overline{x}_L - \underline{x}_L}{1-\theta}) < \varepsilon < (\frac{1}{\theta})y$, *there exist two values,* p_H^* *and* p_L^*, *such that whenever* $p_H \geq p_H^* = y - \theta\varepsilon$, *bidders of type H neither play* $b_H < y + \underline{x}_H + \theta\varepsilon$ *nor ever default; and whenever* $p_L < p_L^* = (1-\theta)\varepsilon + \underline{x}_L - \overline{x}_L$ *bidders of type L fulfill their bids only if the realized cost of the project is low.*

Proof of Lemma A.1: That a bidder i of type H always fulfills his bid when $p_H \geq p_H^*$ follows from the observation that this would be a dominated strategy. Suppose he defaults if the project is costly, i.e. $p_H \leq (y + x_i + \varepsilon) - b_H$, then his payoff in case of victory is:

$$\begin{aligned}(1-\theta)(b_H - (y + x_H)) - \theta p_H \ &\leq (1-\theta)((y + x_H + \varepsilon) - p_H - (y + x_H)) - \theta p_H \\ &= (1-\theta)\varepsilon - p_H \\ &\leq (1-\theta)\varepsilon - (y - \theta\varepsilon) \\ &= \varepsilon - y < 0\end{aligned}$$

Therefore, given that this bidder never defaults, his expected cost in case of victory is $y+x_H+\theta\varepsilon \geq y+\underline{x}_H+\theta\varepsilon$ so that bidding anything below $y+\underline{x}_H+\theta\varepsilon$ generates a negative payoff in case of victory and is thus strictly dominated by bidding $b \geq y + x_H + \theta\varepsilon$. As regards the second part of the lemma, notice that in an FPA no bid is higher than $y + \overline{x}_L + \theta\varepsilon$ (following a simple Bertrand argument). Therefore, if $p_L < p_L^*$ a type L bidder will always default when the project is costly because $p_L < (1-\theta)\varepsilon + \underline{x}_L - \overline{x}_L \leq (y + \underline{x}_L + \varepsilon) - b_L$.

Having proved Lemma A.1, I can now turn to prove Lemma 1. For the existence of a pure strategy monotone equilibrium, following Lemma A.1, we only need to show that the dual auction, de Castro and de Frutos (2010), of the procurement auction under assumption (i) satisfies all the assumptions of the existence theorem in Reny and Zamir (2004), RZ from now on. The dual auction is defined by action $\tilde{b}_j^i = (v_j + r_j - b_j^i)$, signal $\tilde{x}_j^i = (v_j + r_j - (x_j^i + s_j))$, and the payoff function $\tilde{u}_j^i = \tilde{x}_j^i - \tilde{b}_j^i$,where $v_j = \overline{x}_j + s_j$, $r_j = \underline{x}_j + s_j$, and s_j is equal to a_H or a_L depending on whether the bidder is type H or L.

RZ-Assumption 1 (utility function): define $r_L \equiv a_L + \underline{x}_L$ and $r_H \equiv a_H + \underline{x}_H$ and let $l \in [0, \min(a_H + \underline{x}_H, a_L + \underline{x}_L)]$. Then the bid space conforms to that of RZ: $B_j \in \{l\} \cup [r_j, \infty)$. Moreover, notice that in the dual auction formulation $\tilde{u}_j^i = \tilde{x}_j^i - \tilde{b}_j^i$. This payoff function is: (i) measurable, it is bounded in $[\underline{x}_j, \overline{x}_j]$ for each $\tilde{b}_j^i$ and continuous in $\tilde{b}_j^i$ for each $\tilde{x}_j^i$; (ii) define $b^* \equiv \max(a_H + \overline{x}_H, a_L + \overline{x}_L)$, then $\tilde{u}_j^i(\tilde{b}_j^i, \tilde{x}_j^i) < 0$ for all $\tilde{b}_j^i > b^*$ and for any $\tilde{x}_j^i \in [\underline{x}_j, \overline{x}_j]$; (iii) for every bid $\tilde{b}_j^i \geq r_j$, I have that $\tilde{u}_j^i(\tilde{b}_j^i, \tilde{x}_j^i)$ is constant in $\tilde{x}_j^{-i}$ and strictly increasing in $\tilde{x}_j^i$; (iv) $\tilde{u}_j^i(\overline{b}_j^i, \tilde{x}_j^i) - \tilde{u}_j^i(\underline{b}_j^i, \tilde{x}_j^i)$ is constant in $\tilde{x}$.

RZ-Assumption 2 (signals): assume that the private value $\tilde{x}$ is a monotonic function $x : [0,1]^N \to [\underline{x}, \overline{x}]^N$, then the assumption that signals are independent implies that signals' affiliation weakly holds and that for any $\tilde{x}^i$ the support of i's conditional distribution does not change with the other signals. Since Assumption 1 and 2 are satisfied, existence follows.

Having assured existence, the rest of Lemma 1 follows from de Castro and de Frutos (2010).

B. Proof of Theorem 1

The fact that the strategy profile in which all bidders offer the maximum bid equals R is an equilibrium is clear: a unilateral deviation leads to a zero probability of winning as opposed to having probability 1/N of winning a non negative amount. When N=2 this is the unique symmetric BNE. Although I cannot rule out the presence of other symmetric BNE, I can characterize four properties that they must have. The last property implies that for a large enough N all equilibria approximate flat bid functions.

Property 1: Non Decreasing Function. The proof is by contradiction. Assume that the BNE bidding function, b, has an interval over which ti is strictly decreasing. Take two types, x_1 and x_0, with $x_1 > x_0$ such that $b(x_1) < b(x_0)$. Then by b being BNE it follows that:

$[b(x_1) - x_1] \Pr(win|b(x_1)) \geq [b(x_0) - x_1] \Pr(win|b(x_0))$ and

$[b(x_0) - x_0] \Pr(win|b(x_0)) \geq [b(x_1 - x_0)] \Pr(win|b(x_1))$.

Therefore from the first and from the second inequalities I have respectively that:

$$\begin{aligned}\Pr(win|b(x_1)) &\geq \{[b(x_0)-x_1]/[b(x_1)-x_1]\}\Pr(win|b(x_0)) \\ &\geq \{[b(x_0)-x_1]/[b(x_1)-x_1]\}\{[b(x_1)-x_0]/[b(x_0)-x_0]\}\Pr(win|b(x_1))\end{aligned}$$

The above implies: $[b(x_1)-b(x_0)][x_1-x_0] \geq 0$, which is a contradiction.

Property 2: Non Strictly Increasing Function at the Bottom. This property significantly distinguishes the ABA from the FPA: under the stated assumptions no BNE can have the lowest cost type bidding the lowest bid. If $\underline{x}$ is the lowest type and, by contradiction, it is assumed that the equilibrium bid is such that $b(\underline{x}) = \underline{b} < b(x)\ \forall x \neq \overline{x}$ then it is easy to show that a unilateral profitable deviation exists. For instance, for a small $\delta > 0$ a bidder can deviate bidding: $b(x)$ for any $x \neq \overline{x}$ and $\underline{b}+\delta$ for $x = \underline{x}$. His expected revenues are unchanged for any $x \neq \overline{x}$ and they are strictly higher for $x = \overline{x}$ since the probability of winning goes from being zero to being positive. By property 1 and continuity we must have that the bidding function is flat at the bottom.

Property 3: Cost Shading. This property is standard in auction models with imperfect information. Clearly any strategy profile requiring a bidder to bid below its cost is strictly dominated and cannot be an equilibrium. Moreover, for any strategy profile requiring some type, x', below the highest cost type to bid $b(x') = x'$, it is easy to construct a unilateral profitable deviation by picking a small $\delta > 0$ and modifying his strategy exclusively for $b(x') = x' + \delta$. His expected revenues are unchanged for any $x \neq x'$ and they are strictly higher for $x = x'$ since in case of victory his payoff goes from being zero to being strictly positive while the probability remains positive.

Property 4: Restriction on the Lowest Equilibrium Bid. I look at the lowest type, $\underline{v}$, such that for all $x \in [\underline{x}, \underline{v}]$ bidding some constant $\underline{b}$ (the flat bottom of Property 2) with $\underline{v} < \underline{b}$ gives no unilateral incentive to deviate to a higher bid. Hence, assume b^* is a symmetric BNE that is weakly increasing and such that $b^* = \underline{b}$ if $x \leq \underline{v}$. Then, if agent N draws $\underline{v}$ it must be that: $u(\underline{v}, \underline{b}, b^*_{-N}) \geq u(\underline{v}, b, b^*_{-N})$ for any $b > \underline{b}$. That is: $\Pr(win|\underline{b})[\underline{b}-\underline{v}] \geq \Pr(win|b)[b-\underline{v}]$ for any $b > \underline{b}$. The event that $\underline{b}$ wins occurs when $\underline{b}$ is the bid closest to the average bid, conditional on all other players playing b^*. It is useful to define the following probabilities:

$$p \equiv \Pr[(X_1 \leq \underline{v}) \cap (X_2 \leq \underline{v}) \cap \ldots \cap (X_{N-1} \leq \underline{v})],$$

$q_1 \equiv \Pr[(X_1 > \underline{v}) \cap (X_2 \leq \underline{v}) \cap (X_3 \leq \underline{v}) \cap ... \cap (X_{N-1} \leq \underline{v})]$,

...

$q_{N-2} \equiv \Pr[(X_1 > \underline{v}) \cap (X_2 > \underline{v}) \cap ... \cap (X_{N-2} > \underline{v}) \cap (X_{N-1} \leq \underline{v})]$,

$\alpha_M \equiv \Pr[|\underline{b} - \frac{1}{N}\sum_{r=1}^{N} b_r^*| < |b(x_j) - \frac{1}{N}\sum_{r=1}^{N} b_r^*|$ for any $x_j > \underline{v}$ and $j = 1, ..., M\ |q_M = 1]$,

where M=1,...,N-2. I can now rewrite $\Pr(win|\overline{b})$ as: $\Pr(win|\overline{b}) = p(\frac{1}{N}) + [q_1(\frac{1}{N-1}) + q_2\alpha_2$ $(\frac{1}{N-2}) + ... + q_{N'}\alpha_{N'}(\frac{1}{N-N'})]$, where N' is $(\frac{N}{2} - 1)$, or the closest lower integer if N is odd.

Whenever there is at least one bidder drawing a valuation strictly bigger than $\underline{v}$ then the average bid will be strictly bigger than $\underline{b}$. Therefore I can always take a $b' > \underline{b}$ but ε-close to $\underline{b}$, such that conditional on having at least one player drawing $x > \underline{v}$, b' leads to a probability of winning strictly greater than $\underline{b}$. Moreover the payment in case of victory with the bid b' is strictly less than that in case of winning with $\underline{b}$. Define β_M as follows:

$\beta_M \equiv \Pr[|b' - \frac{1}{N}\sum_{r=1}^{N} b_r^*| < |b(x_j) - \frac{1}{N}\sum_{r=1}^{N} b_r^*|$ for any $x_j > \underline{v}$ and $j = 1, 2, ..., M\ |q_M = 1]$, where M=1,2,...,N-2. Therefore I can now rewrite $\Pr(win|b')$ as $\Pr(win|b') = [q_1 + q_2\beta_2$ $+ ... + q_{N-2}\beta_{N-2}]$. Now, given the way b' was chosen, it must be that $[q_1 + q_2\beta_2 + ... + q_{N-2}\beta_{N-2}][b' - \underline{v}] \geq [q_1 + q_2\alpha_2 + ... + q_{N'}\alpha_{N'}][\underline{b} - \underline{v}]$. The left hand side of this inequality is exactly $u(\underline{v}, b', b^*_{-N})$. A necessary condition for b^* to be an equilibrium is $\{p(\frac{1}{N}) + [q_1(\frac{1}{N-1}) +$ $q_2\alpha_2\ (\frac{1}{N-2}) + ... + q_{N'}\alpha_{N'}(\frac{1}{N-N'})]\}[\underline{b} - \underline{v}] \geq [q_1 + q_2\alpha_2 + ... + q_{N'}\alpha_{N'}][\underline{b} - \underline{v}]$. Hence, it must be that $p \geq Nq_1(\frac{N-2}{N-1})$, which can be rewritten using the definitions of p and q_1 as:

$F(\underline{v})^{N-1} - N(\frac{N-2}{N-1})[(1 - F(\underline{v}))F(\underline{v})^{N-2}] \geq 0. \quad (*)$

Therefore, considering the left hand side of the above inequality as a function of $\underline{v}$, say $g(\underline{v})$, then only the values of $\underline{v}$ such that $g(\underline{v}) > 0$ satisfy the necessary condition. The function $g(\underline{v})$ starts at 0 for $\underline{v}$ equal to $\underline{x}$ and converges toward 1for $\underline{v}$ equal to $\overline{x}$. Moreover with $N > 2$ the function has a unique critical point, a minimum that is attained at the value of $\underline{v} = z$, where z is the (unique) value such that the following equation is satisfied:

$F(z) = 1 - \frac{2N^2 - 4N + 1}{N^3 - N^2 + 1}$.

Since the denominator is larger than the numberator with F absolutely continuous, z must always exist. Therefore $g(\underline{v})$ starts at 0, decreases until it reaches a minimum value and then converges to 1 from below. Hence it must be that $g(\underline{v})$ crosses zero from below just once so that the only values of $\underline{v}$ for which $(*)$ is satisfied are those that lie in $(v^*, \overline{x}]$ where v^* is defined to be the value of $\underline{v}$ such that the inequality of $(*)$ would be an equality. Moreover since $v^* < \overline{x}$ the following is true: *For any (absolutely continuous)* F_X *and* $\forall\, \delta > 0$, $\exists\ N^*_{\delta,F}$ *such that* $\forall N \geq N^*_{\delta,F}$ *the following is true:* $|\underline{v}_{\delta,F} - \overline{x}| < \varepsilon$.

To see why this is the case, consider that by the definition of v^* the values of $\underline{v}$ such that $(*)$ holds are the ones for which $g(\underline{v}) > g(v^*) \longrightarrow \underline{v} > v^*$ because g is strictly increasing until $z > v^*$. However the expression defining z is such that, in the limit for N that goes to infinity, $z = \underline{x}$. Therefore it must be the case that also v^* and hence $\underline{v}$ go to $\overline{x}$ as N goes to infinity. Therefore there is always an $N^*_{\delta,F}$ that for any F and for any $\delta > 0$ it is large enough so that the difference between $\underline{v}$ and $\overline{x}$ is less than δ. Finally one can see that using $(*)$ as a threshold for checking that any symmetric BNE must have a highest bid strictly lower than v^* is very conservative: As N grows above 3, the actual maximum bid might be substantially lower than this bound. However, given the very high concavity of $(1 - F(\underline{v}))^{N-1}$ this is not likely to reduce the usefulness of this bound because as N grows the bound reduces the size of the interval $(v^*, \overline{x}]$ very rapidly by bringing v^* closer to $\overline{x}$. Therefore even for small N, v^* will be close to $\overline{x}$. This is the reason why, even for small N, $(*)$ gives a bound that is useful.

C. Proof of Lemma 2

Before proving Lemma 2, I report here how the awarding rule deals with all the special cases that can arise. First, if all prices are equal, the winner is selected with a fair lottery. Second, if there are no prices strictly below $A1$ and above the disregarded bottom 10 percent of prices, then the lowest price equal to or higher than $A1$ wins. Third, a random draw is used to ensure that exactly 10 percent of the top/bottom prices 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 I ignore them since this never occurs in the data.

To prove Lemma 2, notice that an argument identical to that used in the proof of Theorem 1 implies that any candidate type-symmetric equilibrium must have a flat bottom. However, contrary to the *Florida average bid auction*, there cannot be any equilibrium in which this flat bottom is less then R. This follows from the combined effect of the tail trimming and the requirement of the winning price being strictly above $A2$. Indeed, consider a candidate equilibrium where a pair of type-symmetric continuous bidding functions entail a flat bottom below R. Denote the minimum bid of this candidate equilibrium as $\underline{b} < R$. The problem of a bidder i considering deviating from $\underline{b}$ consists of assessing his payoff in two cases: Either all other bidders bid $\underline{b}$ (case 1), or at least one other bidder bids above $\underline{b}$ (case 2).

Under case 1, if bidder i deviates to bid R, then he wins with probability one and earns the highest possible payoff, which is strictly positive since R is non binding. This is because his bid will be the closest from above to $A2$, since in this case $A2 = \underline{b}$. Under case 2, if bidder i does not deviate from $\underline{b}$, he must earn a zero payoff. If, instead, bidder i deviates to a higher bid he earns a weekly higher payoff. Since the flat bottom entails that a mass of bidders bids $\underline{b}$, the argument is because a deviation from $\underline{b}$ to R is always weakly profitable and it is strictly profitable with positive probability.

To conclude the equilibrium description, note that defaults can occur only on the part of L type bidders if the contract cost exceeds R by more than their penalty p_L.

Table 1: Panel Data Summary Statistics

	FPAs				ABAs			
	Mean	SD	P.50	N	Mean	SD	P.50	N
Win Discount	30.67	10.10	31.40	338	13.75	5.22	13.94	892
Number of Bids	8.56	7.27	7.00	338	55.67	46.97	44.00	892
Max-Win Bid	0.66	2.34	0.00	338	3.96	3.30	3.19	892
Win-Second Bid	4.89	5.11	3.83	338	0.46	1.65	0.10	892
Max-Min Bid	18.61	8.92	17.87	338	11.88	5.25	11.60	892
Within Auction SD	6.87	3.11	6.60	338	2.88	1.44	2.71	892

The table reports the statistics for all those ABAs and FPAs for which all bids are observed. The statistics are the mean, the standard deviation, the median calculated across auctions. The number of auctions is reported in the last columns of each panel. *Win Discount* is the winning discount (expressed as a percentage discount over the reserve price). *Number of Bids* is the number of bids admitted to the auction. *Max-Win* is the difference between the highest discount offered and the winning discount. In the FPAs, this quantity is typically equal to zero since the highest discount wins (unless it is eliminated via bid screening). *Win-Second Bid* is the difference between the winning discount and the discount immediately below it. *Max-Min Bid* the within-auction range of all discounts. *Within Auction SD* is the within-auction standard deviation of all discounts.

Table 2: Main Data Summary Statistics

		Winning Bid				Reserve Price				
N.Bids	N.Auct	Mean	SD	Max	Min	Mean	SD	Max	Min	MLT
2	143	4.33	4.12	18.84	0.22	5.91	5.94	28.69	0.28	0.76
3	124	5.47	6.16	30.10	0.55	7.47	8.48	41.17	0.76	0.74
4	141	4.09	3.78	20.76	0.78	5.69	5.14	28.42	1.06	0.72
5	149	6.51	7.54	38.33	0.52	9.25	10.24	48.77	0.79	0.70
6	194	4.70	4.14	26.29	0.53	6.98	6.12	44.44	0.66	0.68
7	135	5.29	4.69	26.59	0.62	7.91	6.91	33.97	0.75	0.68
8	184	5.47	5.38	26.80	0.52	8.30	7.98	38.27	0.71	0.66
9	143	6.07	7.76	31.39	0.36	9.35	11.98	48.89	0.48	0.67
10	120	4.74	4.83	23.37	0.42	7.47	7.65	33.43	0.95	0.65

The table reports summary statistics for the sample of FPAs, dividing it into nine subsamples that differ in the number of bids admitted. The first columns reports the number of bids, the baseline estimates are obtained using the largest subsample. This is the subsample of 6-bidder auctions which has 194 auctions, as shown in the second column. The next two sets of four columns report summary statistics for the winning bid and the reserve price. Both are expressed in €100,000 and are transformed in real 2003 euro by adjusting for the yearly inflation rate. The last column, MLT (money left on the table) is the difference (in €100,000) between between the winning bid and and the next lower bid. This latter variable is calculated using only the subset of auctions that belong to the panel data where I can observe the lowest non-winning price.

Table 3: Efficiency Comparison

	No.Bids	ABA Groups	ABA Lottery	Winner Cost	Cost Ineff.	Share Ineff. Auct.
Baseline FPA	7	No	-	6.45	0%	0%
Scenario 1:						
Same Bidders	7	No	Fair	8.91	38.3%	86.3%
Scenario 2:						
Higher Entry	73	No	Fair	9.06	40.6%	74.4%
Scenario 3:						
Groups & Fair Lottery	73	Yes	Fair	7.37	14.4%	72.2%
Groups & Unfair Lottery	73	Yes	Unfair	7.12	10.5%	50.4%

The first row contains the values for the benchmark FPAs, while the following rows describe the counterfactual scenarios for the ABA. The estimated cost distributions are used to simulate 1,000 auctions. The column *No.Bids* reports the number of bidders in each auction. The following column states whether bidder groups are considered. The next column indicates whether in the lottery used to simulate the allocation of the ABA all bidders have the same probability of winning or not. The columns *Winner Cost* reports the average winner cost across the simulations. The column *Cost Ineff.* reports the first measure of inefficiency: the average of the (percentage) difference between the cost of the winner in the ABA and the winner in the corresponding FPA. The column *Share Ineff. Auct.* is the second measure of inefficiency: the share of auctions in which the winner designated by the ABA has a cost strictly above that of the winner in the corresponding FPA.

Table 4: Robustness Checks

	Scenario 1	Scenario 2	Scenario 3 Fair Lottery	Scenario 3 Unfair Lottery
Baseline Estimates	38.3%	40.6%	14.4%	10.5%
Krasnokutskaya ($N = 2$)	80.6%	78.7%	32.6%	23.9%
Subsample with $N = 2$	65.9%	67.2%	25.2%	19.7%
Turin Administrations	56.6%	58.5%	16.8%	13.0%
Turin & Renegotiations	59.6%	73.5%	24.9%	17.7%
Auction Homogenization	55.0%	57.9%	19.8%	12.9%

The first row reports the baseline estimates for the first measure of inefficiency presented in Table 3. The next row shows the analogous estimates using different estimation methods and subsamples: the second row uses the method of Krasnokutskaya (2011), the third uses the reserve price-based method, but with a sample of 2-bidder auctions, the fourth uses only the 6-bidder auctions held by the county and municipality of Turin, the fifth accounts for renegotiation in the winning bids of the Turin auctions, the sixth uses a sample homogenized bids and reserve prices.

Figure 1: Two ABAs in the Panel Dataset

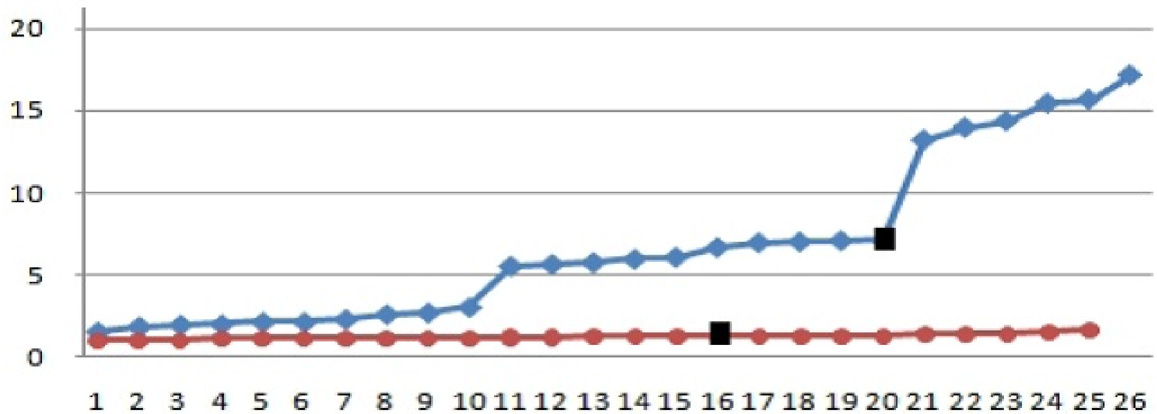

The figure plots all the bids offered in two ABAs in the panel dataset. The bids are reported in terms of discount over the reserve price and are sorted in increasing order of the discount. Each discount offered is denoted as a circle for the 25-bidder auction and as a diamond for the 26-bidder auction. For both auctions, however, the winning bid is denoted with a square. The auctions were selected to be similar along various observable characteristics: the year of the auction, the geographical location of the auctioneer, the object of the contract and the number of bidders.

Figure 2: Estimates of the Cost Components Distributions

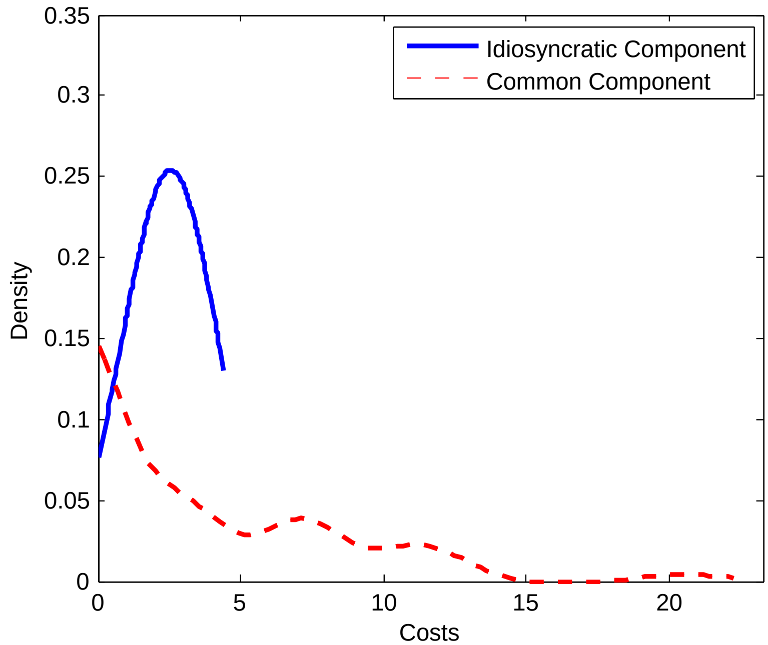

The figure plots the mean of the estimates of the distribution of the common and idiosyncratic cost components. The support are shifted so that zero is the minimum of the support for both distributions.

References

Asker, John. 2010. “A Study of the Internal Organization of a Bidding Cartel.” *American Economic Review*, 100(3): 724–762.

Athey, Susan, Dominic Coey, and Jonathan Levin. 2013. “Set Asides and Subsidies in Auctions.” *American Economic Journal: Microeconomics*, 5(1): 1–27.

Athey, Susan, Jonathan Levin, and Enrique Seira. 2011. “Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions.” *Quarterly Journal of Economics*, 126(1): 207–257.

Bajari, Patrick, Stephanie Houghton, and Steve Tadelis. 2007. “Bidding for Incomplete Contracts: An Empirical Analysis.”

Board, Simon. 2007. “Bidding into the Red: A Model of Post-Auction Bankruptcy.” *Journal of Finance*, 62(6): 2695–2723.

Bonhomme, Stephan, and Jean Marc Robin. 2010. “Generalized Deconvolution with an Application to Earnings Dynamics.” *Review of Economic Studies*, 77: 491–533.

Burguet, Roberto, Juan-Jos Ganuza, and Esther Hauk. 2009. “Limited Liability and Mechanism Design in Procurement.” Barcelona Graduate School of Economics Working Papers 383.

Calveras, Aleix, Juan-Jose Ganuza, and Esther Hauk. 2004. “Wild Bids. Gambling for Resurrection in Procurement Contracts.” *Journal of Regulatory Economics*, 26(1): 41–68.

Che, Yeon-Koo, and Jinwoo Kim. 2010. “Bidding with Securities: Comment.” *American Economic Review*, 100(4): 1929–35.

Chillemi, Ottorino, and Claudio Mezzetti. 2009. “Procurement Under Default Risk: Auctions or Lotteries?”

Conley, Timothy G., and Francesco Decarolis. 2012. "Detecting Bidders Groups in Collusive Auctions."

Decarolis, Francesco. 2013. "Awarding Price, Contract Performance and Bids Screening: Evidence from Procurement Auctions." *American Economic Journal: Applied Economics*, forthcoming.

Decarolis, Francesco, Cristina Giorgiantonio, and Cristina Giovanniello. 2010. "The Awarding of Public Works in Italy: An Analysis of the Mechanisms for the Selection of the Contractors." *Bank of Italy - Quaderni di Economia e Finanza*, 83.

de Castro, Luciano I., and Maria-Angeles de Frutos. 2010. "How to translate results from auctions to procurements." *Economics Letters*, 106(2): 115–118.

Diggle, Peter, and Peter Hall. 1993. "A Fourier Approach to Nonparametric Deconvolution of a Density Estimate." *Journal of the Royal Statistical Society, Series B (Methodological)*, 55(2): 523–531.

Engel, Andreas R., Juan-Jos Ganuza, Esther Hauk, and Achim Wambach. 2006. "Managing Risky Bids." In *Handbook of Procurement.* , ed. Nicola Dimitri, Gustavo Piga and Giancarlo Spagnolo. Cambridge University Press.

European Commission. 2002. "Prevention, Detection and Elimination of Abnormally Low Tenders in the European Construction Industry." *Enterprise and Industry Report.*

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica*, 68(3): 525–574.

Ioannou, Photios G., and Sou-Sen Leu. 1993. "Average-Bid Method. Competitive Bidding Strategy." *Journal of Construction Engineering and Management*, 119(1): 131–47.

Kotlarski, Ignacy. 1966. "On Some Characterizations of Probability Distributions in Hilbert Spaces." *Annali di Matematica Pura e Applicata*, 74: 129–134.

Krasnokutskaya, Elena. 2011. "Identification and Estimation in Highway Procurement Auctions under Unobserved Heterogeneity." *Review of Economic Studies*, 78(1): 293–327.

Krasnokutskaya, Elena, and Katja Seim. 2011. "Bid Preference Programs and Participation in High- way Procurement Auctions." *American Economic Review*, 101(6): 2653–2686.

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

Li, Tong, and Quang Vuong. 1998. "Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators." *Journal of Multivariate Analysis*, 65: 139–165.

Li, Tong, Isabelle Perrigne, and Quang Vuong. 2000. "Conditionally independent private information in OCS wildcat auctions." *Journal of Econometrics*, 98(1): 129–161.

Marion, Justin. 2007. "Are Bid Preferences Benign? The Effect of Small Business Subsidies in High- way Procurement Auctions." *Journal of Public Economics*, 91(7-8): 1591–1624.

Reny, Philip J., and Shmuel Zamir. 2004. "On the Existence of Pure Strategy Monotone Equilibria in Asymmetric First-Price Auctions." *Econometrica*, 72(4): 1105–1125.

Rhodes-Kropf, Matthew, and S Viswanathan. 2005. "Financing Auction Bids." *RAND Journal of Economics*, 36(4): 789–815.

Roberts, James W. 2009. "Unobserved Heterogeneity and Reserve Prices in Auctions."

Spagnolo, Giancarlo, Gian Luigi Albano, and Milo Bianchi. 2006. "Bid Average Methods in Procurement." *Rivista di Politica Economica*, 1(2): 41–62.

Spulber, Daniel F. 1990. "Auctions and Contract Enforcement." *Journal of Law, Economics and Organization*, 6(2): 325–44.

Waehrer, Keith. 1995. "A Model of Auction Contracts with Liquidated Damages." *Journal of Economic Theory*, 67(2): 531–555.

Zheng, Charles Z. 2001. "High Bids and Broke Winners." *Journal Economic Theory*, 100: 129–171.