

Regulating bidder participation in auctions

Vivek Bhattacharya*

James W. Roberts**

and

Andrew Sweeting***

Regulating bidder participation in auctions can potentially increase efficiency compared to standard auction formats with free entry. We show that the relative performance of two such mechanisms, a standard first-price auction with free entry and an entry rights auction, depends nonmonotonically on the precision of information that bidders have about their costs prior to deciding whether to participate in a mechanism. As an empirical application, we estimate parameters from first-price auctions with free entry for bridge-building contracts in Oklahoma and Texas and predict that an entry rights auction increases efficiency and reduces procurement costs significantly.

1. Introduction

■ Given the importance of public procurement for many economies around the world (e.g., 12% of GDP (Gross Domestic Product) is spent on public procurement in OECD (Organization for Economic Cooperation and Development) countries (OECD, 2011)), there is natural interest in improving the efficiency of the procurement process and reducing procurement costs. Under the assumptions, which we maintain, of symmetric and independent private costs, quite standard auction formats, such as first-price (low-bid) auctions with suitably chosen reserve prices will be optimal when the number of bidders is fixed. Assuming a fixed set of bidders will not be appropriate, however, when suppliers have to spend significant resources in “due diligence” to understand how much it will cost them to complete the project, and in most of the procurement settings that have been studied, these types of entry costs appear to be significant (e.g., Li and

*MIT; bhattachv@mit.edu.

**Duke University and NBER; j.roberts@duke.edu.

***University of Maryland and NBER; sweeting@econ.umd.edu.

We are very grateful to Dakshina De Silva, Timothy Dunne, Anuraddha Kankanamge, and Georgia Kosmopoulou for sharing their data with us. Ali Hortaçsu, Tim Hubbard, Jon Levin, Rob Porter, Paulo Somaini, Steve Tadelis, John Turner, John Asker (the co-editor), seminar participants, and three referees provided feedback that improved the article. Bhattacharya would like to thank the Davies Fellowship from the Duke University Department of Economics for financial support. Roberts and Sweeting acknowledge National Science Foundation support (grant no. SES-1260876) for this research. Any errors are our own.

Zheng, 2009; Krasnokutskaya and Seim, 2011). In this article, we examine the performance of a standard first-price auction with costly but unregulated, or “free,” entry relative to an alternative mechanism where the procurer regulates entry by using an initial auction to allocate a fixed number of rights to participate in a second auction for the contract (an “entry rights auction”).

When entry is endogenous, a standard first-price auction is usually thought of as a two-stage game. In the first stage, a set of potential suppliers (e.g., a set of local contractors who purchase or are issued with specifications for the contract) simultaneously and noncooperatively decide whether to pay the cost of entering the auction based on any private information (signals) that they have about their costs of completing the project. In the second stage, the entrants find out their true costs of completing the project and submit bids. The contract is allocated to the firm with the lowest bid, as long as this is less than any reserve price, at a price equal to its bid. We will call this standard model a “first-price auction with free entry” (FPAFE). In the equilibrium of an FPAFE, the marginal entrant in the first stage will be indifferent between entering the auction, which involves paying the entry cost but allows the supplier to potentially win the contract, and staying out.

A natural question is whether these entry decisions are efficient, in the sense of being consistent with the minimization of the expected sum of the winning bidder’s cost of completing the project and total entry costs (social costs). As shown theoretically by Levin and Smith (1994) and Gentry and Li (2012), under different assumptions about the private information that potential suppliers have about their costs prior to entry, entry decisions will be efficient in the sense that a planner who had to choose whether a potential supplier should enter based only on that firm’s private information would choose the same entry rule that the potential supplier would choose in the symmetric equilibrium.¹ However, inefficiencies in entry can still arise from the fact that the entry decisions of different potential bidders are not coordinated, so that there can be outcomes where too many or too few firms enter. For example, when potential suppliers have no private information about their costs when taking entry decisions, as assumed by Levin and Smith (1994) (LS), the symmetric equilibrium involves potential suppliers mixing over entry. As a result, the number of entrants will be stochastic, and Milgrom (2004) shows that efficiency would be increased if the procurer instead randomly selected a fixed number of potential suppliers to enter the auction.

The assumption that potential suppliers have no private information about their costs before entering seems implausible for many procurement settings, as firms are likely to know at least something about their free capacities and their costs of securing some of the necessary inputs. When their private signals are informative about their costs (Roberts and Sweeting, 2013b; Gentry and Li, 2014), free entry will tend to lead to the most efficient firms entering, and an arrangement where the auctioneer randomly selected entrants without using this private information would have a natural disadvantage. However, there could still be scope for the auctioneer to regulate the number of entrants by using an entry rights auction mechanism (ERA) where, based on their private information, the potential suppliers first participate in an auction for a fixed number of rights to perform due diligence and enter the auction for the contract. As long as first-stage bids are monotonic in the bidders’ private signals, the ERA will also tend to select the most efficient firms, but the ERA will have the advantage that the number of entrants will be controlled.

We show, using an example, that whether an FPAFE or an ERA is more efficient can depend in a nonmonotonic way on the precision of the private information that suppliers have about their costs. When signals are imprecise, the ERA is more efficient, as one would expect, given the

¹ The typical Mankiw and Whinston (1986) result that free entry results in excess entry in a homogeneous goods market does not hold because an entrant only takes business from other firms when it is socially optimal for it to do so. The Levin and Smith and Gentry and Li results are conditional on the auctioneer setting a reserve price that is equal to the cost of completing the project outside of the auction if there is no winner. This is the auctioneer-optimal reserve in the Levin and Smith model, but not the Gentry and Li model, although it is the natural “nonstrategic” reserve price in that setting as well. As pointed out by a referee, one interpretation of these results is that the procurer cannot improve efficiency by choosing a reserve strategically.

informational assumptions made by LS and Milgrom in deriving the result mentioned above. When signals are very informative, in which case the model tends toward the Samuelson (1985) (S) model where potential entrants know their costs exactly, the ERA is also more efficient because it is optimal to have only the supplier with the best (lowest cost) signal paying the entry cost; this is the outcome that the ERA ensures will happen, whereas no supplier or more than one supplier may enter in an FPAFE. However, when potential bidders' signals are moderately informative about their costs, an FPAFE can be more efficient. The reason is that with free entry the number of entering suppliers will depend on the private information available to the suppliers but not the auctioneer, and this can be desirable. For example, although on average it might be desirable to have two entrants, three entrants might be preferred if they all have signals that indicate that their costs are likely to be low.

Given this ambiguity in which mechanism is more efficient, we estimate our model using data on a sample of contracts for highway projects involving work on bridges let by the Departments of Transportation (DoTs) in Oklahoma and Texas using FPAFEs in order to understand which mechanism might perform best in a real-world context. To do so, we develop methodologies for both solving and estimating a parametric model of FPAFEs where potential bidders have imperfect information about their costs when deciding whether to enter. We solve the model using the Mathematical Programming with Equilibrium Constraints (MPEC) approach proposed by Su and Judd (2012).² We estimate our model using a simulated method of moments estimator where the moments are computed using importance sampling, following Akerberg (2009). This approach enables us to allow for observed and unobserved heterogeneity in the parameters across auctions, unlike the one previous attempt to estimate this type of model of which we are aware (Marmer, Shneyerov, and Xu, 2011).³

We estimate that entry is moderately selective for these contracts. However, given the other parameters of our model, our model predicts that overall efficiency would have been increased by using the ERA format. For example, for the representative (average) auction in our data, we predict that the sum of the expected cost of the winning bidder and total entry costs would be 2.4% lower in the ERA than the FPAFE and that the total cost to the DoT of letting the contract (procurement cost) would fall by the same proportion. This fall in DoT procurement costs is much greater than what we predict would be achieved by using a more standard design adjustment to an FPAFE, such as adding an optimal reserve price, which for our representative case would only reduce costs by 0.03%, or by adding additional potential entrants.

We are not aware of any previous attempts in the literature to compare ERAs and auctions with free entry using real-world parameters. A related theoretical literature has considered so-called "indicative bidding" schemes where potential bidders submit nonbinding indications of what they are willing to bid, which the seller may use to select a subset of firms to participate in an auction.⁴ In contrast, in an ERA, first-round bids are binding and can result in payments by firms that do not win the contract. Indicative schemes are often used by investment banks and as part of the complicated processes by which defense equipment is purchased (Welch and Fremond, 1998; Foley, 2003; Quint and Hendricks, 2013). Ye (2007) argues that, generically, indicative schemes will not induce strictly monotonic first-stage bids in equilibrium, and so will

² Hubbard and Paarsch (2009) use this approach to solve first-price auction models under the S model assumption that potential bidders know their values when deciding whether to enter. We extend their solution method to allow for noisy signals about values (in our case, costs), creating imperfect selection of low-cost bidders into the auction.

³ Some earlier work, such as Krasnokutskaya and Seim (2011), estimates models of endogenous entry into first-price auctions under the informational assumptions of LS. This implies that there is no selection in the entry process, which simplifies estimation because the distribution of values for bidders will be the same as the distribution of values in the population of potential suppliers.

⁴ There are other designs that also have the flavor of using an auction to select a small number of firms into another round. For example, in the Anglo-Dutch auction described by Klemperer (2002), bidders first participate in an ascending auction that selects a smaller set of bidders who compete in a first-price sealed bid auction. This design has been discussed primarily in settings where bidders are asymmetric (e.g., where there are well-known strong incumbents) and where valuations are likely to have some common component. In our model, suppliers are symmetric and we assume IPV.

not tend to guarantee the selection of the bidders that are most likely to have the highest values or the lowest costs in the second stage. However, he shows that certain types of ERA schemes, like the one that we consider, will induce strictly monotonic first-stage bidding and so guarantee an efficient selection of entrants. Recently Quint and Hendricks (2013) have shown that, when indicative bids are limited to a discrete set and either entry costs or the number of bidders is large, the unique indicative bidding equilibrium can involve weakly monotonic bid functions, and that, in computed examples, indicative schemes can outperform standard auctions with free entry. By guaranteeing fully efficient selection of entrants, ERAs should, in theory, be more efficient than indicative schemes when the number of selected entrants is the same. At the end of the article, we discuss several possible reasons why ERAs may be less common than indicative schemes or FPAFES in practice.

The article is also related to our earlier work on selective entry auction models (Roberts and Sweeting, 2013a, 2013b).⁵ In those articles, we consider selective entry into second-price auctions, where the simpler form of equilibrium bidding eases computation and estimation relative to the first-price auctions considered in this article. FPAFES are the auction format that is most widely used in practice.⁶ In Roberts and Sweeting (2013b), we compare the performance of a free-entry auction with a sequential bidding process, which is more efficient, and can improve the seller's revenues for a wide range of parameters, even though it allows early bidders to deter entry. Here, we consider a different type of design change, which retains the simultaneous bidding feature of most auction formats. This may be attractive, as sequential procedures have the potential weakness that the auctioneer may favor certain bidders when determining the order.

Four comments are in order about both our model and the nature of the results. First, when considering the ERA, we model the first-stage auction for entry rights as an all-pay auction. Following Ye (2007), the advantage from a modelling perspective of using the all-pay format is that equilibrium first-stage bids are guaranteed to be strictly monotonic in the signal that potential bidders have about their costs when entry costs are subsidized, which we also assume. This guarantees that the suppliers that are most likely to have the lowest costs are selected. When the more common discriminatory or uniform price formats are used to sell entry rights, this property may not hold because the expected benefit of being the marginal entrant allowed into the final auction, which determines the bid in these formats, is not guaranteed to be monotonic in the signal. In fact, it is straightforward to find examples where the relationship is not monotonic for the distributions that we consider. However, as Ye shows, when the relationship is monotonic, the expected costs and efficiency under all three first-stage schemes (all-pay, discriminatory, and uniform) will be identical.

Second, we are comparing our results to a specific form of ERA, where the seller announces an exact value for the number of bidders who will be able to enter the second-stage auction in advance, rather than considering the optimal mechanism or even the optimal entry rights auction, where the number of allowed entrants into the second stage might well depend on the bids that are announced in the first stage. We do this partly because it is not straightforward to characterize the procurer-optimal form of entry rights auction in our model, but also because even if it were known, the optimal mechanism would likely have such a complicated form that a public agency would likely not want to use it in practice.⁷ In any case, because our bottom-line conclusion is

⁵ Other work on selective entry in auctions includes Gentry and Li (2014), who focused on identification, and Marmer, Shneyerov, and Xu (2013), who consider nonparametrically testing different models using first-price auction data in the absence of unobserved auction heterogeneity.

⁶ As we assume potential suppliers are symmetric, it is possible to compute expected efficiency and procurement costs for first-price auctions by exploiting the fact that these outcomes should be the same in the second-price format. However, to use the information contained in the distributions of submitted bids and winning bids in estimation, it is necessary to be able to solve first-price auctions with selective entry.

⁷ Lu and Ye (2013) characterize the optimal entry rights auction format in a setting where there is heterogeneity in potential bidders' entry costs, but they have no private information on their values prior to entering. Cremer, Spiegel, and Zheng (2009) characterize the optimal mechanism in this case, which involves a complicated sequential search process.

that there may be significant cost and efficiency advantages to using ERAs, this conclusion would only be strengthened if we considered the optimal form of the ERA instead. In discussing our results, we do, however, compare the efficiency of our ERA mechanism with the efficiency of two hypothetical mechanisms where we allow a planner to have access to the private information of the bidders and to use these signals to choose how many firms enter. Although these hypothetical alternatives are naturally more efficient, we show that the ERA actually performs almost as well as they do for the parameters that we estimate.

Third, our results are numerical. When entry is partially selective, it is very difficult to solve models analytically outside of some special cases.⁸ Moreover, as we show, efficiency and cost comparisons will often depend on the distributions and parameters that are assumed, so it is important to try to be as realistic about these elements as possible. For this reason, we focus our results on parameters that are estimated from real-world data and emphasize the magnitude as well as qualitative direction of our results.

Finally, although we allow for partially selective entry and a model of auction heterogeneity, we maintain the independent private costs/values assumption that has been typically made in both empirical studies of public procurement (e.g., Krasnokutskaya and Seim, 2011) and theoretical work on models of two-stage entry rights or indicative bidding procedures (e.g., Quint and Hendricks, 2013). Understanding the effects of trying to regulate entry where costs have both common value and private value components is left to future work. Similarly, we assume that firms would act competitively in both FPAFEs and ERAs, whereas in some settings it might be important to consider which mechanism would be more susceptible to collusion.

The article proceeds as follows. Section 2 presents the model of imperfectly selective entry for FPAFEs and ERAs, and explains how we solve these models. Section 3 describes the data. Section 4 explains our estimation method. Section 5 discusses the estimates of the model's parameters and compares the performance of the FPAFE and the ERA. Section 6 concludes. The Appendixes contain some additional details on our numerical routines and present some Monte Carlo evidence on the performance of our estimator.

2. A model of endogenous bidder participation in auctions

■ We introduce the general structure of costs and information in our model of procurement and then describe the FPAFE and the ERA. We also present a numerical example illustrating how an efficiency comparison of the two mechanisms will depend on the degree of selection in the entry process.

A procurement agency, which we will call the “procurer” in what follows, wishes to select one of N risk-neutral suppliers to complete a project. In our empirical setting, a project might involve repairing or constructing a bridge, and we will allow for cross-auction variation in both N and the other parameters. The agency may require that it pay no more than a reserve price r to complete the project, and if no firm submits a bid below this amount, the agency incurs a cost c_0 of procuring the work outside of this specific auction (e.g., this could be the expected cost from rerunning the auction at a later date, or negotiating one-on-one with a particular contractor). Supplier i can complete the project at cost C_i distributed $F_C(\cdot)$ with compact support $[\underline{c}, \bar{c}]$ that admits a continuous density $f_C(\cdot)$. We assume that suppliers are *ex ante* symmetric and have independent private costs. To participate in the project-allocation auction of any mechanism,⁹ a supplier must pay an entry cost K . Entrants learn their true costs of completing the project. One interpretation of K is that it includes the cost of research involved in finding out the true cost, but it will also include any costs associated with preparing a bid. An important assumption that we maintain throughout is that a firm cannot participate in the project-allocation stage of

⁸ For example, Ye (2007) considers a case where bidders are uncertain about a component of costs that can only take on discrete high or low values.

⁹ This will be the auction in the FPAFE and the second-stage auction in an ERA.

any mechanism without paying K .¹⁰ Prior to deciding whether to enter, each supplier observes a signal S_i that is correlated with its true cost C_i and that is not correlated with any other supplier's cost. Signals are affiliated with costs in the standard sense: the cost distribution conditional on a signal s first-order stochastically dominates the cost distribution conditional on a signal $\tilde{s} < s$. We assume that the marginal distribution of signals admits a continuous density on a compact interval $[s_{min}, s_{max}]$.

The LS and S models are limiting cases of this model. If $S_i = C_i$, then signals are perfectly informative of costs, and the setup reduces to the S model. When S_i is independent of C_i , signals contain no information about costs, and the setup reduces to the LS model. For many empirical settings, it seems plausible that buyers will have some, but imperfect, information about their costs prior to conducting costly research, consistent with S_i being positively, but not perfectly, correlated with C_i .

□ **Selective free entry into first-price auctions.** The first model of procurement is the FPAFE, which we view as describing the procedure used by the DoTs in our data and most other government agencies. In the first stage, potential suppliers observe their imperfectly informative cost signals. Based on these private signals, they take independent, simultaneous entry decisions and pay K if and only if they enter. In the second stage, the set of entrants compete in a first-price (low-bid) auction with reserve price r . We assume that firms that do not pay K cannot participate in the auction and that agents bid without knowing how many of their competitors actually entered the auction, as is done in Li and Zheng (2009). If an entrant learns that his true cost exceeds the reserve price, he will not bid.

We solve for entry decisions and bid functions that define the unique symmetric Bayesian Nash equilibrium with monotone bidding behavior.¹¹ Potential bidders enter using a cutoff strategy; that is, a supplier enters if and only if he observes a signal $s_i < s^{*}$ for some critical value of s^{*} (Gentry and Li, 2012). Define $H(c)$ to be the probability (not conditional on S_i) that a given supplier either (i) enters the auction and has a cost no smaller than c , or (ii) does not enter the auction; thus, $H(\cdot)$ depends on the value of the signal cutoff used for the entry decision. The equilibrium bid functions $\beta^*(\cdot)$ are given by the solutions to the optimization problem

$$\beta^*(c) \equiv \arg \max_b (b - c) [H(\beta^{*-1}(b))]^{N-1}.$$

The first-order condition associated with this optimization problem gives the differential equation

$$1 + \beta^{*-1'}(b) (b - \beta^{*-1}(b)) (N - 1) \left[\frac{H'(\beta^{*-1}(b))}{H(\beta^{*-1}(b))} \right] = 0, \tag{1}$$

with the upper boundary condition

$$\beta^*(r) = r. \tag{2}$$

¹⁰ As pointed out by a referee, if K is interpreted as containing an information acquisition cost, then there may be incentives for some suppliers to bid without paying the information acquisition cost. To get around this objection, one can either assume that K is simply a cost that must be paid in order to be responsive in the auction (in which case, we could allow nonentering suppliers to find out their costs as well once they have decided not to pay K) or we can assume that there are some penalties outside of our model that deter this type of uninformed bidding. For example, there may be penalties that arise from defaulting on the contract or incurring financial losses, or it may be harder for an uninformed supplier to find suitable subcontractors.

¹¹ Gentry and Li (2012) establish uniqueness of the symmetric equilibrium in a class of affiliated signal models. As usual in entry games, asymmetric equilibria may exist.

The equilibrium critical cutoff values $s^{/*}$ are determined by the indifference condition that any potential supplier who receives a signal of $s^{/*}$ must be indifferent between entering the auction or not paying the entry cost at all. This zero-profit condition is thus written

$$\int_{\underline{b}}^r (b - \beta^{*-1}(b)) f_C(\beta^{*-1}(b)|s^{/*}) [H(\beta^{*-1}(b))]^{N-1} db = K, \quad (3)$$

where $\underline{b} \equiv \beta(c)$, and $f_C(\cdot|s)$ is the conditional density of a supplier's costs, computed using Bayes' Rule, given he receives a signal s .

To solve for the bid functions, we use the Mathematical Programming with Equilibrium Constraints (MPEC) approach (Su and Judd, 2012). In a manner similar to that outlined by Hubbard and Paarsch (2009), who use MPEC to solve an FPAFE model where bidders know their values when they decide whether to enter, we express the inverse bid function as a linear combination of the first P Chebyshev polynomials (we use $P = 25$), scaled to the interval $[\underline{b}, r]$. The choice variables in our programming problem are, therefore, P Chebyshev coefficients, the signal cutoff, and the value of the low bid (\underline{b}). We pick a fine grid $\{x_j\}_{j=1}^J$ with $J = 500$ points on the interval $[\underline{b}, r]$. Then we solve for the bid functions (more precisely, the Chebyshev coefficients) and the signal cutoff using

$$\arg \min_{\{\underline{b}, \beta^{*-1}, s^{/*}\}} \sum_{j=1}^J g(\beta^{*-1}(x_j))^2 \text{ s.t. (2) and (3),} \quad (4)$$

where $g(\beta^{*-1}(b))$ is defined to be the left-hand side of equation (1). This nonlinear programming problem is solved using the SNOPT solver called from the AMPL programming language. We give more details about the numerical methods in Appendix A.¹²

□ **Entry rights auctions.** When entry is regulated and suppliers have some information about what their costs are likely to be, a straightforward way to choose the firms that will be allowed to bid for the contract is to use an initial auction to allocate a fixed number of rights to participate in a second-stage auction for the contract (Ye, 2007). Before the first stage, the procurer announces the number of suppliers that will be selected to participate in a (second-stage) auction for the contract and the reserve price in that auction. In the first stage of our ERA, each supplier i receives its signal s_i and submits and pays a nonnegative bid, given by a function $\gamma^*(s_i)$, for the right to participate in the second-stage auction. The highest n bidders are selected to participate in the second stage.¹³ In the second stage, the procurer pays each of the selected entrants K to cover their entry costs, and the selected firms incur these costs and find out their true costs of completing the contract. The procurer also publicly announces the value of the $(n + 1)^{\text{st}}$ highest first-stage bid.¹⁴ The second-stage auction uses a first-price format with the preannounced reserve price, r . However, as discussed by Ye (2007) any standard auction format used in the second stage (i.e., one where the firm with the lowest bid wins the contract and a firm with the highest possible cost pays nothing), will generate the same expected payoffs for all parties, and therefore the

¹² We assume that suppliers are *ex ante* symmetric, but this procedure can also be used to solve for equilibria when suppliers belong to certain discrete types. However, a procedure for finding all of the equilibria in an asymmetric FPAFE model remains a topic of ongoing research.

¹³ We assume that there is no reserve price in the first-stage auction, as, given the all-pay nature of the auction, a reserve price might lead some suppliers being unwilling to participate in the mechanism at all. However, it is possible that a first-stage reserve price could be used to improve the performance of the ERA beyond the gains that we currently identify.

¹⁴ Publicly revealing the $(n + 1)^{\text{st}}$ highest bid in the first stage, but not the winning bids, ensures that bidders enter the second stage with symmetric beliefs about the distribution of their rival entrants' costs, which helps to ensure the efficiency of the second-stage allocation and the revenue equivalence of different second-stage formats. Otherwise, beliefs about rivals' costs will depend on an entrant's own signal and so they will not be symmetric. As pointed out by a referee, it is unclear whether the procurement cost or efficiency performance of the ERA would change significantly if this information was not revealed, but it would certainly complicate our analysis.

same incentives in the first stage, as long as entrants have symmetric beliefs about their rivals and the first stage has selected the firms with the lowest costs. We assume a first-price format because it reveals an interesting difference in bidding behavior between the FPAFE and the second stage of an ERA with at least two entrants.¹⁵

The first-stage auction is therefore a *binding all-pay auction with an entry subsidy*. Some explanation for why we model the ERA in this way is required. As shown by Ye (2007) in his Proposition 5, in an all-pay auction, there will be a unique, symmetric pure strategy equilibrium where the first-stage bid function, $\gamma^*(s)$, is a strictly decreasing function of the cost signal, so that a firm with a lower cost signal will bid more, when the entry subsidy is sufficiently large.¹⁶ This monotonicity property is important because it guarantees that the selection of entrants into the second stage will be efficient. An entry subsidy equal to the due diligence cost K is sufficient for this property to hold, and, by Ye's Proposition 6, the level of the subsidy does not affect either the procurer's or the bidders' total payoffs from the mechanism as long as equilibrium first-stage bids are strictly decreasing, as more generous subsidies are exactly offset by higher first-stage bids. A corollary of this argument is that none of our efficiency or procurement cost comparisons would change if we considered slightly lower subsidies that also gave rise to monotonic first-stage bid functions.

In contrast, a strictly decreasing equilibrium first-stage bid function might not exist if we considered a first-round format that operated as either a uniform price auction (e.g., the n selected entrants would each pay the $(n + 1)^{\text{st}}$ highest bid) or a discriminatory auction where only the selected entrants pay their bid. As shown by Ye (2007) in his Proposition 3, first-round bid functions will be strictly decreasing in these formats only if, in the first stage, a bidder's expected gain from participating in the second stage conditional on being the marginal (i.e., n^{th}) selected entrant is strictly decreasing in the cost signal. Whether this will be the case will depend on both n and the specifics of the cost and signal distributions considered, and we have found ranges of signals over which this monotonicity property does not hold for some of the parameters that we consider in this article. However, if the property does hold, then both efficiency and the expected payoffs of all parties should be the same as under the all-pay format, so it is still insightful to use the all-pay format even if it is used less frequently than other formats in the real world.

We are interested in the symmetric Bayesian Nash equilibrium where the first-round bidding strategy $\gamma^*(\cdot)$ is monotone in the signal and the second-round bid function $\beta^*(\cdot; \bar{s})$ is monotone in the cost for every possible value of the revealed $(n + 1)^{\text{st}}$ highest first-round bid. Note that in such an equilibrium, revealing the $(n + 1)^{\text{st}}$ highest bid is equivalent to revealing the $(n + 1)^{\text{st}}$ lowest signal \bar{s} . Thus, entrants in the second round use Bayes' Rule to compute the distribution of the costs of any one of their opponents as $F_{C|S \leq \bar{s}}(\cdot)$, with density $f_{C|S \leq \bar{s}}(\cdot)$.

The second-stage bid function $\beta(\cdot; \bar{s})$ solves the differential equation

$$\beta^{*'}(c; \bar{s}) = (n - 1)(\beta^*(c; \bar{s}) - c) \left[\frac{f_{C|S \leq \bar{s}}(c)}{1 - F_{C|S \leq \bar{s}}(c)} \right],$$

with the boundary condition $\beta^*(r; \bar{s}) = r$. To solve for the bid function $\gamma^*(\cdot)$, note that the profit of a bidder with cost c who is invited to enter the auction, when the $(n + 1)^{\text{st}}$ signal is \bar{s} , is, in the case of a second-round auction operating in a first-price format,

$$\Pi(c; \bar{s}) = (\beta^*(c; \bar{s}) - c) [1 - F_{C|S \leq \bar{s}}(c)]^{n-1},$$

where we are exploiting the fact that the entry cost is being fully subsidized by the procurer.

¹⁵ We exploit payoff equivalence so that we can verify our solutions are accurate by considering a second-stage auction that operates in a second-price format where we do not need to solve for Chebyshev-approximated bid functions.

¹⁶ Fullerton and McAfee (1999) also show how an all-pay auction with an entry subsidy induces efficient entry into the related mechanism of a research tournament. See Ye (2007) for an extensive discussion of the differences between the Fullerton and McAfee model and the type of model considered here.

The first-stage bid function $\gamma^*(s)$ should solve

$$\gamma^*(s) \equiv \arg \max_g \int_{\gamma^{*-1}(g)}^{s_{max}} \int_{\underline{c}}^{\bar{c}} \Pi(c; \bar{s}) dF_{C|S=s}(c) dF_{S^{(n:N-1)}}(\bar{s}) - g,$$

where $F_{C|S=s}(\cdot)$ is the conditional distribution of the cost given a signal of s and $F_{S^{(n:N-1)}}(\cdot)$ is the distribution of the n^{th} highest signal of the remaining $N - 1$ bidders, with pdf $f_{S^{(n:N-1)}}(\cdot)$. Solving this maximization problem and imposing that in equilibrium $g = \gamma^*(s)$ gives the differential equation

$$\gamma^{*'}(s) = - \int_{\underline{c}}^{\bar{c}} \Pi(c; s) f_{S^{(n:N-1)}}(s) dF_{C|S=s}(c),$$

with boundary condition $\gamma^*(s_{max}) = 0$. As entry is subsidized, the integrand on the right-hand side is always positive (a selected firm cannot lose money in the second stage because the entry subsidy covers the entry cost), so the first-stage bid function will be monotonically decreasing in s . The total procurement cost will be equal to the price paid in the second round *plus* the sum of the entry subsidies for the n entrants *less* the sum of the first-stage bids.

The value of n , the number of firms selected for the second stage, plays an important role in the ERA, and is chosen by the procurer. So that the first stage is meaningful, it is natural to assume that $1 \leq n < N$.¹⁷ If $n = 1$, then the procurer selects one firm in the first stage and effectively makes that firm a take-it-or-leave-it offer to complete the project at the second-stage reserve price r (we assume that the chosen firm would still have to pay the entry cost even if it faces no competition in the second stage). Although this case may seem somewhat trivial, this is exactly what an efficiency-maximizing procurer would like to do when bidders are well-informed about their costs, because the firm with the best signal will be the one that should be selected to complete the project with high probability. Therefore, for the rest of this section we allow for the possibility that $n = 1$, but in the empirical application we will restrict ourselves to allowing only $2 \leq n < N$. As we find that the ERA outperforms the FPAFE, changing this constraint could only strengthen our conclusions.

□ **Efficiency of bidder entry.** We now compare the efficiency of the ERA and the FPAFE using an example to illustrate how the advantage of fixing the number of entrants depends on how well informed bidders are about their values before they pay the entry cost. We say that a mechanism is more efficient when the expected social costs of completing the project (the cost of the auction winner or c_0 if there is no winner, plus the sum of entry costs) is lower than in the other mechanism. Recall from the discussion in the Introduction that although entry strategies in the FPAFE are optimal from an efficiency perspective whether or not signals are affiliated with true costs (Levin and Smith; 1994; Gentry and Li, 2012), if a potential bidder's entry strategy is only allowed to depend on that firm's private information, there may be an efficiency advantage to fixing the number of entrants, for example, through an ERA. Doing so will help overcome the problem that conditioning only on a bidder's own private information will result in uncoordinated entry, and possibly too many or too few entrants.

It is straightforward to show that an ERA must be more efficient when either signals are completely uninformative, as in the LS model, or completely informative, as in the S model. In the former case, Milgrom (2004, pp. 225–227) shows this result.¹⁸ In the case of the S model,

¹⁷ Considering a case where all suppliers are invited to participate in the second stage ($n = N$) would involve no competition in the first stage, so all bids would be zero. In essence, the procurer would subsidize the entry costs of all the potential suppliers and then run a standard FPA in which everyone knows their true costs. For the parameters that we look at in Section 5, choosing $n = N$ would not be optimal.

¹⁸ The results reflects the fact that when signals are uninformative, the sum of the expected cost of the lowest cost entrant plus total entry costs will be convex in the number of realized entrants either in an FPAFE or when a fixed number of entrants is chosen randomly by the procurer. Under the LS assumptions, the expected cost of the lowest cost firm among n will be the same in both of these cases.

efficiency requires that only the firm with the lowest cost signal enters (assuming this signal is less than c_0), and this is achieved through an ERA with $n = 1$ but cannot be guaranteed in any mechanism where entry decisions are not coordinated.

The situation is more complicated when signals are partially informative. An ERA, where the number of entrants is fixed in advance, does not allow the number of entrants to depend on the private information that suppliers have about their costs, which is correlated with how valuable their entry is likely to be from an efficiency perspective. To illustrate, consider a specific example where there are four suppliers with costs distributed lognormally with location parameter $\mu_C = -0.09$ and scale parameter $\sigma_C = 0.2$ (so that the mean cost is 0.93 and the standard deviation is 0.19). We assume that c_0 and the reserve in the FPA and the second round of the ERA are equal to 0.85. We truncate the costs to the interval $[0, 4.75]$. We set $K = 0.02$ and $S_i = C_i \cdot \exp(\epsilon_i)$, where $\epsilon_i \sim N(0, \sigma_\epsilon^2)$. Initially, we assume that $\sigma_\epsilon = 0.2$.

For these parameters, the equilibrium signal threshold in the FPAFE would be 0.829 and, on average, 1.46 suppliers would enter. An ERA designed to maximize expected efficiency, without knowledge of what suppliers' signals are, would select a single entrant. To illustrate why the FPAFE might be more efficient, suppose that the signals of the four potential entrants are $\{0.78, 0.79, 0.80, 0.90\}$. *Given these signals*, expected efficiency would be maximized by having the three firms with the lowest signals enter (expected social costs are 0.812), which in this example is what happens in the FPAFE, rather than just the firm with the lowest signal, which is what happens in the ERA (expected social costs are 0.823).

Given these parameters, the FPAFE also dominates in expectation (i.e., integrating out over all possible realizations of suppliers' cost signals). To show how this result relates to the precision of the suppliers' information, the top panel of Figure 1 shows the percent decrease in expected social costs from using an ERA rather than an FPAFE, allowing for either $n = 1, 2,$ or 3 in the ERA. Positive numbers reflect the ERA outperforming the FPAFE, and negative numbers reflect the FPAFE doing better. The solid line indicates the gain from using an ERA with n chosen to maximize efficiency, so that it marks the upper envelope of the lines for the individual ERAs. On the horizontal axis, we measure the precision of information by a parameter $\alpha \equiv \sigma_\epsilon^2 / (\sigma_\epsilon^2 + \sigma_C^2)$ which varies from 0 to 1.¹⁹ Holding σ_C fixed, as $\alpha \rightarrow 1$, the model will tend toward the informational assumptions of the LS model, and as $\alpha \rightarrow 0$, it tends toward the informational assumptions of the S model. The other parameters are the same as in the preceding example. The middle panel shows the expected cost of completing the project for each of these mechanisms (the cost is c_0 if no bidder submits a bid below the reserve price), and the bottom panel shows the expected number of entrants in the FPAFE and the efficiency-maximizing ERA.

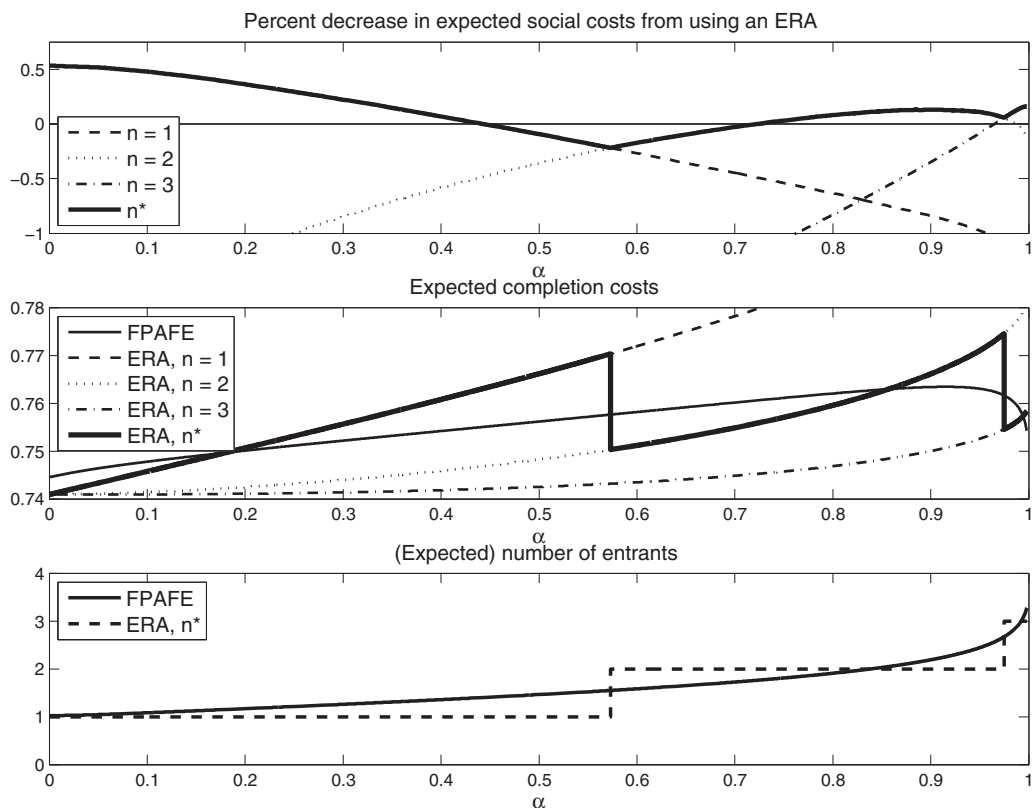
When α is close to 0 or 1, an ERA is more efficient in expectation when n is chosen appropriately. However, for values of α between 0.45 and 0.7, the FPAFE is more efficient than any of the ERAs. Whether the FPAFE does better because it economizes on entry costs or lowers the expected cost of completing the project depends on the value of α , because this affects how many firms the ERA should select for the second stage. When $\alpha = 0.5$, the optimal n is 1 and we are in the case considered above where the additional entry under the FPAFE lowers the expected cost of completing the project by enough to lower social costs. When $\alpha = 0.6$, a different logic applies. In this case, the optimal ERA has $n = 2$, and there is higher expenditure on entry costs and a lower cost to complete the project than in the FPAFE. However, because the firm with the second-lowest signal will be less likely to win in the second-stage auction, especially if it would not have entered the FPAFE, the entry cost effect dominates and the FPAFE is again more efficient.²⁰ The fact that which mechanism is more efficient can depend on the parameters,

¹⁹ If the support of the lognormal distribution were not truncated, then the conditional distribution of the cost given a signal s would be lognormal with location parameter $\alpha\mu_C + (1 - \alpha)\log(s)$ and scale parameter $\sigma_C\sqrt{\alpha}$.

²⁰ As illustration, the supplier with the second lowest—cost signal wins in the second stage of the ERA with probability 0.314 if its cost signal would have led it to enter the FPAFE and with probability 0.178 if it would not have entered the FPAFE.

FIGURE 1

PERFORMANCE OF AN ERA (WITH $n = 1, 2, \text{ OR } 3$) AND AN FPAFE



Notes: We assume $N = 4$, $F_C(c) \sim LN(-0.09, 0.2)$, $K = 0.02$, $r = c_0 = 0.85$, and $S_i = C_i \cdot \exp(\epsilon_i)$. $\alpha \equiv \sigma_c^2 / (\sigma_c^2 + \sigma_\epsilon^2)$. In the top panel, the vertical axis shows the percentage reduction in social costs from using an ERA rather than an FPAFE. Outcomes are calculated using 500,000 simulations.

including the degree of selection, motivates us to estimate our model using real-world data on procurement auctions.

3. Data

■ We estimate our model using a sample of FPAFE procurement auctions for bridge construction projects conducted by the Oklahoma and Texas DoTs from March 2000 through August 2003. These data is part of the sample originally collected and used by De Silva, Dunne, Kankanamge, and Kosmopoulou (2008), and we detail how our sample was chosen in Appendix B. As they describe, the data comes from all areas of Oklahoma, but only the North Texas and Panhandle regions of Texas, so that the geographic, construction, and economic conditions are likely to be similar across the two states.

For each project, the data contain information on the number of “planholders,” which are the construction companies that were interested enough in the project to purchase, at a price of about \$100, the detailed plans for the project developed by the state’s engineer. The list of planholders is publicly available prior to the auction, and we use the number of planholders to define N , the commonly known number of potential suppliers for the project.

Summary statistics for the sample are presented in Table 1. On average, the number of potential suppliers, the number of entrants, the level of unemployment, and the number of entrants (observed bidders in the FPAFE) are significantly (at the 1% level) higher in Texas than

TABLE 1 Summary Statistics for Our Sample of Auctions, by State

	Variable	Mean	Standard Deviation	25 th -tile	50 th -tile	75 th -tile
OK # obs. = 262	Potential suppliers	6.882	1.997	5	7	8
	Number of entrants (bidders)	4.344	1.570	3	4	5
	Unemployment	4.136	0.971	3.2	4.1	4.6
	Engineer's estimate	491,120	543,662	226,338	320,755	523,139
	Winning bid	451,800	520,945	207,146	288,317	471,151
	Relative winning bid	0.917	0.133	0.826	0.888	0.980
TX # obs. = 154	Potential suppliers	7.909	2.027	6	8	10
	Number of entrants (bidders)	4.844	1.657	4	5	6
	Unemployment	5.641	1.214	4.4	5.6	6.5
	Engineer's estimate	1,409,158	1,185,331	550,537	967,817	1,977,169
	Winning bid	1,366,487	1,177,504	530,084	892,737	1,847,180
	Relative winning bid	0.976	0.134	0.892	0.965	1.038
ALL # obs. = 416	Potential suppliers	7.262	2.067	6	7	9
	Number of entrants (bidders)	4.529	1.619	3	4	6
	Unemployment	4.693	1.291	3.8	4.4	5.7
	Engineer's estimate	830,971	949,131	258,485	446,273	914,745
	Winning bid	790,410	936,689	240,553	406,155	887,277
	Relative winning bid	0.939	0.136	0.844	0.914	1.006

Notes: Winning bid and engineer's estimate are in dollars, and the relative winning bid is the winning bid divided by the engineer's estimate.

Oklahoma, but the entry rate (measured by the proportion of planholders that submit bids) is not significantly different across the states. One feature of the data is that the number of entrants can vary significantly across auctions even when the number of planholders is fixed. For example, out of the 51 Oklahoma auctions with six planholders, there are 2, 6, 11, 13, 14, and 5 auctions with 1 through 6 observed bidders, respectively. If entry is costly, which is suggested by the fact that around one third of planholders ultimately chose not to submit bids, this type of variation in the number of entrants suggests that coordinating entry may help to increase efficiency.

Because we will apply the FPAFE model presented in Section 2 to these data, we briefly discuss how far its major assumptions are appropriate in this setting. The assumption of independent private costs is common in the literature studying highway procurement auctions (e.g., Jofre-Bonet and Pendorfer, 2003; Li and Zheng, 2009; Krasnokutskaya and Seim, 2011). We also follow much of the literature in assuming that bidders are *ex ante* symmetric (see, e.g., Li and Zheng, 2009).²¹ Last, our model assumes that there is a commonly known reserve price, which is required to prevent suppliers wanting to submit infinite bids when there is some probability that no other firms enter. Therefore, we assume that the DoT would reject bids that were too high; specifically, we assume that the DoT uses a reserve price equal to 1.5 times the engineer's estimate of the cost of the project, even though the DoTs actual policy for rejecting bids is not public information.²² We choose this relatively high price so that the number of auctions that we have to drop because the DoT accepted bids above this assumed reserve price is very small (two in OK, four in TX).²³ In our data, we also observe that the DoTs reject 28 winning bids that are below this reserve price. However, many of these auctions have winning bids that are well within

²¹ The De Silva, Dunne, Kankanamge, and Kosmopoulou (2008) data does contain an estimate of capacity utilization and bidder distance from the project site, but for the auctions that we look at, these supplier-specific variables have almost no power when trying to predict which supplier wins the auction. For example, in a linear probability model with auction fixed effects, where the dependent variable is a dummy for the winning bidder, these variables are both individually and jointly insignificant and the within-auction R^2 is only 0.0005.

²² An alternative approach (e.g., Li and Zheng, 2009) is to assume that the government acts as an additional bidder if only one firm enters.

²³ If our assumption about the reserve price is correct, we might still expect to see a few bids submitted above the reserve price if a supplier decides to enter and finds out that it has a cost that is greater than the reserve price.

the range that we see the DoTs accepting; for example, the average rejected winning bid is 0.96 times the engineer's estimate, which is close to the average accepted winning bid in the data, and only five of the rejected winning bids are above 1.2 times the engineer's estimate. We assume that these must have been rejected for reasons other than the level of the bid, and so we also exclude these auctions from our estimation sample. In Section 5, we consider whether changing the assumed reserve price could alter our conclusions about how well the ERA and an FPAFE perform.

4. Estimation

■ In this section, we describe how we estimate the model and detail the additional parametric assumptions we make on the specification of the model in order to take it to the data. Our estimation approach is based on matching a set of moments predicted by our model, which we calculate using simulations and importance sampling, to a set of moments observed in our data (e.g., the proportion of auctions where five suppliers enter). Li and Zhang (2013) provide an earlier example of the estimation of a first-price auction model with nonselective entry using a moment-based approach.²⁴

□ **Model setup and specification.** We estimate a fully parametric version of the FPAFE model from Section 2. We normalize the cost parameters by the engineer's estimate of costs, so that a supplier cost of 0.9 means 90% of the engineer's estimate, and an entry cost of 0.03 means that it costs 3% of the engineer's estimate for the entire project to enter the auction.²⁵ Allowing for the a subscript to denote a particular auction, we assume that costs in an auction a are drawn from a truncated lognormal distribution, $F_{C_a}(\cdot)$, with location parameter μ_{C_a} and scale parameter σ_{C_a} . We truncate costs to the interval $[0, 4.75]$, but, as the average winning bid is 0.94, this truncation has essentially no effect on the distribution of costs for plausible parameters. We assume $S_i = C_i \cdot \exp(\epsilon_i)$, where $\epsilon_i \sim N(0, \sigma_{\epsilon_a}^2)$ and ϵ_i independent across suppliers. However, rather than estimating the values of $\sigma_{\epsilon_a}^2$, we estimate α_a where $\alpha_a \equiv \sigma_{\epsilon_a}^2 / (\sigma_{\epsilon_a}^2 + \sigma_{C_a}^2)$.²⁶

As is implicit in the notation, we allow for parametric, auction-specific observed and unobserved heterogeneity in the structural parameters $\theta_a \equiv \{\mu_{C_a}, \sigma_{C_a}, \alpha_a, K_a\}$. Previous research (e.g., Bajari, Hong, and Ryan, 2010; Krasnokutskaya, 2011) has found both types of heterogeneity in the means of the cost distribution for road construction contracts, and it is plausible that the variance of these distributions and the level of entry costs are also heterogeneous. Note, however, that we assume that within an auction bidders are symmetric in that they draw their project costs and signal noise from the same distributions and that they have the same cost of entering the auction. The parameters for auction a are drawn from independent truncated normal distributions with means that depend on observed covariates X_a . Letting $TRN(\mu, \sigma^2, \underline{c}, \bar{c})$ denote a normal distribution with mean μ and variance σ^2 and truncation points $[\underline{c}, \bar{c}]$, we assume that

$$\begin{aligned}\mu_{C_a} &\sim TRN(X_a \beta_{\mu_C}, \omega_{\mu_C}^2, \underline{c}_{\mu_C}, \bar{c}_{\mu_C}) \\ \sigma_{C_a} &\sim TRN(X_a \beta_{\sigma_C}, \omega_{\sigma_C}^2, \underline{c}_{\sigma_C}, \bar{c}_{\sigma_C}) \\ \alpha_a &\sim TRN(X_a \beta_{\alpha}, \omega_{\alpha}^2, \underline{c}_{\alpha}, \bar{c}_{\alpha}) \\ K_a &\sim TRN(X_a \beta_K, \omega_K^2, \underline{c}_K, \bar{c}_K).\end{aligned}\tag{5}$$

²⁴ Specifically, Li and Zhang (2010) use an indirect inference method where they match parameters with those from an auxiliary model. Li (2010) provides additional information on this type of approach.

²⁵ It is natural to assume that larger projects will require more due diligence and may also require the bidders to post larger bonds that guarantee that they will complete the work. Of course, the imposed linearity of the entry cost with the engineer's estimate is an assumption.

²⁶ Estimating the distribution of α_a means that we are directly estimating how the degree of selection differs across auctions. Understanding this variation is less straightforward when estimating separate distributions of σ_{ϵ_a} and σ_{C_a} .

In our setting, X_a is a vector consisting of a constant, the unemployment rate in the county where the project is located in the month that the auction is held (in percentage points), and a dummy variable equal to 1 for auctions in Texas. We denote the coefficients for these covariates as $\beta_{0,\times}$, $\beta_{1,\times}$, and $\beta_{2,\times}$, respectively, where \times can be μ_C , σ_C , α , or K . We allow only unemployment to affect the location parameter of the cost distribution, setting $\beta_{1,\sigma_C} = \beta_{1,\alpha} = \beta_{1,K} \equiv 0$. The Texas dummy coefficient could capture differences in the procurement process across the states, which might affect entry costs or how much suppliers know about their costs, as well as the costs of completing the project.

To make our estimation approach work, we assume that we, the researchers, know the truncation points.²⁷ For the cost parameters, we choose very wide bounds that should have little effect on the parameter distributions for plausible parameters. For example, we truncate K_a at 0.04% and 16% of the engineer’s estimate, where the industry literature (cited in our discussion of the results below) indicates that values in the range 1%–2% are likely. The parameters to be estimated are $\Gamma \equiv \{\beta_{\mu_C}, \beta_{\sigma_C}, \beta_\alpha, \beta_K, \omega_{\mu_C}^2, \omega_{\sigma_C}^2, \omega_\alpha^2, \omega_K^2\}$.

□ **Importance sampling.** We estimate the model using a simulated method-of-moments estimator and we use importance sampling to approximate the moments predicted by a given set of parameters (Ackerberg, 2009). The motivation behind using importance sampling is that it would be too costly to resolve a large number of FPAFE models each time one of the parameters changes. Instead, with importance sampling, we can solve a very large number of games once and then only reweight these outcomes as we change the parameters.

To explain the way that we use importance sampling, let $y^e = f(X_a, \theta_a)$ denote an expected outcome of an auction, such as an indicator for whether the winning bid lies between 0.8 and 0.9 times the engineer’s estimate. Calculating y^e in our setting is expensive because we have to solve the FPAFE and then simulate outcomes. The density of θ , given X_a and Γ , is $\phi(\theta|X_a, \Gamma)$. The expected value of the outcome given X_a and Γ , which we can label $E(y^e|X_a, \Gamma)$, is

$$E(y^e|X_a, \Gamma) = \int f(X_a, \theta_a)\phi(\theta_a|X_a, \Gamma) d\theta_a.$$

As this integral does not generally have an analytic form, one could approximate it for a given value of Γ using simulation by drawing S samples of θ from the distribution $\phi(\theta|X_a, \Gamma)$ and computing $f(X_a, \theta_{as})$ for each draw θ_{as} so that

$$E(y^e|X_a, \Gamma) \approx \frac{1}{S} \sum_{s=1}^S f(X_a, \theta_{as}).$$

However, this calculation would require solving a new set of S auctions whenever one of the parameters in Γ changes. The importance sampling methodology exploits the fact that

$$\int f(X_a, \theta_a)\phi(\theta_a|X_a, \Gamma) d\theta = \int f(X_a, \theta_a) \frac{\phi(\theta_a|X_a, \Gamma)}{\psi(\theta_a|X_a)} \psi(\theta_a|X_a) d\theta,$$

where $\psi(\theta_a|X_a)$ is an importance sampling density that has the same support as $\phi(\theta_a|X_a, \Gamma)$, but does not depend on unknown parameters. In this case, we can calculate $f(X_a, \theta_{as})$ ²⁸ for a large number of draws taken from $\psi(\theta_a|X_a)$, and then calculate an approximation to $E(y^e|X_a, \Gamma)$, where we only have to recompute the weights $\frac{\phi(\theta_{as}|X_a, \Gamma)}{\psi(\theta_{as}|X_a)}$ when Γ changes, that is,

$$E(y^e|X_a, \Gamma) \approx \frac{1}{S} \sum_{s=1}^S f(X_a, \theta_{as}) \frac{\phi(\theta_{as}|X_a, \Gamma)}{\psi(\theta_{as}|X_a)}. \tag{6}$$

²⁷ As noted by Ackerberg (2009), we would not be able to use importance sampling as part of our estimation approach if the truncation points depended on the parameters that we are estimating.

²⁸ For each solved auction (i.e., for a given draw of θ_{as}) we use a single simulation as an unbiased estimator of $f(X_a, \theta_{as})$. We could, of course, use more draws to increase efficiency.

TABLE 2 Importance Sampling Density Parameters ($\tilde{\Gamma}$) for Estimation

	β_0 Constant	β_1 Unemployment	β_2 Texas	ω Standard Deviation	$\underline{\xi}$ Lower Trunc.	\bar{c} Upper Trunc.
μ_C	-0.07	-0.005	0	0.012	-0.4	1.0
σ_C	0.05	0	0	0.025	0.0095	0.995
α	0.5	0	0	0.25	0.1	0.9
K	0.015	0	0	0.01	10^{-4}	0.16

Notes: The parameter $\tilde{\Gamma}$ specifies the distribution, as given by the functional form in (5).

To estimate the parameters, we group the auctions observed in the data into 32 groups, based on the interaction of the number of planholders ($N = 4, \dots, 11$), the state (OK or TX), and whether the unemployment rate in the county is above or below the median level for the state in the month that the auction was held. For each of these groups, we create moments that measure the difference between average observed outcomes (y_a) and expected outcomes ($E(y^e|X_a, \Gamma)$) for the auctions in the group. The outcomes consist of indicator variables for the number of firms submitting bids, indicators for whether the winning bid lies in one of 15 discrete bids (we divide the interval $[0, 1.5]$ into 15 equally sized bins), and the proportion of all suppliers' bids that lie in each of these bins (as some suppliers do not enter, the sum of the proportions across the bins for a given group will be less than 1).

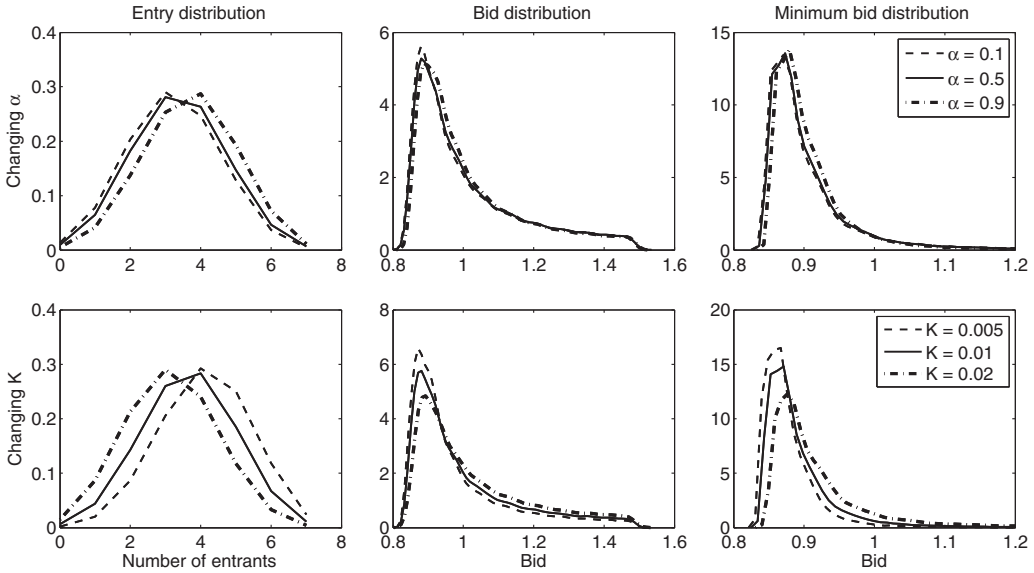
We estimate Γ by minimizing the squared sum, across groups and moments, of these differences, which is a consistent method of moments procedure where the different moments and the different groups receive equal weighting. We use $\psi(\theta_a|X_a) \equiv \phi(\theta_a|X_a, \tilde{\Gamma})$, where $\tilde{\Gamma}$ is given in Table 2 and $\phi(\cdot)$ denotes the truncated normal model given in (5), as our importance sampling density. For each group of auctions, we take NS draws of θ from this distribution and solve the associated FPAFE model, calculating $y^e = f(X_a, \theta_a)$ using simulation, where NS is equal to 100 times the number of auctions in the group. We choose 50 of these simulated auctions for each of our auctions to calculate $E(y^e|X_a, \Gamma)$ using equation (6), as part of the estimation procedure. In Appendix C we provide some Monte Carlo evidence that an importance sampling-based estimator can perform well for this model even using a smaller number of draws (as few as five) for each auction. As is usual when using importance sampling, accuracy is improved by using an importance sampling density which is close to the true density. In practice, we chose the starting parameters reported in Table 2 based on a large number of initial runs, using more diffuse importance sampling densities, which indicated that these parameters values would be close to the parameters that we would estimate in both states.

We calculate standard errors using a bootstrap procedure. For a given bootstrap replication, we redraw auctions from each of our 32 groups with replacement. For each of these auctions, we choose a new set of 50 simulated auctions, from our large sample of NS for that group, to use in the importance sampling calculation. In this way, the reported standard errors should account for the fact that our estimates are dependent on the particular set of importance sampling draws that we use to calculate the moments.

□ **Identification.** Gentry and Li (2014) formally study nonparametric identification in a broad class of selective entry auction models where potential bidders first simultaneously decide whether to enter the model based on signals about their values and then compete in a standard auction, such as a first-price auction. With no unobserved heterogeneity, they show that exogenous sources of variation in the equilibrium level of entry can identify the model. When the sources of variation in entry are sufficiently rich (e.g., a continuous variable that shifts entry costs or a continuously varying reserve price with suitable support), one obtains point identification, whereas when there is only discrete variation (e.g., in the number of potential entrants), one obtains partial identification, although they show that the resulting bounds are often informative.

FIGURE 2

SIMULATED ENTRY, BID, AND WINNING BID DISTRIBUTIONS FOR A REPRESENTATIVE AUCTION FOR VARIOUS VALUES OF α AND K



Notes: The top row plots these distributions as α moves from 0.1 to 0.5 to 0.9. the bottom row plots them as K moves from 0.005 to 0.01 to 0.02. we set the location of the value distribution equal to -0.0963 , the scale parameter equal to 0.0705 , α (when we vary K) equal to 0.4979 , and K (when we vary α) equal to 0.0147 . densities are plotted based on 100,000 simulations of each auction.

One intuition for why the model is identified with no unobserved heterogeneity comes from the fact that when there are few potential entrants or entry costs are very low, all potential bidders should choose to enter with high probability. In this case, we have exogenous entry and more standard results for the identification of value or cost distributions in auction models will hold (Athey and Haile, 2002). The average level of entry costs and the degree of selection (equivalently, the informativeness of suppliers' signals) will then be identified from how the amount of entry and the distribution of bids change as the number of potential entrants rises.²⁹ If not all potential entrants participate in the auction for larger values of N , entry costs must be nontrivial. Furthermore, as shown by Levin and Smith (1994) in the corollary to their Proposition 9, one implication of the random entry, discussed in the Introduction, that occurs with no selection ($\alpha = 1$) is that the expected procurement cost will rise with the number of potential entrants given optimal or nonstrategic reserve prices. When $\alpha < 1$, the sign of the relationship will depend on the parameters (Samuelson, 1985, discusses the $\alpha = 0$ case). In our sample, controlling for state fixed effects and unemployment, there is a significant (at the 10% level) negative relationship between the log of the procurement cost relative to the engineer's estimate and the number of potential entrants, indicating that a model with no selection can be rejected. In estimation, the value of $\hat{\alpha}$ will be chosen to not only match this relationship, but also, given the estimated value of \hat{K} , how the distribution of all bids and the distribution of the number of entrants change with both N and observables such as unemployment.

To illustrate how α and K jointly determine observed outcomes, the two rows of Figure 2 show how the distributions of the number of entrants, all bids and winning bids, change as we change α from 0.1 (a high degree of selection) to 0.5 to 0.9 (a low degree of selection) and K

²⁹ Variation in the number of potential entrants is a common source of identification in entry models (Berry and Tamer, 2006).

from 0.005 to 0.01 to 0.02, with fixed N . We draw these figures holding the other parameters fixed at their mean estimated values for auctions in Oklahoma with an unemployment rate of 3%, and we assume that there are seven potential entrants.³⁰

As α increases, there is more entry on average (potential suppliers place less weight on their signals when forming their posteriors so that suppliers with high [bad] signals are more willing to enter). However, suppliers with very low costs become less likely to enter (it is more likely that they receive a bad signal), and those that do enter tend to face less competition from other low-cost entrants so that their markups tend to increase. This tends to shift the distribution of winning bids slightly to the right. The distribution of all bids also shifts slightly to the right, as more high cost suppliers tend to enter. As entry costs increase, fewer firms enter but, because of selection, the firms that do enter tend to have lower costs. On its own, this would tend to shift the distribution of the bids to the left. However, because they expect to face less competition, entering suppliers will tend to submit higher markups. The increase in markups tends to shift the distribution of bids, as well as the distribution of winning bids, to the right, and, in this case, we can see that this effect actually tends to dominate. As we shall see in Section 5 below, the markups that suppliers with moderately high costs submit in an FPAFE play an important role in our counterfactual comparison with an ERA.

Gentry and Li (2014) show that their identification results extend to a case where there is unobserved heterogeneity that affects both the distribution of values/costs and the level of entry costs, as long as there is still some observed exogenous source of variation in the equilibrium level of entry. Our model, where we allow for independent sources of unobserved heterogeneity to affect the distribution of costs, entry costs, and the degree of selection is not included in the class of models that they consider and our parametric assumptions are likely to be important. To provide some additional transparency into how the moments that we use in estimation identify the parameters of our model, we follow a recent article by Gentzkow and Shapiro (2013). They suggest a method for illustrating how sensitive (locally) estimated parameters are to particular groups of moments that are defined by a structural model.

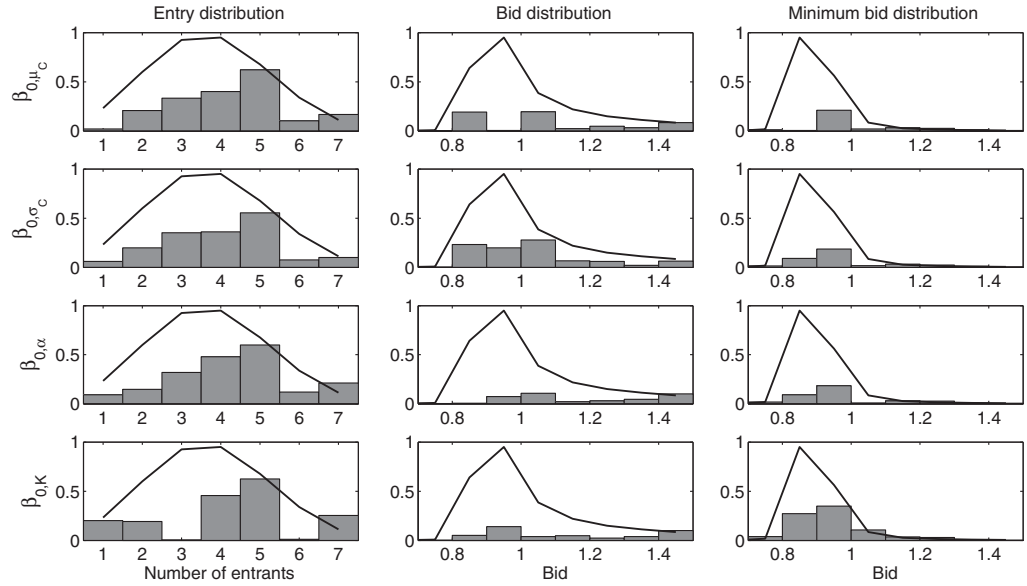
The eight rows of Figure 3 show the absolute values of Gentzkow and Shapiro's scaled sensitivity parameters Λ for moments based on $N = 7$ potential suppliers and low unemployment in Oklahoma, and the parameter estimates that will be presented in Section 5. Gentzkow and Shapiro define their sensitivity parameter as $\Lambda = \Sigma_{\theta\gamma} \Sigma_{\gamma\gamma}^{-1}$, where θ are the parameters, γ are the moments, and Σ represents the variance-covariance matrix of the stacked vector of the parameters and the moments evaluated at the estimated parameters. This scaling means that we can interpret the vertical axis values as telling us how much a one standard deviation increase in a given moment would change the parameter relative to its asymptotic standard deviation.³¹ Figure 3(a) shows the scaled sensitivity parameters for the constant ($\beta_{0,x}$) parameters, and Figure 3(b) shows them for the standard deviation (ω_x) parameters. In each panel, we show the scaled sensitivity parameters for each of the moments. The first column considers the entry moments, with each bar corresponding to the moment that measures the probability of seeing a particular number of entrants. The second column considers the bid distribution, and each bar corresponds to a moment that measures the probability that a supplier's bid lies in the corresponding bin (e.g., 0.7–0.8, 0.8–0.9, etc.). The final column considers the distribution of the winning bid. Figure 3 also overlays the simulated values of the associated moments, at the estimated parameters, to put the horizontal axis into perspective; these moments are scaled so that they fit on the figures and the value of the moment should not be read from the vertical axis. All distributions are conditional on at least one entrant, as this is what we observe in the data.

³⁰ Specifically, we set the location of the value distribution equal to -0.0963 , the scale parameter equal to 0.0705 , α (when we vary K) equal to 0.4979 , and K (when we vary α) equal to 0.0147 . The winning bid distribution is drawn conditioning on at least one potential bidder entering.

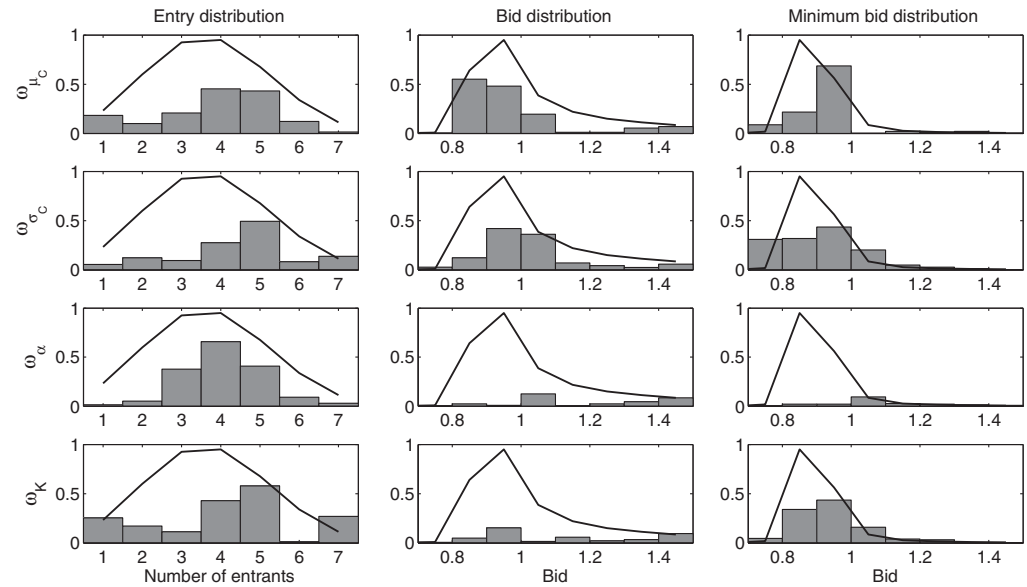
³¹ When interpreting the values on the y-axis, it is important to remember that the combination of $N = 7$, OK, and low unemployment is only one of the 32 N -state-unemployment level combinations that we use in estimation. The patterns are broadly similar but not identical for the other groups.

FIGURE 3

ABSOLUTE VALUES OF THE SCALED SENSITIVITY PARAMETER FROM GENTZKOW AND SHAPIRO (2013)



(a) Changing β_0



(b) Changing ω

Notes: For Oklahoma auctions with low unemployment and seven potential suppliers. Each row corresponds to a different parameter ($\beta_{0,x}$ in (a) and ω_x in (b)), and each column corresponds to a different distribution. The scaled sensitivity parameters for each moment (i.e., each bin of the distribution) are represented by the gray bars and the value of the scaled sensitivity parameter can be read from the vertical axis. The solid lines show the relative value of the different moments evaluated at the different parameters.

TABLE 3 Parameter Estimates

	β_0 Constant	β_1 Unemployment	β_2 Texas	ω Standard Deviation	OK Mean	TX Mean
μ_C	-0.0868 (0.0324)	-0.0023 (0.0064)	0.0154 (0.0132)	0.0142 (0.0091)	91.1% (0.84%)	92.1% (1.89%)
σ_C	0.0687 (0.0196)		-0.0117 (0.0213)	0.0304 (0.0132)	6.45% (0.97%)	5.61% (0.87%)
α	0.4979 (0.0972)		0.1115 (0.0943)	0.1284 (0.0764)	0.4979 (0.0770)	0.6055 (0.0752)
K	-0.0018 (0.0406)		0.0105 (0.0191)	0.0189 (0.0109)	1.47% (0.26%)	1.88% (0.28%)

Notes: Standard errors are in parentheses and are calculated using a non-parametric bootstrap described in the text. The state averages are (going down the rows) the average of costs, the mean standard deviation of the cost distribution, the mean α and the mean entry cost, relative to engineer's estimate, for auctions in each state.

The sensitivity parameters show two intuitive patterns. First, the parameters tend to be most sensitive to the moments that capture the most observations in the data (e.g., four suppliers enter, rather than one or seven suppliers). Second, all of the parameters are somewhat sensitive to all of the moments, reflecting how, as explained above, a parameter such as the level of entry costs will affect not only the cost distribution of the suppliers that enter and how much entrants markup their bids. That said, the figures suggest that for the four parameters related to α and K , the entry moments are particularly important in the sense that their sensitivity parameters are larger. On the other hand, for ω_{μ_C} and ω_{σ_C} , the moments describing the distributions of either all bids or winning bids play a relatively greater role. This is intuitive, as, for example, increasing both the number of auctions where suppliers tend to have high costs and the number of auctions where they tend to have low costs, which is what happens as ω_{μ_C} increases, should tend to spread out the distribution of the winning bid without necessarily affecting the average amount of entry. Of course, one could take these results further to try to identify more optimal moments for estimation (e.g., interactions of the amount of entry and the winning bid), but we view this type of extension as a topic for further research.

5. Results

■ In this section, we present the estimation results and our counterfactual analysis of the impact of switching from the observed FPAFE format to an ERA format.

□ **Parameter estimates.** The left-hand columns of Table 3 present the parameter estimates. To interpret the location and scale parameters together, the right-hand columns report the average across the auctions in each state of (going down the rows) the mean of the cost distribution, the standard deviation of the cost distribution, the value of α , and the entry cost. The mean of the cost distributions are similar in the two states, between 91% and 92% of the engineer's estimate. As the average *winning* bids are higher than this, it suggests that the markups may be quite substantial. We will return to why this happens, and how the ERA can help to reduce markups, in discussing our counterfactual results.

Mean entry costs are 1.5% of the engineer's estimate of the cost of the project in Oklahoma and 1.9% in Texas. Although in many industries it is hard to know what entry costs are really reasonable, in this setting we are able to compare them with estimates from manuals that are used in the industry. Halpin (2005) estimated that the cost of researching and preparing bids is around 0.25% to 2% of total project costs. Park and Chapin (1992) estimated that these costs are typically about 1% of the total bid. Our estimates therefore appear quite reasonable.

Our estimates of entry costs are also lower than those that have been estimated based on models with no selection. For example, based on highway paving contracts in California,

Krasnokutskaya and Seim (2011) estimates mean entry costs to be approximately 3% of the engineer's estimate, and Bajari, Hong, and Ryan (2010) estimates them to be 4.5% of the engineer's estimate. There are some intuitions for why allowing for selection will tend to lower estimated entry costs. In any model, the level of the entry cost will be identified from the fact that the marginal entrant must be indifferent between getting some expected surplus if it enters and not paying the entry cost. Holding the entry probability of other suppliers and the distribution of costs fixed, the expected surplus of the marginal entrant, and therefore the estimated entry cost, will tend to be lower with selection for two reasons. First, with selection, the other suppliers who choose to enter will tend to have lower costs than if their costs were randomly drawn, reducing the marginal entrant's surplus. Second, as pointed out by one of our referees, because a bidder's surplus is usually convex in its own type, the marginal entrant's expected surplus will tend to be lower when it is conditioning on an informative signal than when it has no signal.

The estimates of α suggest that entry is partially selective in this setting, with a mean α of around 0.50 in Oklahoma and 0.61 in Texas; these mean values of α correspond to the scale parameter σ_c of the error distribution being approximately σ_c in Oklahoma and 1.3 times σ_c in Texas. The degree of selection is therefore in the range where our example suggests that it is possible that the FPAFE may be more efficient than an ERA, even though entry into FPAFEs is clearly volatile in our empirical setting.

□ **Regulating bidder entry with an entry rights auction.** In this subsection, we use our parameter estimates to quantify how procurement costs and efficiencies would change if the DoTs switched from using FPAFEs to using ERAs. Throughout what follows, the procurement cost refers to the net amount that the procurer expects to pay to get the project completed. In an FPAFE, this is just the expected winning bid, or c_0 if there is no winner. In the ERA, it is the expected winning second-round bid (or c_0) plus the sum of entry costs less the expected sum of first-round bids. Social costs refer to the expected sum of entry costs plus the cost of the supplier completing the project (or c_0), and a mechanism is more efficient if and only if it has lower social costs. We assume that the cost of completing a project outside the auction is 1.5 times the engineer's estimate and, unless we state otherwise, this is used as the reserve price in the FPAFE and the second stage of the ERA.

In presenting the results, we assume the DoTs would choose the number of entrants in the ERA to minimize the procurement cost (n_{cost}^*), but we will comment below on how the results are very similar if we assume instead that the number of entrants is chosen to minimize social costs (n_{eff}^*). The choice of n_{cost}^* balances several competing effects. In equilibrium, the procurer pays the full entry cost of an additional entrant, and increasing the number of entrants will reduce the amount that any supplier bids in the first-stage auction. On the other hand, when there are more entrants, a firm with lower costs may be selected to complete the project and bidding in the second-stage auction will tend to be more competitive. With the exception of the additional entry cost, the size of these different effects will depend on the degree of selection, as an additional entrant will be more likely to win when signals are less informative.

Tables 4 and 5 present the results. Table 4 compares the outcomes of interest (social costs, procurement costs, and the profits for an individual supplier), and Table 5 shows the average cost of the winner, the average winning bid, and the average amount of entry under each mechanism. For the ERA, it also shows the average total expenditure on entry costs and the average revenue from first-stage bids. In both tables, we use a range of different parameter values. The baseline case (top row) corresponds to the mean parameters for auctions in Oklahoma (unemployment rate 4.14%) and approximately the mean number of potential suppliers in that state (seven). In the remaining rows, we change the number of potential suppliers up and down by approximately one standard deviation and each parameter, in turn, to be equal to either the 10th or the 90th percentiles of its distribution. The parameter that is changed from the baseline is listed in italics.

For all of the parameters listed, we predict that an ERA would both be more efficient and lower procurement costs. For the baseline parameters, both social costs and procurement costs

TABLE 4 Social Costs, Procurement Costs, and Mean Profits for an Individual Supplier Relative to the Engineer's Estimate

Case	N	Parameters			Social Costs			Procurement Costs			100 × Avg. Supplier Profits			
		μ_c	σ_c	α	K	FPAFE	ERA	% Diff.	FPAFE	ERA	% Diff.	FPAFE	ERA	% Diff.
<i>Baseline</i>	7	-0.0963	0.0705	0.4979	0.0147	0.886	0.865	-2.41%	0.914	0.892	-2.41%	0.399	0.391	-2.01%
<i>Low N</i>	5	-0.0963	0.0705	0.4979	0.0147	0.891	0.874	-1.91%	0.923	0.905	-1.95%	0.648	0.636	-1.85%
<i>High N</i>	9	-0.0963	0.0705	0.4979	0.0147	0.883	0.859	-2.74%	0.908	0.883	-2.75%	0.280	0.277	-1.07%
<i>Low μ_c</i>	7	-0.1145	0.0705	0.4979	0.0147	0.871	0.850	-2.47%	0.899	0.876	-2.56%	0.397	0.384	-3.27%
<i>High μ_c</i>	7	-0.0781	0.0705	0.4979	0.0147	0.901	0.880	-2.36%	0.930	0.908	-2.37%	0.404	0.397	-1.73%
<i>Low σ_c</i>	7	-0.0963	0.0335	0.4979	0.0147	0.926	0.902	-2.54%	0.940	0.915	-2.65%	0.201	0.192	-4.48%
<i>High σ_c</i>	7	-0.0963	0.1081	0.4979	0.0147	0.848	0.829	-2.21%	0.890	0.870	-2.25%	0.599	0.577	-3.67%
<i>Low α</i>	7	-0.0963	0.0705	0.3339	0.0147	0.884	0.861	-2.64%	0.915	0.891	-2.62%	0.441	0.439	-0.45%
<i>High α</i>	7	-0.0963	0.0705	0.6620	0.0147	0.889	0.871	-2.12%	0.914	0.893	-2.29%	0.348	0.328	-5.75%
<i>Low K</i>	7	-0.0963	0.0705	0.4979	0.0024	0.839	0.837	-0.21%	0.870	0.867	-0.34%	0.442	0.418	-5.42%
<i>High K</i>	7	-0.0963	0.0705	0.4979	0.0303	0.934	0.896	-4.11%	0.961	0.923	-3.95%	0.380	0.390	2.63%

Notes: We compare the FPAFE to an ERA with the number of entrants chosen to minimize procurement costs. Note that lower numbers for social costs mean that a mechanism is more efficient. The “% Diff.” columns correspond to the percent change in the quantity when moving from an FPAFE to an ERA. Average outcomes are estimated using 500,000 simulations.

TABLE 5 Breakdown of the Expected Values of Major Quantities in an FPAFE and an ERA

Case	FPAFE			ERA			Total Entry Costs	Total 1 st Stage Bids
	Comp. Cost	Win Bid	No. Entrants	Comp. Cost	Win Bid	n_{cost}^*/N		
<i>Baseline</i>	0.837	0.914	3.381	0.835	0.891	2/7	0.029	0.028
<i>Low N</i>	0.845	0.923	3.125	0.844	0.902	2/5	0.029	0.025
<i>High N</i>	0.831	0.908	3.528	0.829	0.884	2/9	0.029	0.029
<i>Low μ_C</i>	0.821	0.899	0.888	0.820	0.875	2/7	0.029	0.027
<i>High μ_C</i>	0.852	0.930	0.921	0.851	0.907	2/7	0.029	0.029
<i>Low σ_C</i>	0.877	0.940	0.905	0.873	0.900	2/7	0.029	0.013
<i>High σ_C</i>	0.797	0.890	0.907	0.800	0.882	2/7	0.029	0.042
<i>Low α</i>	0.835	0.915	3.311	0.831	0.881	2/7	0.029	0.019
<i>High α</i>	0.838	0.914	3.479	0.841	0.902	2/7	0.029	0.038
<i>Low K</i>	0.828	0.870	4.592	0.830	0.843	3/7	0.007	0.013
<i>High K</i>	0.848	0.961	2.856	0.835	0.891	2/7	0.061	0.028

Notes: The cases correspond to the parameters listed in Table 4. The “Comp. Cost” columns report the expected completion cost, which is the winner’s cost when the project is awarded to a bidder and $c_0 = 1.5$ when there is no winner. The “Win Bid” columns are the average bid of the winning bidder in the auction for the contract (this is identically the same as the procurement cost in the FPAFE). The “Total 1st Stage Bids” column lists the sum of the expected first-round bids in the ERA. The column n_{cost}^*/N lists the procurer-optimal number of entrants selected in the ERA. Average outcomes are estimated using 500,000 simulations.

are 2.4% lower under the ERA. The magnitudes are broadly similar across the parameters, except when entry costs are low (*Low K*) when both the absolute and the proportional differences are small. Aggregating across all of the auctions in both states, by taking a draw of the parameters for each of the auctions in our sample, we predict that using ERAs would lower social costs by \$7.78 million (2.49%) and lower procurement costs by \$8.03 million (2.50%). For 3.5% of these draws, we predict that the FPAFE would be more efficient, but in all of these cases entry costs are very low, α has values between 0.39 and 0.77, the efficiency advantage of the FPAFE is small, and we also predict that the ERA would lower procurement costs, although the procurement cost differences are small in these cases.

Although improvements of around \$8 million may seem fairly small, it is important to remember that the sample we are using (contracts involving bridge work in Oklahoma and a small section of Texas) are a very small fraction of highway contracts let by US DoTs in the years that we study: in 2007, the public sector spent \$146 billion to build, operate, and maintain highways in the US with three quarters of this amount spent by state and local governments (Congressional Budget Office, 2011). Moreover, the proportional effects that we are finding are large relative to those produced by auction design changes that are often considered. For example, for the baseline parameters setting a strategic reserve to minimize procurement costs in the FPAFE would only lower procurement costs by 0.03%, whereas we predict that an ERA would lower procurement costs by 2.4%.³²

Interestingly, the ERA also reduces procurement costs by more than adding additional potential suppliers, which is often viewed as providing an upper bound on what an optimal design change can achieve.³³ This can be seen in Table 4, where adding two more potential suppliers in

³² We calculate the optimal reserve price by computing procurement costs on a fine grid of reserve prices with 0.01 spacing, using 500,000 simulations to calculate expected costs in each case. Efficiency in the FPAFE falls when a reserve price other than 1.5 (the value of c_0) is used. We have also found that an optimal reserve price in the second stage of the ERA can only improve its performance by a trivial amount.

³³ Bulow and Klemperer (1996) show that in an IPV model with a fixed number of symmetric bidders, an additional bidder is more valuable to the seller than being able to implement an optimal auction design. This theoretical result does

an FPAFE to the baseline case (changing it to the *High N* case) lowers procurement costs by 0.6% of the engineer's estimate, compared to a 2.4% reduction from switching to an ERA.

Table 5 shows that in most cases the ERA is more efficient because there is both less entry and the average cost of the winner is lower. The fact that it can have both of these advantages reflects the uncoordinated nature of entry in the FPAFE: even though more firms may enter an FPAFE on average, there may be instances where no firms, or only one firm, enter(s) and at least one additional entrant, who will always be in the ERA as $n_{\text{cost}}^* = 2$ for these parameters, might decrease the expected cost of completing the project substantially.³⁴ However, most of the ERA's efficiency advantage comes from lower entry costs (e.g., this accounts for 91% of the reduction in social costs in the base case). This fact helps to explain why the one case in Table 4 where the social costs are almost identical under the two mechanisms is where entry costs are very low (*Low K*, where K equals just 0.2% of the engineer's estimate). In this case, the additional entry into the FPAFE lowers the expected cost of completing the project although increasing entry costs by a relatively small amount. However, even in this *Low K* case, the ERA is still more efficient.

A natural question is how well the ERA that we consider does relative to a mechanism that was designed to maximize efficiency. As already noted, one weakness of our ERA is that the number of selected entrants does not depend on suppliers' signals. Another weakness, relative to a sequential search procedure, is that the procurer cannot get bids from more firms if the ones for whom it initially chooses to pay the entry cost turn out to have high costs. Based on our results so far, it is unclear whether these weaknesses are important. To investigate this further, we can compare social costs under our ERA and under two more efficient but hypothetical alternatives.³⁵ In the first alternative, which we can call the "signal-dependent ERA" (S-D ERA), we allow the procurer to choose the number of entrants with knowledge of suppliers' signals but without using any information on the costs of the selected entrants. In the second alternative, which we can call the "sequential search procedure" (SSP), we consider a procurer who knows the suppliers' signals and then approaches suppliers sequentially, in ascending order of their signals. The procurer asks an additional supplier to enter if, based on its signal and the lowest cost to complete the project of the suppliers that have entered already, the procurer expects that its entry will reduce the cost of completing the contract by enough to offset the additional entry cost.

For the baseline parameters, we find that our ERA performs very well compared to these hypothetical alternatives. Social costs for the FPAFE, ERA, S-D ERA, and SSP, are 0.886, 0.865, 0.863, and 0.857, respectively. The ERA with fixed n therefore achieves 91% and 72% of the efficiency gains of switching from the FPAFE to the S-D ERA and SSP, respectively, while maintaining the attractive elements of simultaneous bidding and being quite straightforward to implement. Qualitatively, these comparisons also hold for the other parameters in Table 4: with the exception of the *Low K* case, where all mechanisms have very similar social costs, the fixed n ERA achieves at least 78% of the gains that would result from switching from the FPAFE to the S-D ERA, and at least 60% of the gain from switching from the FPAFE to the SSP.

As can be seen in Table 4, the benefits of the increased efficiency of the ERA often accrue entirely to the procurer, with the profits of the average supplier tending to fall slightly (the percent changes in the supplier's profits are nontrivial because supplier profits are quite small under either mechanism), whereas procurement costs tend to fall by slightly more than the gain in efficiency.

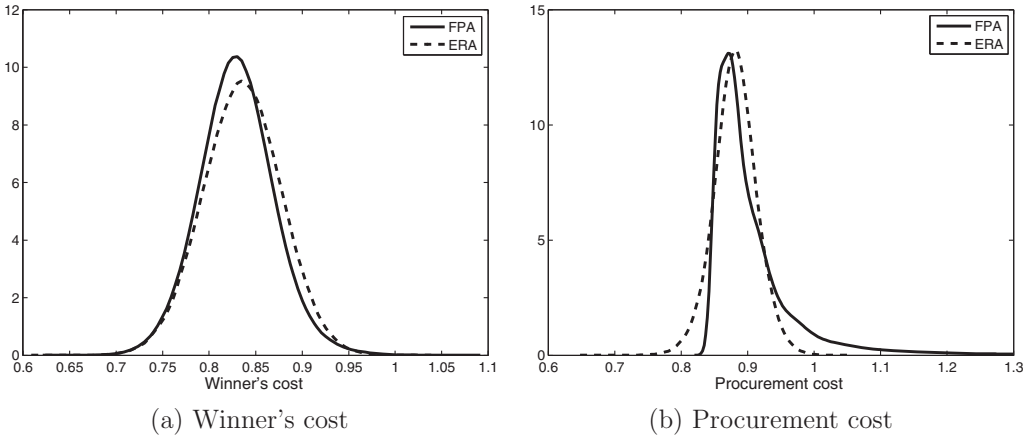
not hold with endogenous entry, but its logic is often used to motivate the importance of trying to generate additional interest in the assets being auctioned (e.g., Klemperer, 2002).

³⁴ For the baseline parameters, there is no entrant into the FPAFE in 1% of simulated auctions, in which case the procurement cost is 1.5. This would happen only if both of the selected firms had costs above 1.5 in the ERA and this happens with almost zero probability, and never in our simulations. If we condition our comparison on cases where at least one firm enters the FPAFE, so that we ignore simulations where the outcome of the FPAFE is particularly inefficient, the expected winner's cost is very similar for the ERA and the FPAFE, but the overall differences in social costs, procurement costs, and bidder profits are almost identical to those shown in Table 4.

³⁵ The alternatives are hypothetical in the sense that we ignore the fact that it might be difficult or very expensive for the procurer to get suppliers to report their signals or costs truthfully.

FIGURE 4

KERNEL DENSITIES FOR (A) THE WINNER'S COST OF COMPLETING THE CONTRACT AND (B) TOTAL PROCUREMENT COSTS IN BOTH THE FPAFE AND THE ERA, USING THE BASELINE PARAMETERS



Notes: The costs to the procurer in the ERA includes entry costs paid less the revenue from the first-stage bids. Densities are plotted based on 500,000 simulations. Simulations of the FPAFE where no suppliers entered are ignored.

Aggregating over the whole sample, total supplier profits fall by \$242,000 (2.82%) under the ERA. An exception is when entry costs are high (*High K*), as the bidders also benefit under the ERA, but even in this case the procurer benefits more than the bidders do.

To understand why procurement costs fall, recall that in an FPAFE, the procurer only pays the winning bid whereas in an ERA, it pays the winning bid and the entry costs of all of the suppliers admitted to the second stage, but also receives the first-stage bids of all of the suppliers. Table 5 shows that in general, first-stage revenues and the entry costs roughly offset from the procurer's perspective so that the fall in procurement costs comes primarily from the fact that winning bids tend to be significantly lower in the second stage of the ERA than in the FPAFE.

Winning bids are lower because the winner's markup tends to be lower in the ERA, not because the winner's cost is lower. Figure 4 plots the distribution of the winner's cost and the procurement cost for both mechanisms at the baseline parameters. The winner cost distributions are very similar in shape for the two mechanisms, but the distribution of procurement costs for the FPAFE has a more pronounced right tail, which is driven by the fact that winners with relatively high costs have high markups in the FPAFE.³⁶


The difference in markups for the baseline parameters can be seen more clearly in Figure 5, which compares the FPAFE bid function and the average of the bid functions from the second stage of the ERA with $n_{\text{cost}}^* = 2$.³⁷ The diagrams include the pdfs of the cost distribution of a typical entrant into each mechanism to indicate which parts of the bid curve are the most empirically relevant (Figure 4 shows the pdf for the winning bidder). The markups in the ERA are slightly larger than in the FPAFE for low-cost bidders, but they are smaller for higher costs, which the majority of bidders, and even a large share of winning bidders, have. This leads to an

³⁶ The left end of the procurement cost distributions in Figure 4(b) are also of note. In an FPAFE, there is a single bid function for given parameter values (θ_i) because entrants do not know how many other planholders are bidding. This bid function maps costs to the interval $[\underline{b}, r]$, and the sharp cutoff in the kernel density for the FPAFE corresponds to \underline{b} . On the other hand, bid functions in the ERA depend on the realization of \bar{s} , and this causes \underline{b} to vary depending on signals even when the parameters of the auction are fixed. As a result, there is no sharp cutoff in the density for the ERA.

³⁷ The bid function in the second stage of the ERA depends on the value of the highest losing bid in the first stage. Here, the average bid function in the ERA is defined by $\bar{\beta}(c) \equiv \int \beta(c; \bar{s}) f_{\bar{s}}(\bar{s}) d\bar{s}$, where $f_{\bar{s}}(\cdot)$ is the density of the $(n + 1)^{\text{st}}$ lowest signal.

FIGURE 5

COMPARING THE BID FUNCTION FOR AN FPAFE (HEAVY SOLID LINE) WITH THE AVERAGE BID FUNCTION IN AN ERA WITH $n = 2$ AT THE BASELINE PARAMETERS



Notes: The bid function for the FPAFE is shown with a heavy solid line, and the bid function for the ERA is shown with a heavy dashed line. The base parameters are $(\mu_C = -0.0963, \sigma_C = 0.0705, \alpha = 0.4979, K = 0.0147, N = 7, \text{ and a reserve price } r = 1.5)$. The dotted line is the 45° line, for comparison. The lighter lines show the densities for the cost of a typical entrant into each mechanism.

average markup for the winning bidder in the FPAFE of 9.4% (ignoring cases where there is no sale), whereas the average markup for the winning bidder in the second stage of the ERA is only 6.8%.

Markups are lower in the ERA even though for the baseline parameters a bidder should expect to face 2.9 rival bidders in the FPAFE and only one rival bidder in the second stage of the ERA with $n_{\text{cost}}^* = 2$. This reflects how the lack of coordination in entry affects bidding incentives. In the FPAFE, a bidder with a high cost knows that he is most likely to win when no other firms enter, in which case he should be submitting a bid at the reserve price; even though it is unlikely that no other firms enter, the possibility leads him to submit a high bid in equilibrium.³⁸ On the other hand, in an ERA with $n_{\text{cost}}^* = 2$, a bidder knows that he can only win if he submits a lower bid than his competitor. Therefore, as a bidder's cost increases, he will tend to bid more aggressively (i.e., with a lower markup), as can be seen in Figure 5 where the average ERA markup is monotonically decreasing in the bidder's cost for all costs below 1.2.³⁹

We can also compare the FPAFE and ERA, assuming that the number of selected firms in the ERA is chosen to maximize efficiency (n_{eff}^*). Up to integer constraints, n_{eff}^* balances the cost of an additional entrant against the possible gain that will come from lowering the cost of completing the project. Given our estimates, it turns out that using n_{eff}^* or n_{cost}^* does not change the results very much because, given the integer constraint, n_{eff}^* and n_{cost}^* are usually the same. This

³⁸ That high-cost bidders bid close to the reserve may suggest that the level of the reserve plays an important role in determining why average winning markups in the FPAFE are so large. However, Figure 4 shows that the distribution of the winner's cost is usually not in the region where bids are close to the reserve even if it is often in the range where markups are high. For the baseline parameters, reducing the reserve price to 1.1 only reduces the average winner's markup to 8.3% (compared to 9.4% when the reserve is 1.5), and increasing it to 4.75 (the maximum of the support of the cost distribution) only increases it to 10.9%.

³⁹ Above 1.2, the markup increases, reflecting the shape of the upper tail of the cost distribution. However, it is very rare for the winner in the second stage of an ERA to have a cost above 1.2, so this has little effect on procurement costs.

is the case for all of the parameters in Tables 4 and 5. When we draw parameters for all of the auctions in our data, there are some cases with low entry costs where the optimal n s are higher (e.g., in 12% of cases n_{cost}^* is greater than 2 and it takes on values as high as 6). In 17% of cases, we find that $n_{\text{eff}}^* > n_{\text{cost}}^*$. However, although this is a significant proportion of cases, we find that using n_{eff}^* rather than n_{cost}^* only reduces social costs slightly, by less than \$0.1 million in aggregate (recall that the advantage of the ERA with n_{cost}^* over the FPAFE is \$7.8 million). In none of our experiments have we found a case where $n_{\text{eff}}^* < n_{\text{cost}}^*$. The intuition here is that allowing an additional entrant reduces the bids that firms make in the first stage of the ERA, and a procurer that seeks to minimize procurement costs, therefore, has an incentive to exert market power by reducing the number of entry slots. This incentive appears to typically outweigh the incentive to reduce second-stage markups by increasing competition. This is sensible given that, even with $n = 2$, markups in the second stage of the ERA tend to be small.

Given the sizable gains of using an ERA that we estimate above, a potential puzzle is why they are not used more often for the sort of public procurement we study here. There are several possible explanations. First, it may be the case that the reduced supplier profits are actually harmful to the procurer in the long run, especially when it wants to let many contracts over a number of years. In this case, it may benefit when there are a large number of firms in the industry who are potentially interested in bidding, although, as noted above, the benefits from increasing the number of potential suppliers in an FPAFE are quite small in our setting. Kagel, Pevnitskaya, and Ye (2008) show that in a laboratory setting students tend to overbid in ERAs, leading to bankruptcies. If this is true in the field as well, this might provide an additional route by which ERAs might diminish long-run competition.

Second, in empirical settings like ours, many suppliers are relatively small companies that face liquidity constraints, and these may be tightened by having to make binding bids into ERAs.⁴⁰ A system that favors bigger companies may be undesirable, as one of the goals of public procurement is often to favor smaller or minority-owned businesses (e.g., Krasnokutskaya and Seim, 2011; Athey, Coey, and Levin, 2013). On the other hand, to the extent that suppliers are also hurt by uncoordinated entry when entry is costly, one could argue that these same concerns could also provide motivation for regulating entry.

Third, designing an ERA appropriately, in particular choosing n , may require the procurer to have more information, for example, about the distribution of costs or the entry cost, than running a FPAFE. In our empirical setting, when we take a draw of the parameters for each of the auctions in our sample, $n = 2$ is the procurer-optimal choice for 88% of these draws, and it is always the best choice unless K is small.⁴¹ Therefore, in general, adopting a policy that two suppliers will be selected should be effective.

Fourth, Quint and Hendricks (2013) suggest that ERA may provide sellers with perverse incentives to try to sell off contracts that, once they undertake research, bidders will find to be worthless. In some settings this may be an important consideration, but it may be less of a concern when a state agency is letting a large number of contracts and so should want to maintain a reputation for integrity.

Finally, it may of course be that some of the standard assumptions used to model procurement auctions, which we have maintained here, such as independent private values, simultaneous entry decisions, and independent signals (e.g., Athey, Levin, and Seira, 2011; Krasnokutskaya and Seim, 2011; Li and Zheng, 2011) are incorrect and affect the comparison of the different designs. For example, if bidder signals are correlated or entry decisions are made sequentially, entry into the FPAFE may be less volatile, mitigating some of the gains to using the ERA. Alternatively,

⁴⁰ Following the same logic, liquidity constraints or risk aversion might tend to reduce first-stage bids, increasing procurement costs.

⁴¹ Exactly how small K needs to be to make having more entrants optimal depends on the other parameters, but particularly the level of α . As α increases, it becomes more desirable to select more entrants.

it may be that our assumption that firms act competitively in both mechanisms is not correct: if firms are more able to collude in an ERA, then this might offset the reductions in procurement costs that our model predicts.

6. Conclusion

■ In procurement settings, where it is expensive for potential suppliers to learn their costs of completing a project, unregulated entry can lead to volatile amounts of supplier participation. If suppliers have either very little information or almost perfect information about their costs prior to entering, then fixing the number of entrants, and choosing them through some type of entry rights auction procedure, will tend to raise efficiency. However, when suppliers are partially informed about their costs, fixing the number of entrants in advance can be inefficient because it ignores the fact that suppliers' private information can determine how valuable additional entry would be.

We compare the performance of a standard first-price auction with free entry and an entry rights mechanism in the empirical setting of highway construction projects involving bridge work in Oklahoma and Texas, where we find suppliers' information to be partially informative and entry costs to be significant. For the vast majority of projects in our sample we predict that using an entry rights auction procedure would significantly increase efficiency and reduce procurement costs, and that using ERAs would have much larger effects on procurement costs than alternative procedures such as setting an optimal reserve price in a standard auction or even increasing the number of potential bidders. This reflects the fact that an entry rights procedure can both economize on entry costs by coordinating entry, but also encourage even high-cost suppliers to bid aggressively by ensuring that they will face competition.

Appendix A

In this Appendix, we provide some additional details on how we solve the FPAFE model developed in Section 2.

The choice of grid $\{x_i\}_{i=1}^N$ has a minor effect on the solution to the programming problem. Hubbard and Paarsch (2009) use an N -point Gauss-Lobatto grid on $[\underline{b}, r]$,⁴² but we have found in our experiments that an evenly spaced grid is often more efficient. Throughout the article, we use $N = 500$ grid points and $P = 25$ polynomials in the expansion; these choices solve the optimization problem reliably and efficiently.

Bid functions are monotonic in the cost of the bidder, so we follow Hubbard and Paarsch (2009) and impose that $\beta^{*i}(x_i) \geq \beta^{*i}(x_{i-1})$ for $2 \leq i \leq N$ when solving the problem in (4). Furthermore, we impose the rationality constraint that an agent never bids less than his cost: $\beta^{*i}(x_i) \leq x_i$ for all i . Finally, we replace the constraint in equation (3) by the condition that the two sides cannot differ by more than 0.1% of K . The integral in equation (3) is replaced by an approximation using the trapezoidal rule on the N -point grid defined previously.

Although we need to calculate bid functions to use in our estimation procedure, we are able to verify the accuracy of our solutions by exploiting the fact that a second-price auction, which is easier to solve (Roberts and Sweeting, 2013b, use this format to analyze timber auction data), should give the same equilibrium entry thresholds and the same expected procurement cost as a first-price auction with the same parameters. In the many examples that we have looked at, both the entry thresholds and the expected revenues typically agree to several decimal places.

Appendix B

In this Appendix we provide details on how we select the sample of procurement auctions that we use in estimation. De Silva, Dunne, Kankaname, and Kosmopoulou (2008) study a range of different types of highway procurement auctions from January 1998 through August 2003. However, we focus only on auctions after March 2000, due to an important change in policy in Oklahoma: prior to March 2000, they did not disclose the engineer's estimate to bidders. Throughout this article, we restrict our attention to auctions for new projects (not reauctions of failed sales) involving bridge work with between four and 11 planholders. From the remaining set of bridge construction contracts let during this period (483 in OK and 177 in TX), we also employ some additional selection rules removing auctions where (i) there was no

⁴² A Gauss-Lobatto grid on $[-1, 1]$ is defined to be the points $y_k = \cos[k\pi/(N-1)]$ for $k \in \{0, 1, \dots, N-1\}$. To define a grid on other intervals, we simply scale these points linearly.

winner or the DoT rejected all bids (drops 50 in OK, one in TX); (ii) the engineer's estimate was less than \$100,000 or more than \$5 million (drops 121 in OK, 15 in TX); (iii) the data is incomplete in the sense that it records that the number of recorded planholders differs from the number of suppliers that are listed (drops one in OK); and (iv) the winning bid was extremely high or extremely low. Specifically, we drop auctions where the winning bid was less than 70% of the engineer's estimate, which drops three auctions in Texas and 47 auctions in Oklahoma. Although these auctions have similar observable characteristics to the rest of the sample, 16 of these Oklahoma auctions were won by two companies that only won four of the remaining auctions in the sample, suggesting that they may have been somewhat unusual. We also drop two auctions in Oklahoma and four in Texas with winning bids greater than 150% of the engineer's estimate, as we assume in estimation that the DoTs set a reserve price at this level even though no reserve price was officially announced for these auctions. After these selection rules, we have a sample of 262 auctions in Oklahoma and 154 in Texas.

Appendix C

In this Appendix, we describe some Monte Carlo experiments designed to test the performance of a simulated method of moments estimator where we use importance sampling to approximate the moments. We are particularly interested in how many importance sampling draws per auction are required for accurate estimates (in our actual estimation, we use 50). In all of our experiments we create data using the specification in (5). We draw the number of potential suppliers N uniformly at random from the set $\{4, 6, 8\}$. We set the vector of observed characteristics for auction a to be $X_a \equiv (1, x_a)$ where x_a is a random draw from the set $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$. The parameters in the specification are assumed to be $\beta_{\mu_C} = (0, 0.65)'$, $\beta_{\sigma_C} = (0.05, 0)'$, $\beta_\alpha = (0.5, 0)'$, $\beta_K = (0.04, 0.05)'$, $\omega_{\mu_C} = 0.05$, $\omega_{\sigma_C} = 0.02$, $\omega_\alpha = 0.05$, and $\omega_K = 0.015$. We denote this "true" parameter vector as Γ_0 . The truncation points are taken to be $\underline{c}_{\mu_C} = -0.4$, $\bar{c}_{\mu_C} = 1$, $\underline{c}_{\sigma_C} = 0.005$, $\bar{c}_{\sigma_C} = 0.995$, $\underline{c}_\alpha = 0.1$, $\bar{c}_\alpha = 0.9$, $\underline{c}_K = 0.0001$, and $\bar{c}_K = 0.16$.⁴³ We truncate the cost distribution to $[0, 4.75]$, where the probability that a cost draw will be above the upper support is almost always negligible for our draws of the parameters, and we set the reserve price equal to 4.75. For any set of parameters drawn from this distribution, we can solve for the equilibrium bid functions and entry decisions, and simulate outcomes.

As in our empirical application, we estimate the parameters by matching a set of moments describing the outcomes from auctions with their expectations, calculated using importance sampling, with separate moments for each $\{N, x_a\}$ combination. Specifically, we match the distribution of the number of entrants (i.e., indicator variables for whether one, two, three, etc. firms enter), the distribution of winning bids (using indicator variables for whether the winning bid is in a particular bin, where we divide the interval $[0.4, 2.65]$ into bins of width 0.075), the distribution of all bids (we divide the interval $[0.4, 4.2]$ into bins of width 0.1) and the number of suppliers that enter but do not submit bids because they have valuations greater than the reserve. We estimate the parameters by minimizing the squared sum of the differences between the observed and predicted moments, where each moment is weighted equally.

We report the results from three experiments. In each experiment, we use 500 auctions as data, but we vary the importance sampling density and the number of simulated draws we use for each auction. Table C1 shows the mean and the standard deviation of each parameter when we run each experiment 100 times.

In the first experiment (A), we use the true distribution of the parameters as the importance sampling density, which is optimal but obviously infeasible in an empirical application where the goal is to estimate the true parameters, and use five simulations for each auction.⁴⁴ In the second experiment (B), we make the importance sampling density more diffuse, by setting $\omega_{\mu_C} = 0.1$, $\omega_{\sigma_C} = 0.04$, $\omega_\alpha = 0.1$, and $\omega_K = 0.03$, and other parameters remain the same. We continue to use five simulations per auction.⁴⁵ In the final experiment (C), we use the same importance sampling density as in (B) but we use 30 draws per auction. In all of our experiments, we use the true parameters to begin our search but we have verified that using a wide-range of starting values does not change our estimates because the objective function is generally very well behaved. The results of these experiments are shown in columns (A)–(C) of Table C1, respectively.

In all three experiments we recover the coefficients accurately, and, when we use a more diffuse importance sampling density and 30 draws (C), we recover the parameters almost as well as when we use the true distribution of the parameters (A). There is some evidence that we tend to overestimate the variance ω parameters in (B) when we use the diffuse density and only five draws, but even in this case, the magnitude of any bias appears to be quite small.

Based on these results, in our empirical application we choose to use 50 draws for each auction. We also choose the importance sampling density based on a large number of initial runs, using quite diffuse importance sampling densities. This indicated that the parameters of the importance sampling density that we actually use (see Table 2) were likely to be close to the parameters that we would end up estimating.

⁴³ These truncation points are chosen to ensure that we can solve the FPAFE model accurately.

⁴⁴ Note rather than solving a fresh set of 2,500 auctions each time we perform a separate run of the experiment, we solve 25,000 auctions drawn from the distribution up front, and then draw, with replacement, from this pool when creating our data and importance sampling draws for each run.

⁴⁵ In this case, we created solved 75,000 auctions and draw from this pool for each run.

TABLE C1 Results of Monte Carlo Experiments

Parameter	Variable	True Value	Experiment		
			A	B	C
Location parameter (μ_C)	Constant	0	0.0000 (0.0100)	0.0015 (0.0186)	-0.0043 (0.0170)
	x_a	0.65	0.6507 (0.0174)	0.6533 (0.0282)	0.6602 (0.0247)
	ω	0.05	0.0488 (0.0044)	0.0575 (0.0073)	0.0504 (0.0059)
Scale parameter (σ_C)	Constant	0.05	0.0494 (0.0074)	0.0463 (0.0152)	0.0416 (0.0126)
	x_a	0	-0.0002 (0.0113)	0.0085 (0.0204)	0.0195 (0.0194)
	ω	0.02	0.0192 (0.0035)	0.0295 (0.0094)	0.0225 (0.0077)
Degree of selection (α)	Constant	0.5	0.4994 (0.0205)	0.4985 (0.0354)	0.4984 (0.0231)
	x_a	0	0.0000 (0.0308)	-0.0102 (0.0607)	-0.0130 (0.0571)
	ω	0.05	0.0479 (0.0078)	0.0803 (0.0190)	0.0561 (0.0164)
Entry cost (K)	Constant	0.04	0.0388 (0.0047)	0.0361 (0.0118)	0.0387 (0.0085)
	x_a	0.05	0.0516 (0.0074)	0.0559 (0.0173)	0.0529 (0.0121)
	ω	0.015	0.0143 (0.0025)	0.0209 (0.0063)	0.0145 (0.0038)

Notes: Each column reports the mean and standard deviation of the parameter estimates from 100 runs, where there are 500 auctions used as data in each run. In column A, the importance sampling density is the true distribution and we use 5 simulations per auction in the data. In columns B and C, we use a more diffuse importance sampling density (see text) and either 5 (B) or 30 (C) simulations per auction.

References

- ACKERBERG, D. "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation." *Quantitative Marketing and Economics*, Vol. 7 (2009), pp. 343–376.
- ATHEY, S., COEY, D., AND LEVIN, J. "Set-Asides and Subsidies in Auctions." *American Economic Journal: Microeconomics*, Vol. 5 (2013), pp. 1–27.
- ATHEY, S. AND HAILE, P. "Identification of Standard Auction Models." *Econometrica*, Vol. 70 (2002), pp. 2107–2140.
- ATHEY, S., LEVIN, J., AND SEIRA, E. "Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions." *Quarterly Journal of Economics*, Vol. 126 (2011), pp. 207–257.
- BAJARI, P., HONG, H., AND RYAN, S. "Identification and Estimation of a Discrete Game of Complete Information." *Econometrica*, Vol. 78 (2010), pp. 1529–1568.
- BERRY, S. AND TAMER, E. "Identification in Models in Oligopoly Entry." In R. Blundell, W. Newey, and T. Persson, eds., *Advances in Economics and Econometrics, Ninth World Congress, Vol. 2*. New York: Cambridge University Press, 2006.
- BULOW, J. AND KLEMPERER, P. "Auctions versus Negotiations." *American Economic Review*, Vol. 86 (1996), pp. 180–194.
- CONGRESSIONAL BUDGET OFFICE. "Economic and Budget Issues Brief: Spending and Funding for Highways." Washington, DC, 2011.
- CREMER, J., SPIEGEL, Y., AND ZHENG, C. "Symposium on Auction Theory: Auctions with Costly Information Acquisition." *Economic Theory*, Vol. 38 (2009), pp. 41–72.
- DE SILVA, D., DUNNE, T., KANKANAMGE, A., AND KOSMOPOULOU, G. "The Impact of Public Information on Bidding in Highway Procurement Auctions." *European Economic Review*, Vol. 52 (2008), pp. 150–181.
- FOLEY, J.A. "Multi-Round Auctions for Institutional Real Estate Assets: Theory and Practice," Master's Thesis, MIT, 2003.
- FULLERTON, R.L. AND MCAFEE, R.P. "Auctioning Entry into Tournaments." *Journal of Political Economy*, Vol. 107 (1999), pp. 573–605.

- GENTRY, M. AND LI, T. "Auctions with Selective Entry." Working Paper, Department of Economics, London School of Economics, 2012.
- . "Identification in Auctions with Selective Entry." *Econometrica*, Vol. 82 (2014), pp. 315–344.
- GENTZKOW, M. AND SHAPIRO, J. "Measuring the Sensitivity of Parameter Estimates to Sample Statistics." Working Paper, Booth School of Business, University of Chicago, 2013.
- HALPIN, D. *Construction Management*, 3d ed. New York: Wiley, 2005.
- HUBBARD, T. AND PAARSCH, H. "Investigating Bid Preferences at Low-Price, Sealed-Bid Auctions with Endogenous Participation." *International Journal of Industrial Organization*, Vol. 27 (2009), pp. 1–14.
- JOFRE-BONET, M. AND PESENDORFER, M. "Estimation of a Dynamic Auction Game." *Econometrica*, Vol. 71 (2003), pp. 1443–1489.
- KAGEL, J., PEVNIYSKAYA, S., AND YE, L. "Indicative Bidding: An Experimental Analysis." *Games and Economic Behavior*, Vol. 62 (2008), pp. 697–721.
- KLEMPERER, P. "What Really Matters in Auction Design." *Journal of Economic Perspectives*, Vol. 16 (2002), pp. 169–189.
- KRASNOKUTSKAYA, E. "Identification and Estimation in Highway Procurement Auctions under Unobserved Auction Heterogeneity." *Review of Economic Studies*, Vol. 78 (2011), pp. 293–327.
- KRASNOKUTSKAYA, E. AND SEIM, K. "Bid Preference Programs and Participation in Highway Procurement Auctions." *American Economic Review*, Vol. 101 (2011), pp. 2653–2686.
- LEVIN, D. AND SMITH, J. "Equilibrium in Auctions with Entry." *American Economic Review*, Vol. 84 (1994), pp. 585–599.
- LI, T. "Indirect Inference in Structural Econometric Models." *Journal of Econometrics*, Vol. 157 (2010), pp. 120–128.
- LI, T. AND ZHANG, B. "Affiliation and Entry in First-Price Auctions with Heterogeneous Bidders: An Analysis of Merger Effects." Working Paper, Department of Economics, Vanderbilt University, 2013.
- LI, T. AND ZHENG, X. "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions." *Review of Economic Studies*, Vol. 76 (2009), pp. 1397–1429.
- . "Information Acquisition and/or Bid Preparation: A Structural Analysis of Entry and Bidding in Timber Sale Auctions." *Journal of Econometrics*, Vol. 168 (2011), pp. 29–46.
- LU, J. AND YE, L. "Efficient and Optimal Mechanisms with Private Information Acquisition Costs." *Journal of Economic Theory*, Vol. 148 (2013), pp. 393–408.
- MANKIW, N. AND WHINSTON, M. "Free Entry and Social Inefficiency." *RAND Journal of Economics*, Vol. 17 (1986), pp. 48–58.
- MARMER, V., SHNEYEROV, A., AND XU, P. "What Model for Entry in First-Price Auctions? A Nonparametric Approach." Working Paper, Department of Economics, University of British Columbia, 2011.
- . "What Model for Entry in First-Price Auctions? A Nonparametric Approach." *Journal of Econometrics*, Vol. 176 (2013), pp. 46–58.
- MILGROM, P. *Putting Auction Theory to Work*. Cambridge, England: Cambridge University Press, 2004.
- OECD. "Size of Public Procurement Market." In *Government at a Glance 2011*. Paris, France: OECD Publishing, 2011. <http://www.oecd.org/gov/governmentataglance2011> Accessed August 13, 2014.
- PARK, W. AND CHAPIN, W.B. *Construction Bidding: Strategic Pricing for Profit (2 ed)*. New York: Wiley, 1992.
- QUINT, D. AND HENDRICKS, K. "Indicative Bidding in Auctions with Costly Entry." Working Paper, Department of Economics, University of Wisconsin, 2013.
- ROBERTS, J. AND SWEETING, A. "Bailouts and the Preservation of Competition." NBER Working Paper no. 16650, 2013a.
- . "When Should Sellers Use Auctions?" *American Economic Review*, Vol. 103 (2013b), pp. 1830–1861.
- SAMUELSON, W. "Competitive Bidding with Entry Costs." *Economic Letters*, Vol. 17 (1985), pp. 53–57.
- SU, C. AND JUDD, K. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica*, Vol. 80 (2012), pp. 2213–2230.
- WELCH, D. AND FREMOND, O. *The Case-by-Case Approach to Privatization: Techniques and Examples*. Washington, DC: World Bank, 1998.
- YE, L. "Indicative Bidding and a Theory of Two-Stage Auctions." *Games and Economic Behavior*, Vol. 58 (2007), pp. 181–207.