

COMPARING PUBLIC PROCUREMENT AUCTIONS*

BY FRANCESCO DECAROLIS¹

Università Bocconi, Italy

This article contrasts two auction formats often used in public procurement: first price auctions with ex post screening of bid responsiveness and average bid auctions (ABAs), in which the bidder closest to the average bid wins. The equilibrium analysis reveals that their ranking is ambiguous in terms of revenues, but the ABA is typically less efficient. Using a data set 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 and used to quantify the efficiency loss under counterfactual ABAs.

1. INTRODUCTION

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 led to the proliferation of auction formats that deviate from the well-known first price auction (FPA).

This study contrasts from both a theoretical and an empirical perspective such auction formats frequently used in public procurement. The first format consists of supplementing a conventional 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 FPA are common. For example, in the context of the public procurement of roadwork contracts by the California Department of Transportation (DoT), Bajari et al. (2014) report that in 4% 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 United States, 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, Switzerland, and Taiwan. Moreover, its usage has

*Manuscript received December 2013; revised February 2017.

¹ This article supersedes an earlier version circulating under the title "When the Highest Bidder Loses the Auction: Theory and Evidence from Procurement Auctions," which was the main 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. I thankfully acknowledge the many comments received on earlier drafts from Susan Athey, Marianne Bertrand, Peter Cramton, Jeremy Fox, Amit Gandhi, Ken Hendricks, Jean-Francois Houde, Jakub Kastl, Jonathan Levin, Greg Lewis, Gustavo Piga, Rob Porter, Marc Rysman, Jesse Shapiro, and Giancarlo Spagnolo. I also benefited from aid with the data collection of the Bank of Italy and the Italian Authority for Public Contracts. Please address correspondence to: Francesco Decarolis, Department of Economics, Università Bocconi, Via Roentgen N. 1, 20136 Milan, Italy. Phone: 02-5836.5329. Fax: 02-5836.5329. E-mail: francesco.decarolis@unibocconi.it.

also been suggested by both the civil engineering literature (Ioannou and Leu, 1993) and major institutions (European Commission, 2002).

In the first part of this article, 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 FPAs: 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 analyzing 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 facts 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 winner's default, they are not equivalent. In particular, 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, and answering it is essential to understand the adequate functioning of a procurement system.

In the second part of the article, I address this question by analyzing a data set of Italian public procurement auctions held alternately under the ABA or the FPA with screening. This data set, 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 lottery-like nature of the bids offered in ABAs implies that they 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 data sets where the econometrician does not observe all the bids, but observes at least the reserve price, along with the winning bid. 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 cost.

Although the estimated dispersion in the idiosyncratic cost component suggests substantial inefficiencies in ABAs, the counterfactual ABA estimates qualify this effect. They show that ABAs fail to select the lowest bidder in two thirds of the auctions and that the average production cost is one sixth higher than in the FPAs. I arrive at this result after adjusting the counterfactual ABA for the presence of both higher bidder participation and collusion among subsets of the bidders, two phenomena that are shown to be characterizing features of the ABAs in the data. As argued below, these estimates are based on assumptions that make them best interpreted as a lower bound on the inefficiency of ABAs.

This article has three main contributions. The first contribution is to bridge the vast theoretical literature on the perverse effect of FPAs when bidders can default² with the analysis of two alternative formats that are frequently encountered in real-world public procurement.

² An incomplete list of the main studies on FPAs in this context includes Spulber (1990), Waehrer (1995), Zheng (2001), Rhodes-Kropf and Viswanathan (2005), Board (2007), Che and Kim (2010), Burguet et al. (2012), and Chillemi and Mezzetti (2013).

In particular, this article contributes to the literature on auctions by analyzing the ABA. The evidence on its relevance in Italy, along with the description of numerous cases of similar regulations in other countries, points to the importance of understanding this format. The only two previous studies in the literature that analyzed it, Engel et al. (2006) and Spagnolo et al. (2006), characterized its properties, but under restrictive assumptions on the number of bidders and their cost and information. Thus, the equilibrium characterization of the ABA in Theorem 1 is a key result of this article. This result has contributed to make ABAs part of the economics body of knowledge³ and to spur subsequent research aimed at understanding how bidding works in this format. Among these related studies, Conley and Decarolis (2016) consider the same Independent Private Value (IPV) environment of this article, but introduce the possibility for bidders to collude; Chang et al. (2015) analyze the ABA in a common value environment; Galavotti et al. (2015) consider the IPV case, but through the framework of *level-k* thinking; and Eun (2015) analyzes bidding in a closely related procurement format used in Korea. Regarding the FPA with screening, this article is mostly related to the recent work of Bajari et al. (2014), who analyze bidding and adaptation costs in U.S. procurement auctions.

The second contribution is to complement the theoretical comparison of ABAs and FPAs with screening with a quantitative analysis of their performance in a major procurement market. Auctions are typically very persistent institutions so that format changes are rarely observed. Only a few other studies, mostly involving public auctions in the United States, have presented this type of comparative analysis: Athey et al. (2011) compare open versus sealed bid auctions used for the sale of timber harvesting contracts, Lewis and Bajari (2011) compare first price versus scoring rule with time incentive auctions for the procurement of roadwork contracts, and Marion (2007) compares first price versus first price with small-business bid subsidy for roadwork contracts.⁴ For Italian public procurement, two complementary studies, Decarolis (2014) and Branzoli and Decarolis (2015), present a difference-in-differences analysis of the effects of switching from ABAs to FPAs with screening on observable quantities: the winning bid and the ex post contractual performance in the first study and the entry and subcontracting choices in the second study. In this article, I structurally estimate firms' underlying costs from a sample of FPAs with screening and quantify the efficiency loss in counterfactual ABAs.⁵ The findings in the aforementioned papers of Decarolis (2014) and Branzoli and Decarolis (2015), along with the results in Conley and Decarolis (2016) on collusion in ABAs, are used in this article to guide the complex construction of an appropriate counterfactual ABA. Methodologically, the article contributes by showing how the structural estimation can be performed by extending the method of Krasnokutskaya (2011) when the only observed bid is the winning bid, but the reserve price is also observed.⁶

The third contribution concerns the policy implications stemming from the article's findings. The inefficiency estimated for the ABA suggests that its continued use in public roadwork contracts procurement is wasteful. Nevertheless, an effective solution is unlikely to be either a naive adoption of FPAs, because of the risk of costly defaults, or of FPAs with screening, because of the presence of screening costs. Adequate solutions, instead, should involve the simultaneous adoption of an efficient auction format, like an FPA, *and* of effective methods to reduce the default risk, which combine elements of a centralized bid screening system, stricter qualification

³ See the discussion of the ABA from this article presented in chapter 6 of the textbook *Game, Strategies, and Decision Making*, 2nd ed., by Joseph Harrington (Worth Publishers, 2014).

⁴ Krasnokutskaya and Seim (2011) study the same bid preference system studied by Marion (2007), but their data are exclusively from the FPAs with bid subsidy. Athey et al. (2013) study a closely related question analyzing timber auction run alternately with or without set-asides for small business.

⁵ Thus, similarly to Athey et al. (2011, 2013), I only use one of the two auction formats observed in the data to estimate bidder costs.

⁶ The method, also used in Asker (2010), originates from the work of Li et al. (2000). As explained in greater detail below, while both Asker (2010) and Krasnokutskaya (2011) exploit the availability of multiple bids from the same auctions to achieve identification, I will exploit the observability across auctions of one bid along with one additional instrument, which in my application will be the reserve price.

criteria, insurance policies (performance bonds), past reputation, and higher penalties in case of default.⁷

2. THEORETICAL ANALYSIS

This section presents a stylized IPV bidding model. The main result is a characterization of equilibrium bidding in ABAs and its comparison in terms of revenues and efficiency to two models of FPAs: with and without bid screening. I conclude with a discussion of alternative models rationalizing why ABAs and FPAs with screening can outperform conventional FPAs.

2.1. *Baseline FPA.* 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 $c^i = y + x^i$, whereas 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 , whereas 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 $[\underline{x}_j, \bar{x}_j]$, where $0 \leq \underline{x}_j < \bar{x}_j < \infty$.

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.⁸ Thus, disregarding dominated strategies, the expected payoff for a bidder of type $j = \{L, H\}$ bidding b_j can be written as

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

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

ASSUMPTION (1). $\frac{\bar{x}_L - \underline{x}_L}{1 - \theta} < \varepsilon < y$, 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, whereas for type L bidders the penalty is low enough that they always optimally default if the cost is high when the format is an FPA. This greatly simplifies the game by allowing me to write bidder expected payoffs in the FPA conditional on bidding b_j as

$$(2) \quad \begin{cases} [b_H - x_H - a_H] \Pr(\text{win}|b_j), & \text{if the bidder is type } H, \\ [b_L - x_L - a_L](1 - \theta) \Pr(\text{win}|b_j), & \text{if the bidder is type } L, \end{cases}$$

⁷ Spulber (1990) and Calveras et al. (2004) analyze the relative merits of these methods.

⁸ 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 $(1 - \theta)\varepsilon > 0$.

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 nonbinding 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 $b_j : [\underline{x}_j, \bar{x}_j] \rightarrow R_+$ 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' actions. I begin by showing the perverse features of FPAs in this environment. The game described above is isomorphic to an FPA with asymmetric bidders. Thus, under 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 $(\underline{x}_L - \bar{x}_H) < (a_H - a_L) < (\bar{x}_L - \underline{x}_H)$, despite type Hs 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.*

The restriction that $(\underline{x}_L - \bar{x}_H) < (a_H - a_L) < (\bar{x}_L - \underline{x}_H)$ ensures that the supports of type Ls 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 a public procurer, it is evident why alternative mechanisms are often preferred to the FPA for public procurement.

2.2. *Alternative Auction Format I: ABA.* The two alternative mechanisms that I analyze are an ABA and an FPA with bid screening. I start from the ABA. 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. To further simplify the exposition, I will present Theorem 1 for the case of the classical independent private value paradigm, that is, I assume 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 $\xi \in [\bar{x}, r]$ 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 $\xi \in [\bar{x}, r]$ is arbitrarily small for N large enough (see the proof in the Appendix).*

To understand this theorem, consider first the special case where $\bar{x} = r$. 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. Moreover, when the reserve price is not binding, $\bar{x} < r$, a multiplicity of equilibria exists: Every constant bid function taking a value in $[\bar{x}, r]$ is an equilibrium. This characterization extends to the main model with potential defaults: Its equilibria entail all bidders offering a common bid and 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 have both a random allocation across all bidders and a high winning price. 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 explain. The Italian ABA, which I refer to simply as ABA, determines the winner as follows: Disregard the top and bottom 10% 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% 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.⁹

LEMMA 2. *In the unique equilibrium, bids equal r , type H bidders never default, and types L default only if the contract cost exceeds r by more than their penalty p_L (see the proof in the Appendix).*

When all bidders offer r , no individual bidder can deviate without being excluded with certainty by the 10% 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.

2.3. Alternative Auction Format II: FPA with Screening. The last mechanism that I consider is an FPA augmented by bid screening. This is an FPA in which the procurer, after having received the bids, can eliminate all the bids it judges “too good to be true.” This elimination could take many forms, but, for reasons explained below, I assume that the screening process is imperfect as it entails (i) disqualifying all type L bidders, but also (ii) probabilistically penalizing type H bidders. The latter is captured by including in bidder H 's expected profit a term for the expected loss from an otherwise winning bid being eliminated after screening. Similarly to Bajari et al. (2014), I impose a reduced form penalty that measures the skewness of the bid relative to a reference price. Setting the latter to be equal to the auction reserve price is a natural choice

⁹ Details on how the rules deal with other types of bid ties and special cases are presented in the Appendix.

as, in the data application, firms' bid take the form of a discount over the publicly announced r . Hence, after dropping all H and L subscripts given that only H types are relevant in this format, the payoff function of bidder i can be written as follows:

$$(3) \quad \pi^i(b^i, x^i) = (b^i - P(b^i|r) - (y + \theta\epsilon + x^i))\Pr(b^i < b^j, \forall j \neq i),$$

where $P(b^i|r)$ represents the adjustment of expected revenues that I define to be the penalty from screening. $P(b^i|r)$ encompasses what bidders view as the risk of elimination due to the procurer mistrusting bids that are too low relative to r , and I specify it to be

$$(4) \quad P(b^i|r) = \delta(r - b^i),$$

where $\delta > 0$.¹⁰ For reasonably small values of delta, the equilibrium bid function is

$$(5) \quad \beta(x) = \frac{1}{(1 + \delta)} \left\{ y + \theta\epsilon + x + \delta r + \frac{\int_x^{\bar{x}} (1 - F_x(u))^{N-1} du}{(1 - F_x(x))^{N-1}} \right\}.$$

This equilibrium entails an ex ante efficient allocation, but the FPA with screening awards the contract at a higher price than the FPA. Moreover, the higher the reference point, r , and the enforcement parameter δ , the higher the equilibrium prices will be. Indeed, by strategically moving these two parameters, a procurer could achieve a broad spectrum of results. At one extreme, setting δ to zero reduces the game to a conventional FPA, whereas, at the other extreme, a sufficiently high δ produces an equilibrium outcome equivalent to a posted price at r , analogously to what also the ABA produces in equilibrium.

It is useful to conclude the analysis of this format with three remarks on aspects of the penalty function, as applied to my data. First, while it might be desirable in general to allow for the possibility that also type L bidders are only imperfectly screened out, this model cannot be identified from data on only the winning bid of each auction.¹¹ Data and institutional features discussed in the next section, however, make assuming that only types H bid a reasonable approximation for the observed FPAs with screening. Second, if screening is costly, subjecting type H to screening is wasteful for the procurer. Thus, screening once and for all would seem preferable. In the Italian case, however, the next section explains that this is both forbidden by law and unlikely to be ideal since a bidder's risk depends on factors specific to both time-varying bidder financial conditions and bidder-procurer specific pairs. Third, note that screening type H might be desirable in a richer model where they also pose some default risk, albeit smaller than that of type L . However, since in the data defaults are never observed, the only empirically relevant equilibria of this game would nevertheless be those where types H do not default. But this would make this more complex model observationally equivalent to the simpler model above, thus suggesting that the simpler model is sufficient to rationalize the evidence.

2.4. Discussion: Comparison of the Alternative Formats. The comparison of FPAs with screening and ABAs requires factoring 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 profits under the FPA with screening may or may not exceed those under the ABA. Thus, a revenue comparison between these two formats leads

¹⁰ This penalty function is in the spirit of Bajari et al. (2014). Its justification follows their logic with one notable exception. In their paper, they deal with a vector of input prices as the bid object. To avoid corner solutions and allow the first-order condition to represent optimal choice, they require a strictly convex penalty function. However, in the one-dimensional bid space presented here, convexity is not needed, as positive costs and the assumption that δ is small ensure interior optimality.

¹¹ Unless the bidder type is observable to the econometrician.

to an ambiguous result. 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.

3. THE MARKET

The market that I analyze is that for the execution of public work contracts awarded by counties and municipalities (public administrations, PAs) in the north of Italy. I focus on road construction and maintenance contracts, which represent about a quarter of all public works procured (in terms of both the value of the contracts and the number of auctions).

This market exhibits at least four of the key elements characterizing the stylized model introduced above. First, firms face cost uncertainty when bidding because awarded contracts are fixed price, but firms' total cost will be fully observed only 10 months after bidding.¹²

Second, while defaults are possible in principle, none are present in my data. An interesting feature of the data, however, is the presence of price renegotiations (i.e., cost overruns).¹³ Due to certain market regulations, these renegotiations can be revealing of whether the earlier model where only the never defaulting type H bid in FPAs with screening fits the data. In fact, there exist observable differences between firms that can potentially create asymmetries in the default-risk type, but that are—by law—not contractible for the PA: (i) firms' 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) firms' 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. Results exploiting both are reported in the next section.

Third, there is a reserve price that, although formally binding, is nonbinding in practice. This reserve price is set using formulas that greatly overestimate contract costs. Indeed, the discounts offered often exceed 50% of the reserve price and, on average, equal 31% in FPAs. Moreover, an aspect that will be of particular importance is that the administrations must use the same set of formulas to compute the reserve price regardless of whether ABAs or FPAs are 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.¹⁴ A few differences exist across 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.¹⁵ Then every firm qualified to bid 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.¹⁶ If the awarding rule is the ABA rule, the winner is selected following the rule described in the previous section.

¹² On average, bids are submitted four months before the work begins, and then the work lasts for six months.

¹³ Albeit the contracts are formally fixed price, more than half of them have overruns paid by the procurer.

¹⁴ The system is described in Decarolis (2014) and Decarolis and Giorgiantonio (2015).

¹⁵ 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 for which negotiations are typically used and to contracts involving projects of high technical complexity, for which scoring rule auctions are typically used.

¹⁶ Prequalification 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.

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% of the FPAs within my sample because it fails this screening. By law, this process has to be performed for every FPA, and past outcome of the screening cannot be used to decide future screening outcomes. Thus, since all FPAs must be conducted with screening, in the rest of the article, 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% of the contract value. In contrast, the Miller Act in the United States mandates that the winning bidder posts a 100% performance bond that guarantees the execution of the contract by a third party, the *surety*, in case of a default. Nevertheless, the relative importance of ABAs and FPAs has shifted through time: ABAs have been an extensively 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. Then, between November 2008 and May 2011, the ABA was forbidden for contracts above € 1 million. After that, however, 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.¹⁷ 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 led 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.

3.1. *International Comparison.* The interest in the auction formats that I study stems from the fact that similar experiences with ABA and FPA occurred in many countries. To document this fact, I surveyed the public procurement regulations of various countries looking for rules similar to the Italian ABA. Table 1 reports the results. The countries listed in the left-hand column use auctions with awarding rules that identify and automatically eliminate “abnormal”

¹⁷ This manipulation is similar to that of the trim mean determining London Interbank Offered Rate that emerged in the 2012 scandal.

TABLE 1
RULES FOR IDENTIFICATION AND ELIMINATION OF ABNORMAL BIDS

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

NOTES: The classification is based on the rules for public procurement for works, goods, and services that I surveyed. The three columns separate instances where algorithms identifying abnormally low bids are used to (a) automatically eliminate such bids (left column), (b) perform further investigations on the responsiveness of these bids (central column), (c) perform some extra controls on these bids but neither the type of control nor the algorithm is publicly disclosed (right column). This latter case is based on what is reported in Bajari et al. (2014) regarding the use by the California DoT of an algorithm to identify abnormal bids that, however, is not publicly disclosed. To see an example of how a case is classified consider the case of Florida. Its Department of Transportation (DoT) regulation allows for four rules to procure contracts. One of them (subarticle 3-2.1) states that when the bidders are only three or four, the bid closest to the average is selected, but when five or more contractors bid, the low bid and the high bid are excluded, and the bid closest to the average of the remaining bids is selected. Similarly, a major procurement method used in Taiwan entails awarding the auction to the bidder whose bid is the closest to the average of all submitted bids. The Supporting Information Appendix reports a more complete discussion of all the cases listed in this table.

bid(s). Generally, this means that prices lower than some threshold defined as a function of the bids (often their average) are considered unreasonable and are automatically eliminated. The contract is awarded to the lowest noneliminated price. This often implies that the highest bidder who would win in an FPA is instead eliminated with probability one. The exact rules for automatic elimination differ between the various countries as illustrated in the note to Table 1. They range from complicated rules like the Italian one to more straightforward ones that simply award the contract to the bidder closest to the average bid.

The countries in the middle column, instead, present rules mandating the identification of excessively low prices through algorithms. The elimination of such bids is not automatic. Generally, the auction is an FPA augmented by an ex post screening stage in which further checks on the reliability of bids flagged by the algorithm as potentially anomalous are undertaken before the contract can be awarded. For instance, the German Federal Procurement Agency (BESCHA) requires additional explanations from the lowest price bidder whenever his price is more than 20% below the second lowest price. Finally, the last column reports that the California DoT uses an algorithm to identify abnormal bids but it does not disclose it. More details on these cross-county regulations are reported in the Supporting Information Appendix.¹⁸

4. 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 million. In this section, I describe two different data sets. 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

¹⁸ Another interesting aspect of this international comparison is that several countries experienced transitions between ABAs and FPAs similar to that of Italy. In Colombia, for instance, the replacement of FPAs with ABAs recently occurred with law 1150 of 2007. The same happened in the Nagano prefecture (Japan) with a reform that became effective on April 2003 and in the Guangdong region (China) with a series of reforms that started around 2004. Instead, the opposite route was recently taken by Taiwan and Peru that recently reintroduced forms of FPAs after having used ABAs for many years. These cycles suggest that there are failures in both mechanisms that undermine their stability.

TABLE 2
 PANEL DATA SUMMARY STATISTICS

	FPAs				ABAs			
	Mean	SD	P.50	<i>N</i>	Mean	SD	P.50	<i>N</i>
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

NOTES: 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, and 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.


bidding under the two formats. The second data set, which I refer to as the Main data, covers 1,013 FPAs and is the main data set for the structural estimation.

4.1. *Panel Data.* The auctions in this data set are those for which I was able to obtain the entire set of bids submitted in each auction.¹⁹ Their role in the following analysis is twofold. First, they allow me to analyze bid screening. Only in FPAs are bids eliminated via screening, and it is 10% of them. Most cases involve the exclusion of one bid only, but instances of two and three eliminations in the same auction also happen. Albeit the small size of the Panel data is inadequate for the structural estimation by itself, I use it to complement the Main data to relate bids' departure from r with their probability of being screened out.

The second role of the Panel data is to illustrate features of the ABAs and FPAs relevant for the construction of the counterfactuals. As the summary statistics in Table 2 illustrate, the two auction formats clearly differ in terms of both bids and participation. Regarding the former, the discount that the winning bidder offers relative to the reserve price is on average 36.67% in FPAs, while it is only 13.75% in ABAs. In a paper complementary to this one, Decarolis (2014) 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% and 12%. 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.

With regard to the efficiency of the two formats, the data show that ABAs resemble random lotteries. The allocation, however, typically differs from the perfectly fair and random lottery implied by Lemma 2. To exemplify this point, I report in Figure 1 the entire set of bids offered in two ABAs from the data set. The discount (over the reserve price) that was offered is reported on the vertical axis, whereas the horizontal axis lists all bidders in increasing order of their discount. One auction has 25 bidders and has bids represented by a circle, whereas the other has 26 bidders and has bids represented by a diamond. For both auctions, I indicate with a square the winning bid. These auctions share many features: same year, same county (but

¹⁹ 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 copies were purchased from Telemat spa, a company that sells them to firms interested in bidding for public contracts.



NOTES: The figure plots all the bids offered in two ABAs in the panel data set. These auctions were 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), but received very different bids, as the figure shows. 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, the winning bid is denoted with a black square.

FIGURE 1

TWO ABAS IN THE PANEL DATA SET [COLOR FIGURE CAN BE VIEWED AT WILEYONLINELIBRARY.COM]

different municipality), and nearly identical description of the job and number of bidders. The bid patterns, however, 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% and one around a discount of 6%, plus six bids higher than all others.

A related study by Conley and Decarolis (2016) 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.²⁰ 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 noncolluded bidders offering the intermediate discount of 6%. 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. For instance, for each one the two auctions in the example, the other auctions held by the same PA have a similar modal bid. 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. Table 2 confirms 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.

Aside from the differences in the bid distributions between the two formats, the most striking difference revealed by Table 2 is in terms of the number of bids submitted. ABAs have about seven times the number of bids of FPAs. Using the same identification strategy as in Decarolis (2014), Branzoli and Decarolis (2015) find that the adoption of FPAs reduces the number of

²⁰ Bid coordination can be achieved in two ways: (i) through cooperation between separate firms or (ii) through the creation of *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). Both types of behavior are considered “collusion” by Italian criminal law. Since establishing which firms are *shill* is de facto hard, I will also refer to both types of behavior as collusion.

TABLE 3
MAIN DATA SUMMARY STATISTICS

	Mean	SD	P.50	N
Winning price	266.29	143.82	225.50	1013
Final price	388.69	194.08	340.15	1013
Reserve price	393.64	208.83	336.55	1013
Number of bids	7.35	3.84	7.00	1013

NOTES: The table reports summary statistics for the sample of FPAs upon which the main part of the structural estimation is based. The three prices reported (winning price, final price after all renegotiations, and reserve price) are all expressed in € 1,000. All auctions have no more than 17 bids submitted and a reserve price between € 150,000 and € 1 million.

bidders by 35–65 bidders, or at least 50% of the average number of bidders in ABAs. Although the lower winning discount in ABAs is one reason for why this format attracts a higher number of bidders than FPAs, a second reason is clearly the benefit of placing multiple coordinated bids. Moreover, the sharp decline under FPAs is also consistent with the deterrence effect of bid screening on unreliable bidders. Indeed, the data show that bidders located further away from the procurer are less likely to participate in FPAs.

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% of the reserve price and often exceed 40%, suggest that collusion is not a concern in these auctions. The second implication is that both high entry and extensive collusion are important elements for ABAs. Thus, the efficiency comparison that I conduct will consider not only the theoretical benchmark of a fair lottery among the same set of bidders of the FPAs but also the case of higher entry and allocations via an unfair lottery induced by collusion.

4.2. Main Data. The Main data consist of 1,013 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 and the number of bidders. This is a homogeneous set of road work contracts that, in addition to all the restrictions described above, also have a reserve price between 150,000 and 1 million euro and the same number of potential bidders. This homogeneity is of crucial importance for the structural estimation as the model is estimated under the assumption that each auction is the repetition of the same game, with only cost draws changing from auction to auction. In this respect, the choice of the sample is a first-order task. In a highly regulated context like that of public procurement, this entails paying close attention to all the legal details. The sample assembled takes advantage of an extensive analysis of the regulations that allowed me to filter out regions, counties, and even single PAs that at various points in time modified features of their procurement regulations in ways that make them unsuitable for the structural analysis performed here. Details of this regulatory analysis have been documented in a separate paper (Decarolis and Giorgantonio, 2015).²¹

Table 3 reports summary statistics. The reserve price is about € 400,000. Despite the winning price representing a substantial discount over the reserve price, renegotiations typically bring back the final price very close to the original reserve price. The number of bids is 7 on average and has a relatively low standard deviation, equal to 3.8. Based on these statistics, the Main data also look very similar to the Panel data described above.

Two important aspects of the data regard entry and renegotiations. Regarding entry, although the sample of auctions in the Main data exhibits variation in the number of bids submitted, I

²¹ It is important to notice that “homogenization,” another procedure often taken to control for observable heterogeneity and involving a regression stage to purge bids of variation linked to observables, is less satisfactory, as potential model misspecifications will bias the estimates. I defer a more specific discussion of homogenization to the section describing the results.

assume that in all these auctions the number of potential bidders is the same and equal to 17. The assumption of an equal number of potential bidders is based on the similar regulatory requirements that firms have to satisfy to bid in road work contracts where r is between € 100,000 and € 1 million. Moreover, fixing the number of potential bidders at 17 bidders is driven by the drop in the density of the distribution of submitted bids at this value. In the empirical model that follows, I will assume that bidders know the number of potential bidders but not that of actual bids. This is in accordance with the publicity of the regulatory requirements on potential entry and with the prohibition for the PA to disclose who has bid while the bidding process is still open.

Regarding renegotiations, it was mentioned earlier that, due to institutional features, firms' capital and distance to the procurer are a source of asymmetries in the default-risk type that PAs cannot contract upon. Evaluating the association between the extent of contract renegotiation and these proxies for default-risk type is thus an ideal way to check whether more risky firms obtain greater renegotiations. Through a web scraper, I calculate the distance between the bidder and the PA measuring it at the zip code level. This measure exhibits a strong variation across the firms bidding in these auctions: Its average is 78 miles, while the standard deviation is 134 miles. Through the Italian registry of firms (Infocamere), I obtain data on firms' underwritten capital. Subscribed capital has a mean of € 538,000 and a standard deviation of € 4 million. Using these two proxy variables, I evaluate whether, in the FPAs, far away and/or low capital firms are more successful in their renegotiations. However, I fail to find any evidence in support of this hypothesis and interpret this finding as supportive of the idea that FPAs can be characterized through the simple model described earlier where only type H bidders participate. The fact that all bidders appear to be *ex ante* identical in their ability to renegotiate will also allow a convenient way to incorporate renegotiations in the empirical model, as explained below.

5. STRUCTURAL ANALYSIS

This section presents the empirical model used to combine data in the Main and Panel data sets to quantify the inefficiency implied by the use of ABAs. As argued above, data from ABAs cannot be directly used for this exercise. The estimation of the FPAs is based on a slight extension of the model of the FPA with screening presented in Section 2. The extension, entailing adding both renegotiations and entry, is inconsequential in terms of the qualitative properties of the FPAs equilibrium but serves to reduce the gap between the stylized model and the data. The core of the structural analysis regards how to separately estimate the commonly observed and idiosyncratic components of firm costs, given the available data where the only bid observed is the winning bid.

5.1. Empirical Model of FPA Bidding. Consider the model of the FPA with screening of Section 2. First, I extend that model to account for an entry stage. Since in the data only qualified bidders are entitled to bid, I assume that the number of potential bidders, \mathcal{N} , is known and fixed. The previous discussion has argued that ABA and FPA are different in terms of what drives entry: In a typical ABA, firms bid around the *focal bid*. Although this requires very little effort from the firm in terms of learning its project cost, the competitive nature of FPAs requires performing a careful evaluation of future costs. This effort can be modeled as an entry cost. If this cost is known and fixed, and if only firms paying this cost end up learning their production costs and placing a bid, then the number of active bidders, N , is a random variable, $N \sim \text{Binomial}(q)$,²² where q is the probability that a bidder pays the entry cost.

The symmetric equilibrium of this entry and bidding game is characterized by an entry probability and bidding function, $(q^*, \beta(x))$. The bidding function is analogous to that in Equation (5),

²² This implies, using the standard terminology in the literature, that the number of actual (those paying the entry cost) and active (those placing the bid) firms is the same. Li (2005) considers a more complex environment where, due to the presence of a binding reserve price, active bidders are a subset of actual bidders. The literature has analyzed several other types of entry models (see, among others, Marmer et al., 2013).

but with $q^*\mathcal{N}$ replacing N . As shown by Levin and Smith (1994), for the independent private value environment that I analyze the equilibrium entry probability is such that firms are ex ante willing to randomize between entering and not entering.²³

Regarding renegotiations, at the time of bidding, firms have perfect foresight of a , the amount of future renegotiation. However, firms internalize in their ex ante payoff only a fraction of that: $(1 - \alpha)a$. This captures the idea of adjustment/transaction costs such that a share of each dollar successfully renegotiated is (potentially) wasted.²⁴

Under these extensions, the equilibrium of the FPA with screening is such that a winning bid, \bar{b}_w , is linked to the underlying costs y and x according to²⁵

$$(6) \quad \bar{b}_w = \frac{1}{1 + \delta} \left\{ y + x_{(n:n)} + \delta r - (1 - \alpha)a + \frac{[1 - F_B(b_w)]}{(\mathcal{N}q^* - 1)f_B(b_w)} \right\},$$

where y is the commonly observed cost and $x_{(n:n)}$ is the lowest private cost draw among the N bidders. Let me also define b_w as the bid that this bidder would have made if the commonly observed cost y , the overrun, and the penalty for bid skewing had all been set equal to zero. F_B and f_B are the corresponding cumulative and probability density functions of b_w . This formulation follows the logic of the Guerre et al.'s (2000) bid inversion.

To link this theoretical model to the data, I make the following statistical assumptions (I denote 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 , $R = Y + Z$.*

ASSUMPTION (II). *The cost, reserve price and overrun components are independently distributed according to the joint probability distribution function: $\Pr(Z < z_0, Y < y_0, X_1 < x_{10}, \dots, X_N < x_{N0}, A < a_0) = F_Z(z_0)F_Y(y_0)F_A(a_0)\prod_{i=1}^N F_X(x_{i0})$, where F_A, F_Z, F_Y , and F_X are the marginal distributions of the overrun, A , the shock, Z , the commonly observed cost, Y , and privately observed cost, X . The supports of these marginal distributions are, respectively, $[\underline{a}, \bar{a}]$, $[\underline{z}, \bar{z}]$, $[\underline{y}, \bar{y}]$, and $[\underline{x}, \bar{x}]$, with $-\infty < \underline{a} < \bar{a} < \infty$, $0 < \underline{z} < \bar{z} < \infty$, $0 < \underline{y} < \bar{y} < \infty$, and $0 < \underline{x} < \bar{x} < \infty$. 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 Y or to a high X unless for the same auction another variable conveying information about Y is observable. Assumption (ii) states the independence of A, Z, Y , 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.

²³ This entry model is also consistent with the evidence in the Main data showing that the distribution of both the winning bid and the reserve price remains stable as the number of submitted bids grows. Furthermore, institutional features requiring firms interested in bidding to incur costs early on in terms of project documents that need to be purchased and, for some auctions, physical inspections of the worksite that have to be undertaken suggest that a model where firms first decide to incur such costs and only afterward learn their production costs can be a reasonable approximation of the environment. Finally, a relevant source of entry costs in these auctions often stressed by market participants is also represented by the time consumed by dealing with the highly bureaucratized procedures required to prepare the formal documentation compulsory for participation. According to market participants the length of time required to prepare such documentation can entail starting the process even before having fully assessed the production cost of the contract tendered off.

²⁴ This model is in the spirit of Bajari et al. (2014), but their framework is more sophisticated, incorporating detailed data on the types of adjustments taking place during the life of the contract.

²⁵ To simplify the notation, I use here y to denote the term $(y + \theta\epsilon)$ appearing in Equation (5) since the components of this latter term cannot be separately identified.

5.2. *Identification.* The identification is semiparametric, and its main idea is based on results from Krasnokutskaya (2011). Her deconvolution estimator, however, cannot be directly applied to Equation (6) to separately retrieve the distributions of X and Y because of two separate problems. The first is that her method requires observing at least two bids per auction. Here, I show how this can be overcome through the use of the reserve price. The second is that, by applying the deconvolution estimator to a sample of (\bar{b}_w, r) pairs for each auction, identification is possible for the distribution of y , but not for that of x . This is due to the impossibility of separating the distribution of x from that of a after the deconvolution step. Thus, the identification argument must proceed in the reverse order: first transforming the bids, purging them of all elements unrelated to y and x (purging step), and then performing the deconvolution of x and y (deconvolution step).

Purging step: The first step is to get an estimate δ from the data, $\hat{\delta}$. When bidding, firms internalize their expected penalty cost by accounting for the monetary loss of not getting the contract despite having a winning bid, scaled by the likelihood of this event happening. Given the form of the penalty from Equation (4), the equilibrium bid entails a penalty as a monetary value proportional to the distance between the bid and the reserve price. This is driven by the probability of being screened, $Pr(b_{it} \text{ screened out}) = \beta_{Screen}(r_t - b_{it}) + \epsilon_{it}$, so that we can rewrite the penalty function as $P(b|r) = \text{Monetary Penalty} * Pr(b \text{ screened out})$. The identification of the second term on the right-hand side of the equation exploits the fact that in the Panel data I observe all bids, including those that get eliminated through the screening. So from variation in $(r_t - b_{it})$, it is possible to identify the effect of the distance between the bid and the reserve price on the probability of being eliminated. For the first term on the right-hand side, I use a feature of the regulations and set this monetary term to be equal to 2% of r . This is the amount of the monetary fine faced by a firm that is declared the winner, but that refuses to accept the job after the auction takes place.

Once δ is identified, notice that to identify α the equilibrium bid can be expressed as

$$(7) \quad \bar{b}_w(1 + \delta) - \delta r = (1 - \alpha)a + \left[x_{(n:n)} + \frac{[1 - F_B(b_w)]}{(\mathcal{N}q^* - 1)f_B(b_w)} \right],$$

By Assumption (ii), A , Y , and X are independent. Thus, given the identified parameter δ and data on \bar{B}_w , R , and A , variation in A identifies α .

Deconvolution step: Provided with estimates for α and δ and data on renegotiations, for each winning bid in the Main data, I can now obtain what I define to be purged winning bids as $\bar{b}_w \equiv \bar{b}_w(1 + \delta) - (\delta r - (1 - \alpha)a)$. Thus, the equilibrium bidding function becomes

$$(8) \quad \bar{b}_w = y + \left\{ x_{(n:n)} + \frac{[1 - F_B(b_w)]}{(\mathcal{N}q^* - 1)f_B(b_w)} \right\}.$$

The identification of q^* is directly obtained as the sample average N divided by \mathcal{N} . Then, to separately identify x and y , note that the variation of the (purged) winning bid and reserve price across auctions identifies the distribution of the common cost, whereas their within-auction variation identifies the private cost. A proof is presented in Krasnokutskaya (2011) and is built upon the idea of treating the common cost component as auction-specific unobserved heterogeneity.²⁶ The specifics of how this result applies in my context are discussed below.

²⁶ This idea builds on the work of Li and Vuong (1998) on measurement error. Li et al. (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.

First note that, as shown by Equation (8), the separability of firm costs is preserved in equilibrium. Thus, the purged winning bid, \bar{B}_w , can be written as $\bar{B}_w = Y + B_w$, where by B_w I indicate the purged of penalty and renegotiation winning bid conditional on the common cost Y being equal to zero. The pair $(\bar{B}_w; R)$ can therefore be thought of as a pair of convolutions $(Y + B_w; Y + Z)$. Since by Assumption (ii), the idiosyncratic cost X is independent of Y and Z , then B_w , which is a nonlinear transformation of X , is independent of Y 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 Y , B_w , and Z subject to a location normalization.²⁷ Then, Fourier transformations permit identifying the three marginal probability density functions of Y , 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 et al. (2000) shows.

5.3. *Estimation.* The Main data consist of m auctions for which $(n_i, \bar{b}_{wi}, r_i, a_i)_{i=1}^m$ are recorded: n_i is the number of bidders, \bar{b}_{wi} is the winning bid, r_i is the reserve price, and a_i is the cost overrun. The Panel data consist of m' auctions for which $(n_i, \bar{b}_{i1}, \dots, \bar{b}_{in_i}, e_{i1}, \dots, e_{in_i}, r_i, a_i)_{i=1}^{m'}$ are recorded: The \bar{b}_{ij} terms are all the bids, and the e_{ij} terms are dummy recording whether the bid was eliminated through the screening ($e = 1$) or not ($e = 0$). The estimation method, which closely follows that of Asker (2010) and Krasnokutskaya (2011), mimics the logic of the identification argument and consists of the following two-step procedure:

Purging step: First I estimate δ . For this I exploit the Panel data to estimate: $Pr(e_{it}) = \beta_{Screen}(r_t - b_{it})$. Combining the estimate, $\hat{\beta}_{Screen}$, with the monetary penalty of 2% of r gives $P(b|r) = 0.02r\hat{\beta}_{Screen}(r - b)$, so the estimator for delta is $\hat{\delta} = 0.02r(\hat{\beta}_{Screen})$.

Then, to estimate α , I estimate the empirical analogue of the linear regression in Equation (7) via the following ordinary least squares (OLS): $\bar{b}_{wi}(1 + \hat{\delta}) - \hat{\delta}r_i = \beta_{Overrun}a_i + \epsilon_i$, where ϵ equals the sum of the two production cost components and the strategic markup term. It follows from Assumption (ii) that $\hat{\alpha} = (1 - \hat{\beta}_{Overrun})$ is an unbiased estimator of α . With estimates for α and δ , I calculate the purged winning bids as defined above and proceed to the deconvolution step.

Deconvolution step: Since the implementation of the estimators entails a few passages, I break the discussion in two substeps.

Step D.1: Estimation of the probability density functions of Y and B_w . The first task is estimating the joint characteristic function of a winning bid and the reserve price. This is done nonparametrically 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 Y , Z , and B_w :

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

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

²⁷ The normalization that I use is $E(B_w) = 0$, but other normalizations would be possible.

Finally, the estimated probability density functions of Y , B_w , and Z are obtained through an inverse Fourier transformation.²⁸

Step D.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 method is used for this task.²⁹ 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 et al. (2000) can be applied to this sample of simulated bids. This entails first nonparametrically estimating the cdf and pdf of B_w^s .³⁰ 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 (8) implies that for every simulated winning bid we can use Equation (8), with Y 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 nonparametrically estimate the relative cdf and pdf and then to convert them into the corresponding parent distributions \hat{F}_X and \hat{f}_X .


5.4. Discussion. 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 unfeasible. 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.³¹ In my application, 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 nonbinding reserve price, the estimates reveal that it is nonbinding in more than 95% of the simulated FPAs. 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 could be used for that. For instance, in the United States, the auction data sets released by the DoT of many states report the engineers' project cost estimate. This quantity might work well with

²⁸ More specifically, these densities are estimated as $\hat{g}_u(q) = (2\pi)^{-1} \int_{-T_u}^{T_u} d_{T_u}(t) \exp(-itq) \phi_u(t) dt$, where $u \in \{Y, B_w, Z\}$, where d_{T_u} is a damping 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.

²⁹ 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 Italian Authority for Public Contracts 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.

³⁰ This is accomplished using the empirical analogue for the cumulative density function of B_w^s : $\hat{F}_{B_w^s}(b_w^s) = \frac{1}{M} \sum_{j=1}^M 1(B_{w_j}^s \leq b_w^s)$. The kernel estimator: $\hat{f}_{B_w^s}(b_w^s) = \frac{1}{M} \sum_{j=1}^M \frac{1}{h_g} \left[\frac{35}{32} \left(1 - \left(\frac{B_{w_j}^s - b_w^s}{h_g} \right)^2 \right)^3 1\left(\left| \frac{B_{w_j}^s - b_w^s}{h_g} \right| < 1 \right) \right]$ with bandwidth $h_g = (M)^{-\frac{1}{6}} (2.978)(1.06)(SD(B_w^s))$ is used to estimate the probability density of B_w^s .

³¹ Roberts (2013) considers a similar environment where only the winning bid and the reserve price are observable to the econometrician. His solution to the problem of separately estimating the common and private cost components when only the reserve price and the winning bid are observable was developed independently from the one presented in this article and is based on a control function method, which requires different assumptions than the ones used in this study.



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

FIGURE 2

ESTIMATES OF THE COST COMPONENTS DISTRIBUTIONS [COLOR FIGURE CAN BE VIEWED AT WILEYONLINELIBRARY.COM]

the proposed method because it is both linked to firm costs and is nonbinding for bidders. The idea of using an “instrument,” like the reserve price or the engineer’s estimate, to prove the identification of models with unobserved heterogeneity in the presence of incomplete data is explored at length in Hu et al. (2013). They show that identification is not merely the result of the separability and independence assumptions of the deconvolution approach used above, but that they can both be relaxed if a second instrument is available. It is also interesting to note that in my data their estimator could potentially be applied. However, since the procedures that can implement their method assume a discrete state space (i.e., the unobserved heterogeneity follows a discrete distribution), this would require discretizing the bids accordingly. But, since it is not obvious what the ideal discretization would be, I present results based on the method described above that does not require discretizations.³²

6. RESULTS

6.1. Baseline Estimates. I proceed along the steps described above and obtain estimates for δ and α , which are 0.0005 and 0.0988, respectively.³³ Both parameters have the expected sign and a magnitude that is plausible. In particular, the estimate of α implies that every euro renegotiated is worth only 90 cents, a finding similar to that in Bajari et al. (2014). Provided with these estimates, I then obtain estimates for the distributions of the commonly observed and idiosyncratic cost components. These are shown in Figure 2. Since the location of the two distributions is undetermined, I plot the distributions fixing the lowest bound of both of their supports at zero. This figure shows that the distribution of the common cost is characterized by

³² Finally, it is worth mentioning that for more complex environment where unobserved auction heterogeneity is present along with endogenous entry and asymmetric bidders, the literature has successfully employed parametric estimation techniques; see Krasnokutskaya and Seim (2011) and Athey et al. (2011).

³³ The probit regression for the probability of a bid excluded via screening on the size of the discount relative to the reserve price leads to an estimate of β_{Screen} of 0.03 (SD = 0.01). The OLS regression used to estimate $\beta_{Overrun}$ delivers a point estimate of 0.91 (SD = 0.01).

TABLE 4
EFFICIENCY COMPARISON

	No. Bids	ABA Groups	ABA Lottery	Winner Cost	Cost Ineff.	Share Ineff. Auct.
Baseline FPA	7	No	-	2.10	0%	0%
Counterfactual ABA	73	Yes	Unfair	2.44	16%	66%
Subscenarios for the ABA:						
S.1 Same bidders	7	No	Fair	3.01	43%	86%
S.2 Higher entry	73	No	Fair	3.06	45%	89%
S.3 Groups & fair lottery	73	Yes	Fair	2.49	18%	75%
S. 4 Groups & unfair lottery	73	Yes	Unfair	2.38	13%	54%

NOTES: The top panel of the table reports the values for the benchmark FPAs (first row) and of the counterfactual ABA (second row). The four rows in the bottom portion of the table describe the subscenarios described in the text. The column *No. Bids* reports the number of bidders. 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 column *Winner Cost* reports the average winner cost across the simulations. The column *Cost Ineff.* reports the (percentage) difference between the cost of the ABA winner and that of the corresponding FPA winner. The column *Share Ineff. Auct.* reports the share of auctions in which the ABA winner has a cost strictly above that of the corresponding FPA winner.

less variation than the one of the private cost. Under the additively separable structure of total cost and 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 accounts for 40% of the total cost variation. This confirms the relevance of the variation in the private cost that 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 a counterfactual ABA.³⁴

Starting from the estimated cost distributions, I first create a set of 1,000 simulated 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 a counterfactual ABA that incorporates the specificities of this format in terms of entry and collusion. To explain how I obtain the baseline estimates reported on the top rows of Table 4, it is useful to begin from the subscenarios reported in the lower part of the table. The idea of these subscenarios is to gradually move from a counterfactual ABA that is the right object if one seeks to consider the ideal benchmark of Lemma 2—a frictionless ABA—to a more realistic but less theoretically grounded ABA incorporating frictions.

In the first subscenario (S.1), I consider the case of a frictionless ABA. 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 4, on average the cost of the winner in the ABA is 43% 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

³⁴ Note also that for the given cost estimates, the zero expected profit condition of the entry model implies an entry cost of about € 5,000, which is reasonable given both the accounting costs of participation (paying for the project documents, the entry fee, and the salary of the technicians involved in preparing the price bid and the technical documentation) and the opportunity cost of forgone business opportunities (associated with having resources invested in a particular auction process).

that the ABA allocates to the lowest bidder, the share of inefficiently allocated auctions is 86%. 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 subscenario (S.2) 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 other 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 4 shows, the first inefficiency measure worsens, going to 45%. The second performance measure also worsens, with 89% 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 or higher than that of each FPA bidder.

The third subscenario (S.3) starts capturing how bidders' collusion in ABAs interacts with the efficiency of the allocation. This step is key for a realistic characterization of the ABA but entails the obvious difficulty of incorporating a hard to observe behavior. Here, my approach is to 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 (2016). In the 276 ABAs that were presented in the court case, out of the 73 bidders participating on average, 43 were noncooperating firms, whereas the remaining 30 belonged to groups of cooperators. The six groups into which these 30 firms are divided have size 11, 6, 6, 3, 2, and 2. In S.3, which I refer to as the "fair lottery" case, I proceed as follows: (i) draw costs for each bidder, (ii) assign at random the participation of firms into groups, (iii) randomly draw the winner across the 73 bidders, and (iv) use as the cost of the winner the lowest cost among all those of the bidders in the same group of the winner. Essentially, this amounts to having an ABA where groups allocate the contract efficiently between themselves but do not go as far as coordinating their bids to manipulate the bid distribution. This later assumption is relaxed in the later subscenario, S.4. There, I consider the case of an "unfair lottery" in the sense that the winning probability is not identical across bidders. This is fundamental, as the court cases document how winning probabilities are indeed distorted by the activities of the groups competing to manipulate the bid distribution. Here, I resort again to what was reported in Conley and Decarolis (2016) about the probability of winning of each of the six groups described above. That is, each group is assigned a winning probability equal to the relative winning frequency of this group in the court case. That is, a winning probability of 36%, 13%, 10%, 2%, 4%, and 1%, where the order goes from the largest to the smallest cartel. Moreover, within each group all firms are assumed to have the same probability of winning.

The last two rows of Table 4 report the results of S.3 and S.4. As expected, group bidding alleviates the ABA inefficiency. In the fair lottery case, the amount of extra cost of the ABA relative to the FPA reduces to 18%. Nevertheless, the share of auctions that select a winner with a higher cost than in the FPA remains high, 75%. The reduction of the inefficiency is even stronger in the unfair lottery case: The extra cost of the ABA declines to 13%, and the share of inefficiently allocated auctions declines to 54%. This happens because in this latter scenario the largest groups are highly likely to win, and, conditional on winning, they give the contract to their lowest cost member.

It is now possible to return to the counterfactual ABA in the top portion of Table 4. The values reported here are obtained as a weighted average of those in S.3 and S.4 where the weights are 57% on S.3 and 43% on S.4. These weights are still based on Conley and Decarolis (2016) but from a different part of their study relative to the analysis of the court case described above. In fact, in a second part of their paper they used statistical tests—validated on the sample of court case auctions—to detect collusion in a broader data set that very closely overlaps with the Main data set in this study. In this data set, collusion is likely but is not known from a court investigation. Their analysis reveals that groups of closely connected firms are widespread but that active manipulations of the bid distribution are most likely present in only 43% of the

TABLE 5
ROBUSTNESS CHECKS

	Winner Cost FPA	Winner Cost ABA	Cost Ineff.	Share Ineff. Auct.
Baseline estimates	2.10%	2.44%	16%	66%
Auction homogenization	2.03%	2.35%	16%	61%
Turin administrations	2.53%	2.50%	16%	67%
Generalized deconvolution	1.75%	2.12%	21%	65%
Small noise ($\pm 50k$)	1.81%	2.25%	24%	66%
Big noise ($\pm 100k$)	1.32%	1.79%	35%	67%

NOTES: The first row reports the baseline estimates from the first two rows of Table 4. The next rows report the same set of estimates, performed according to each of the five robustness checks described in the main text. Additional details regarding the estimates using generalized deconvolution and auction homogenization are reported in the Supporting Information Appendix.

auctions. Therefore, I maintain the proportion they estimated for the share of auctions with distorted probabilities to weight S.3 and S.4 to construct the final ABA counterfactual. The corresponding estimates reported in the second row of the table show that the overall amount of inefficiency is 16% in terms of the excess cost of the ABA winner relative to the FPA winner or 66% of the ABA have a winner with higher cost than in the corresponding FPA. Compared to S.1, it is clear how incorporating entry and collusion influenced the results. Accounting for entry (marginally) increased the estimated efficiency. However, the inclusion of collusion drastically reduced it, although not enough to eliminate it.³⁵

Finally, the efficiency 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. As argued earlier, the bidders that select into these highly competitive FPAs are likely more efficient than those entering ABAs. Thus, the idiosyncratic cost distribution of FPA bidders gives a very conservative estimate of the potential cost dispersion among the less homogeneous and more inefficient ABA bidders. In addition to this selection argument, the fact that the allocation might not be perfectly efficient within groups contributes to suggest that my estimates are best interpreted as a lower bound on the inefficiency of ABAs.

6.2. Robustness. To assess the reliability of the above results, I conducted a series of robustness checks, which are summarized in Table 5. For each robustness check, I report the baseline estimates for the FPA and the ABA. The first two robustness checks deal with observed heterogeneity. In the fourth row, I control for observable auction heterogeneity by “homogenizing” the auctions. As standard, I perform this 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 procurer) and then using the regression residuals to perform the analysis described earlier. The estimates are in the same ballpark as the baseline estimates. Even more importantly, the results of the OLS regression (reported in the Supporting Information Appendix) indicate that all coefficients are not statistically significant³⁶ and that the explained variance explained by the model is only 5%. This is relevant, as it is evidence in support of the fact that the construction of the sample was already made in ways that ensured a high level of homogeneity in the auctions.

The second row of the table involves changing the sample by performing the analysis exclusively on the auctions held by the Turin county and municipalities. As mentioned earlier (and described at greater length in Decarolis, 2014), these two administrations were the first to

³⁵ It should be stressed here that key for this result is the assumption that collusive groups in the ABA allocate efficiently the contract within the group. This is indeed in line with the anecdotal evidence in Conley and Decarolis (2016) and is supported by the evidence on the large use of subcontracting in Branzoli and Decarolis (2015). Assuming efficiency is also useful to the extent that I seek to characterize a lower bound for the inefficiency of the ABAs in this environment.

³⁶ With the exception of some of the year fixed effects.

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 heterogeneous bidding behavior in auctions held by different administrations. But even in these case the estimates are fairly close to the baseline ones.

The third row of the table deals with an alternative approach to estimate unobserved heterogeneity. It shows the estimates obtained by replacing the Li and Vuong (1998) deconvolution estimator used in step D.1 of the estimation procedure with the Bonhomme and Robin (2010) “generalized deconvolution” estimator. The generalized deconvolution estimates are based on a slightly different set of assumptions, namely, the assumption of finite supports of the distributions is not required for this estimator but required by Li and Vuong (1998). Overall, the estimates are not far from those in the baseline estimates: The cost inefficiency amounts to 21% and the share of inefficient auctions is 65%. Since in my data the assumption of finite supports for the cost distributions seems not too restrictive and since the slower rate of convergence of the generalized deconvolution estimator is not ideal with small sample sizes, I prefer to use the Li and Vuong estimator for my estimates.³⁷

The last two robustness checks consist of adding a mean preserving spread to the estimated idiosyncratic cost distribution to assess the sensitivity of the estimates. I add a draw from a uniform distribution on $[-u; +u]$, where u equals € 50,000 in the first robustness check and € 100,000 in the second. This is a sanity check for the analysis, as inefficiency should increase with more noise in the private cost. In both cases, relative to the original estimates, the winner’s cost declines in the FPA in a way that is not fully matched by the ABA, thus bolstering the estimated inefficiency. The estimated cost inefficiency increases from 16% to 24% (“small noise”) and then to 35% (“big noise”). The share of inefficient auctions, instead, is less sensitive to these modifications.

Overall, the set of robustness checks presented here is reassuring on the magnitude of the estimated inefficiency of ABAs. Across all counterfactuals but those involving the mean preserving spreads, the range of the cost inefficiency is between 16% and 24%. Furthermore, the proportion of inefficiently allocated auctions is even more tightly estimated, with estimates ranging from 61% to 67%.³⁸

7. CONCLUSIONS

This article 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 FPA. 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 data set of Italian procurement auctions for public works, I estimate the inefficiency associated with the ABA. Counterfactual estimates of the efficiency loss under the ABA, accounting for the presence of colluding bidders in this format, show that this mechanism fails to select to lowest bidder in two thirds of the auctions and that the average production cost is one sixth higher than in the FPAs. 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 strongly points out the limits of the alternative solution represented by the FPA with bid

³⁷ See a more detailed discussion in the Supporting Information Appendix.

³⁸ Several other potential robustness checks of interest could be conducted, for instance, in terms of the potential number of bidders, which the earlier analysis holds fixed at 17. For small changes around this value, additional estimates not reported here indicate that the results remain qualitatively similar.

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 a 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.

APPENDIX: PROOFS

A.1. *Proof of Lemma 1.* To prove Lemma 1, I first introduce the following lemma.

LEMMA A.1. *Assuming $(\frac{\bar{x}_L - \underline{x}_L}{1-\theta}) < \varepsilon < (\frac{1}{\theta})y$, 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 $p_L < p_L^* = (1 - \theta)\varepsilon + \underline{x}_L - \bar{x}_L$ 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, that is, $p_H \leq (y + x_i + \varepsilon) - b_H$; 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 $b \geq y + \underline{x}_H + \theta\varepsilon$. With regard to the second part of the lemma, notice that in an FPA no bid is higher than $y + \bar{x}_L + \theta\varepsilon$ (following a simple Bertrand argument). Therefore, if $p_L < p_L^*$, a type L bidder will always default when the project is costly because $p_L < (1 - \theta)\varepsilon + \underline{x}_L - \bar{x}_L \leq (y + \underline{x}_L + \varepsilon) - b_L$.

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, I 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), from now on RZ. 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 $v_j = \bar{x}_j + s_j$, $r_j = \underline{x}_j + s_j$, 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 $[\underline{x}_j, \bar{x}_j]$ for each \tilde{b}_j^i and continuous in \tilde{b}_j^i for each \tilde{x}_j^i ; (ii) define $b^* \equiv \max(a_H + \bar{x}_H, a_L + \bar{x}_L)$; then $\tilde{u}_j^i(\tilde{b}_j^i, \tilde{x}_j^i) < 0$ for all $\tilde{b}_j^i > b^*$ and for any $\tilde{x}_j^i \in [\underline{x}_j, \bar{x}_j]$; (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{b}_j^i ; and (iv) $\tilde{u}_j^i(\tilde{b}_j^i, \tilde{x}_j^i) - \tilde{u}_j^i(\underline{b}_j^i, \tilde{x}_j^i)$ is constant in \tilde{x}_j^i .*

RZ-ASSUMPTION 2 (SIGNALS). *Assume that the private value \tilde{x} is a monotonic function $x : [0, 1]^N \rightarrow [\underline{x}, \bar{x}]^N$; 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 Assumptions 1 and 2 are satisfied, existence follows.

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

A.2. Proof of Theorem 1. The fact that the strategy profile in which all bidders offer the maximum bid 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 nonnegative 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 A.1 (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 A.2 (NONDECREASING FUNCTION). The proof is by contradiction. Assume that the BNE bidding function, b , has an interval over which it 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

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

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

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

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

PROPERTY A.3 (NONSTRICTLY 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 $b(\underline{x}) = \underline{b} < b(x) \forall x \neq \underline{x}$, 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 $x \neq \underline{x}$ and $\underline{b} + \delta$ for $x = \underline{x}$. His expected revenues are unchanged for any $x \neq \underline{x}$, and they are strictly higher for $x = \underline{x}$ since the probability of winning goes from being zero to being positive. By Property A.2 and continuity, we must have that the bidding function is flat at the bottom.

PROPERTY A.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 A.3) 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, $\Pr(\text{win}|\underline{b})[\underline{b} - \underline{v}] \geq \Pr(\text{win}|b)[b - \underline{v}]$ 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: Let p be the probability that all the competing $N - 1$ bidders draw a value no higher than \underline{v} and let

q_1, q_2, \dots, q_{N-2} be the probabilities that, respectively, exactly one, two, \dots , $N - 1$ rival bidders draw a value strictly higher than \underline{v} . Furthermore, let α_M be

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

where $M=1, \dots, N - 2$. I can now rewrite $\Pr(\text{win}|\bar{b})$ as $\Pr(\text{win}|\bar{b}) = p(\frac{1}{N}) + [q_1(\frac{1}{N-1}) + q_2\alpha_2(\frac{1}{N-2}) + \dots + q_N\alpha_N(\frac{1}{N-N})]$, 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:

$$\beta_M \equiv \Pr \left[\left| b' - \frac{1}{N} \sum_{r=1}^N b_r^* \right| < \left| b(x_j) - \frac{1}{N} \sum_{r=1}^N b_r^* \right| \text{ for any } x_j > \underline{v} \text{ and } j = 1, 2, \dots, M | q_M = 1 \right],$$

where $M = 1, 2, \dots, N - 2$. Therefore, I can now rewrite $\Pr(\text{win}|b')$ as $\Pr(\text{win}|b') = [q_1 + q_2\beta_2 + \dots + q_{N-2}\beta_{N-2}]$. Now, given the way b' was chosen, it must be that $[q_1 + q_2\beta_2 + \dots + q_{N-2}\beta_{N-2}][b' - \underline{v}] \geq [q_1 + q_2\alpha_2 + \dots + q_N\alpha_N][\underline{b} - \underline{v}]$. 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 $\{p(\frac{1}{N}) + [q_1(\frac{1}{N-1}) + q_2\alpha_2(\frac{1}{N-2}) + \dots + q_N\alpha_N(\frac{1}{N-N})]\}[\underline{b} - \underline{v}] \geq [q_1 + q_2\alpha_2 + \dots + q_N\alpha_N][\underline{b} - \underline{v}]$. 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

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

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 1 for \underline{v} equal to \bar{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 numerator 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 $(v^*, \bar{x}]$, where v^* is defined to be the value of \underline{v} such that the inequality of (*) would be an equality. Moreover, since $v^* < \bar{x}$ 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: $|\underline{v}_{\delta, F} - \bar{x}| < \varepsilon$.*

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 $g(\underline{v}) > g(v^*) \rightarrow \underline{v} > v^*$ 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 \bar{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 \bar{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 $(v^*, \bar{x}]$ very rapidly by bringing v^* closer to \bar{x} . Therefore, even for small N , v^* will be close to \bar{x} . This is the reason why, even for small N , $(*)$ gives a bound that is useful.

A.3. 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% of prices, then the lowest price equal to or higher than $A1$ wins. Third, a random draw is used to ensure that exactly 10% 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% 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 than 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 nonbinding. 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 weakly 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 .

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's website.

REFERENCES

- ASKER, J., "A Study of the Internal Organization of a Bidding Cartel," *American Economic Review* 100(3) (2010), 724–62.
- ATHEY, S, D. COEY, AND J. LEVIN, "Set Asides and Subsidies in Auctions," *American Economic Journal: Microeconomics* 5(1) (2013), 1–27.
- , J. LEVIN, AND E. SEIRA, "Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions," *Quarterly Journal of Economics* 126(1) (2011), 207–57.
- BAJARI, P., S. HOUGHTON, AND S. TADELIS, "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," *American Economic Review* 104(4) (2014), 1288–319.
- BOARD, S., "Bidding into the Red: A Model of Post-Auction Bankruptcy," *Journal of Finance* 62(6) (2007), 2695–723.
- BONHOMME, S., AND J. M. ROBIN, "Generalized Deconvolution with an Application to Earnings Dynamics," *Review of Economic Studies* 77 (2010), 491–533.
- BRANZOLI, N., AND F. DECAROLIS, "Entry and Subcontracting in Public Procurement Auctions," *Management Science* 61(12) (2015), 2945–62.

- BURGUET, R., J.-J. GANUZA, AND E. HAUK, "Limited Liability and Mechanism Design in Procurement," *Games and Economic Behavior* 76(1) (2012), 15–25.
- CALVERAS, A., J.-J. GANUZA, AND E. HAUK, "Wild Bids. Gambling for Resurrection in Procurement Contracts," *Journal of Regulatory Economics* 26(1) (2004), 41–68.
- CHANG, W.-S., B. CHEN, AND T. C. SALOMON, "An Investigation of the Average Bid Mechanism for Procurement Auctions," *Management Science* 61(6) (2015), 1237–54.
- CHE, Y.-K., AND J. KIM, "Bidding with Securities: Comment," *American Economic Review* 100(4) (2010), 1929–35.
- CHILLEMI, O., AND C. MEZZETTI, "Optimal Procurement Mechanisms: Bidding on Price and Damages for Breach," *Economic Theory* 55(2) (2013), 335–55.
- CONLEY, T. G., AND F. DECAROLIS, "Detecting Bidders' Groups in Collusive Auctions," *American Economic Journal: Microeconomics* 8(2) (2016), 1–38.
- DECAROLIS, F., "Awarding Price, Contract Performance and Bids Screening: Evidence from Procurement Auctions," *American Economic Journal: Applied Economics* 6(1) (2014), 108–32.
- , AND C. GIORGIANTONIO, "Local Public Procurement Regulations: The Case of Italy," *International Review of Law and Economics* 43(C) (2015), 209–26.
- DE CASTRO, L. I., AND M.-A. DE FRUTOS, "How to Translate Results from Auctions to Procurements," *Economics Letters* 106(2) (2010), 115–18.
- DIGGLE, P., AND P. HALL, "A Fourier Approach to Nonparametric Deconvolution of a Density Estimate," *Journal of the Royal Statistical Society, Series B (Methodological)* 55(2) (1993), 523–31.
- ENGEL, A. R., J.-J. GANUZA, E. HAUK, AND A. WAMBACH, "Managing Risky Bids," in N. Dimitri, G. Piga, and G. Spagnolo, eds., *Handbook of Procurement* (Cambridge: Cambridge University Press, 2006), 322–46.
- EUN, D., "Procurement with a No Loss Constraint: Evidence and a New Mechanism from Korea," Mimeo, 2015.
- EUROPEAN COMMISSION, "Prevention, Detection and Elimination of Abnormally Low Tenders in the European Construction Industry," Enterprise and Industry Report, 2002.
- GALAVOTTI, S., L. MORETTI, AND P. VALBONESI "Sophisticated Bidders in Beauty-Contest Auctions," Mimeo, 2015.
- GUERRE, E., I. PERRIGNE, AND Q. VUONG, "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica* 68(3) (2000), 525–74.
- HU, Y., D. McADAMS, AND M. SHUM, "Identification of First-Price Auctions with Non-separable Unobserved Heterogeneity," *Journal of Econometrics* 174(2) (2013), 186–93.
- IOANNOU, P. G., AND S.-S. LEU, "Average-Bid Method. Competitive Bidding Strategy," *Journal of Construction Engineering and Management* 119(1) (1993), 131–47.
- KOTLARSKI, I., "On Some Characterizations of Probability Distributions in Hilbert Spaces," *Annali di Matematica Pura e Applicata* 74 (1966), 129–34.
- KRASNOKUTSKAYA, E., "Identification and Estimation in Highway Procurement Auctions under Unobserved Heterogeneity," *Review of Economic Studies* 78(1) (2011), 293–327.
- , AND K. SEIM, "Bid Preference Programs and Participation in High-Way Procurement Auctions," *American Economic Review* 101(6) (2011), 2653–86.
- LEVIN, D., AND J. L. SMITH, "Equilibrium in Auctions with Entry," *American Economic Review* 84(3) (1994), 585–99.
- LEWIS, G., AND P. BAJARI, "Procurement Contracting with Time Incentives: Theory and Evidence," *The Quarterly Journal of Economics* 126(3) (2011), 1173–211.
- LI, T., "Econometrics of First-Price Auctions with Entry and Binding Reservation Prices," *Journal of Econometrics* 126(1) (2005), 173–200.
- , I. PERRIGNE, AND Q. VUONG, "Conditionally Independent Private Information in OCS Wildcat Auctions," *Journal of Econometrics* 98(1) (2000), 129–61.
- , AND Q. VUONG, "Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators," *Journal of Multivariate Analysis* 65 (1998), 139–65.
- MARION, J., "Are Bid Preferences Benign? The Effect of Small Business Subsidies in High-Way Procurement Auctions," *Journal of Public Economics* 91(7–8) (2007), 1591–624.
- MARMER, V., A. SHNEYEROV, AND P. XU, "What Model for Entry in First-Price Auctions? A Nonparametric Approach," *Journal of Econometrics* 176(1) (2013), 46–58.
- RENY, P. J., AND S. ZAMIR, "On the Existence of Pure Strategy Monotone Equilibria in Asymmetric First-Price Auctions," *Econometrica* 72(4) (2004), 1105–25.
- RHODES-KROPF, M., AND S. VISWANATHAN, "Financing Auction Bids," *RAND Journal of Economics* 36(4) (2005), 789–815.
- ROBERTS, J. W., "Unobserved Heterogeneity and Reserve Prices in Auctions," *RAND Journal of Economics* 44(4) (2013), 712–32.

- SPAGNOLO, G., G. L. ALBANO, AND M. BIANCHI, "Bid Average Methods in Procurement," *Rivista di Politica Economica* 1(2) (2006), 41–62.
- SPULBER, D. F., "Auctions and Contract Enforcement," *Journal of Law, Economics and Organization* 6(2) (1990), 325–44.
- WAEHRER, K., "A Model of Auction Contracts with Liquidated Damages," *Journal of Economic Theory* 67(2) (1995), 531–55.
- ZHENG, C. Z., "High Bids and Broke Winners," *Journal of Economic Theory* 100 (2001), 129–71.