Optimal Procurement With Quality Concerns*

Giuseppe Lopomo
Duke†

Nicola Persico
Northwestern & NBER‡

Alessandro T. Villa
FRBC§

September 28, 2021

Abstract

Adverse selection in procurement arises when low-cost bidders are also low-quality suppliers. We propose a mechanism called LoLA which, under some conditions, is the best incentive-compatible mechanism for maximizing either the seller's or the social surplus in the presence of adverse selection. The LoLA features a floor (or minimum) price, and a reserve (or maximum) price. Conveniently, the LoLA has a dominant strategy equilibrium that, under mild regularity conditions, is unique. We perform a counterfactual experiment on Italian government procurement auctions: we compute the gain that the government could have made, had it used the optimal mechanism (which happens to be a LoLA), relative to a first-price auction, which is the format the government actually used. Finally, we provide software applications for computing the optimal procurement mechanism.

Keywords: Auctions, Procurement, Mechanism Design, Adverse Selection

JEL Codes: D44, H57

*We thank Jim Anton, and seminar participants at the EIEF and Duke University. Excellent research assistance by Chenxi Xu is gratefully acknowledged. Special thanks to Francesco Decarolis for generously sharing the data on which Section 5 is based, and to Giancarlo Spagnolo for helpful feedback.

†The Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708, USA. Email: `glopomo@duke.edu`

‡Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL 60208, USA. Email: `nicola@nicolapersico.com`

§Federal Reserve Bank of Chicago 230 S. LaSalle St. Chicago, IL 60604. Email: `alessandro.villa@chi.frb.org`

1 Introduction

When the quality of a good or service is non-contractible, a buyer holding a standard procurement auction faces an adverse selection (or "lemons") problem: the sellers who bid aggressively may be the low-quality ones. This problem is pervasive in procurement settings: cheap suppliers may provide low quality (maybe because they use shoddy materials and less-qualified labor), whereas high-quality contractors may have high costs and thus be unwilling to bid aggressively. In this case, we say that the buyer has *quality concerns*.

To deal with the adverse selection problem, many procurement agencies use versions of the "average bid auction" (ABA). The ABA format disqualifies "extreme" bids, i.e., bids that fall in extremely high or low quantiles of the bid distribution. The rationale for disqualifying low bids is to weed out low-quality bidders.[1] However, the ABA format has severe limitations. Theoretically, the ABA is unsatisfactory because it gives rise to multiple coordination equilibria where all the bidders coordinate on the same bid (any deviations being disqualified as "extreme"), which is bad for the buyer if the bid is high. Practically, the ABA has been shown to foster collusion. Both concerns were documented by Albano et al. (2016), Decarolis (2014, 2018), and Conley and Decarolis (2016).

This paper derives the optimal mechanism for buying a good or service when there is an adverse selection problem. We call it a "*lowball lottery auction*" (LoLA). A LoLA is a sealed-bid auction that features a "floor price" p_L, such that:

1. If at most one bidder bids below the "floor price" p_L, the lowest bidder supplies the good and is paid the second-lowest bid.
2. If two or more bidders bid below p_L, one of these bidders is randomly selected to supply the good and is paid p_L.

In a LoLA, the buyer commits to pay no less than the floor price p_L. From a bidder's perspective, price competition is less intense if the floor price is higher. When p_L is set at a sufficiently high level, price competition is completely eliminated, and the winning bidder is selected randomly. At the other extreme, when p_L is set below the lowest possible cost, the LoLA becomes a standard second-price auction.

We show that, under mild regularity assumptions, the buyer's expected surplus is maximized by a LoLA among all interim IC and IR mechanisms. To our knowledge, this

[1] See Decarolis and Klein (2011), p. 2.

is the first time that a floor price emerges as part of an optimal selling mechanism. In real-world procurement settings, we are aware of only one case where minimum prices have been used.[2]

Intuitively, a floor price is most helpful when the buyer's quality concerns come from the lower-cost suppliers: in this case, the floor price can make it less likely that most-aggressive bidders – who, presumably, are also the lowest-cost ones – win the auction. Setting the buyer-optimal floor price p_L entails a trade-off: lowering p_L saves the buyer some money, but it increases the quality concerns associated with selecting a cheaper supplier. We will show that if the quality concerns are more severe, in a sense that will be made formal later, then the optimal floor price p_L^* is higher. If the auction designer maximizes social welfare rather than buyer revenues, then the optimal mechanism remains a LoLA but, under fairly general conditions, one with a higher optimal p_L^*. This is intuitive because a benevolent designer does not internalize the buyer's monetary savings from lowering p_L.

The buyer may also choose to augment the LoLA with a "reserve price" that excludes any bid above a certain threshold. A LoLA with a reserve price is reminiscent of the ABA in that both high and low bids are curbed. But in a LoLA the reserve and floor prices are exogenous, whereas in an ABA the disqualification thresholds are a function of the bid distribution. And, whereas the ABA has a continuum of symmetric pure-strategy equilibria, none of which is in (even weakly) dominant strategies (see Decarolis 2014), under mild conditions, the LoLA has a unique equilibrium, and this equilibrium is in weakly dominant strategies.

To illustrate the gains from the optimal mechanism we perform a counterfactual experiment on Italian government procurement auctions. Using information generously provided by Francesco Decarolis, and making some assumptions about how quality enters the government's objective function, we compute the gain that the government could have made, had it used the optimal mechanism (which happens to be a LoLA), relative to a first-price auction, which is the format the government actually used. We find that, in a reasonably calibrated model, these savings can be nontrivial.

Finally, we created two software applications and made them publicly available. These applications compute the buyer-optimal procurement mechanisms in the presence of quality concerns, whether or not the optimal mechanism is a LoLA.

[2]This is the case of some Japanese procurement auctions, see Chassang and Ortner (2019).

The theoretical literature on optimal procurement in the presence of quality concerns is sparse. When there is no lemons problem, the first-price auction and the second-price auction are both socially optimal and maximize the buyer's surplus (Myerson 1982). When the lemons problem is sufficiently severe, Manelli and Vincent (1995) show that it is optimal to select the wining bidder randomly. Both results obtain as polar cases in our setting because, indeed, both mechanisms are LoLAs for suitably chosen values of p_L. Manelli and Vincent (2004) study several functional-form examples with two players, in which certain sequential mechanisms maximize the social surplus in a "lemons" environment. Our implementation, in contrast, is through a sealed-bid auction. Of course, if the functional form in one of their examples satisfies our assumptions, their optimal mechanism and ours must yield the same allocation and payoffs.[3]

The formal literature on (non-optimal) procurement in the presence of quality concerns goes back to, at least, Dini et al. (2006) and Albano et al. (2006). The latter have shown that a mechanism in the spirit of the ABA admits a continuum of equilibria in which the bidders coordinate to keep prices high. Decarolis (2014) documented empirically the severity of the lemons problem in first-price auctions compared to ABAs. The drawbacks of the ABA format, i.e., multiple coordination equilibria and vulnerability to collusion, are documented empirically by Conley and Decarolis (2016). Decarolis (2018) compares the performance of ABA and first price auctions. When contracts are allocated using the ABA, Decarolis (2018) shows that bidders bid extremely close to each other, which can be interpreted as evidence of an "approximately random" allocation. The winner's quality seems to be better when winners are chosen "randomly," suggesting that these auctions suffer from adverse selection.[4] In a dynamic model of bidder collusion, Chassang and Ortner (2019) document theoretically and empirically that, counterintuitively, introducing minimum prices can lower the winning-bid distribution.[5]

In sum, our first and main contribution relative to the literature is that we characterize the *optimal* procurement mechanism in the presence of adverse selection. The optimal mechanism was not known before, except in the extreme case where the adverse selection was so severe that random assignment was optimal. Our proposed mechanism is similar

[3]This is the case for the functional form studied in their Theorem 2. It should be noted that Manelli and Vincent's (2004) analysis is not a special case of ours because some of their examples do not satisfy our assumptions.

[4]Specifically, Decarolis (2018) shows that delays and cost overruns tends to be lower in the ABA than in a first price auction (where contracts are allocated to the lowest bidder).

[5]Calzolari and Spagnolo (2006) also study repeated procurement in the presence of quality concerns.

enough to the ABA format that, we think, it could be perceived as "natural" by practitioners and, thus, implemented in practice. A second contribution is the calibration exercise with Italian procurement data: we show that the LoLA is in fact the optimal mechanism in that setting, and quantify the gain over the existing procurement protocol. We view the calibration method as the main contribution of this exercise, because the method has external validity beyond the specific setting of Italian auctions. A third contribution is a pair of software applications that we have created and made available for the computation of the optimal mechanism (which may or may not be a LoLA). We believe that this contribution is mainly pedagogical; we hope that, by making it easy to compute the optimal auction, we can promote the adoption of the LoLA in real-world settings.

The paper proceeds as follows. The next section contains a simple illustrative example. Section 3 lays out the model. Section 4 derives the optimal mechanism and some comparative static results. Section 5 analyzes the Italian procurement auctions. Section 6 describes the software we created. Section 7 concludes.

2 An illustrative example

This section provides a functional form example to build intuition for the general results to follow.

A buyer faces two suppliers. Each supplier's production cost c_i is privately known and is an i.i.d. random variable distributed uniformly on $[0, 1]$. The buyer's willingness to pay for supplier i's product is given by:

$$v(c_i) \equiv 4\, c_i - 2\, c_i^2 \tag{1}$$

The function $v(\cdot)$ is increasing and concave on $[0, 1]$, which means that the buyer's use value increases with production cost, albeit at a decreasing rate. The increasingness captures the lemons problem: more-reliable suppliers have higher costs. The concavity means, intuitively, that the lemons problem is more severe where the function $v(\cdot)$ increases more steeply, i.e., at lower values of c.

A LoLA coincides with a second-price auction except when both bidders bid less than p_L, in which case either wins with equal probability and pays p_L. In a LoLA, it is a dominant strategy to bid one's cost; this will be proved in Theorem 1. Figure 1 shows the outcome of the LoLA with a floor price $p_L \in (0,1)$, for any realization of the suppliers' costs.

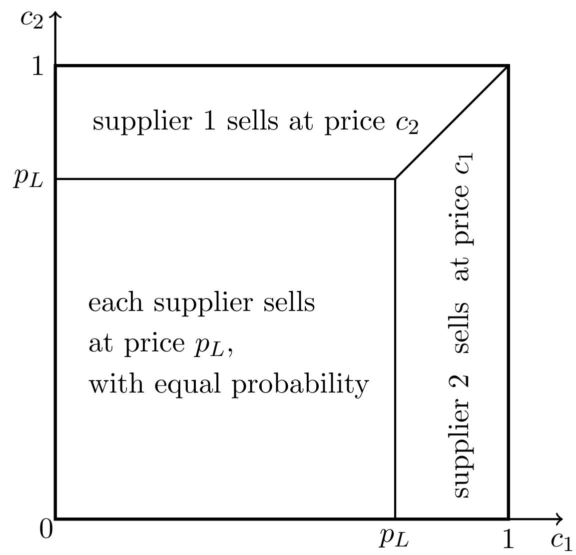

Figure 1: Outcome of the LoLA with floor price p_L

Note that setting $p_L = 0$ yields the first price auction, and $p_L = 1$ yields the random assignment mechanism. In the inner-square region, there is no competition between bidders. This happens to be the region where, intuitively, the lemons problem is worse, because the function $v(\cdot)$ is steeper. Thus, in the LoLA, the buyer gives up the monetary benefits of competition precisely in the region where the lemons problem is most severe, but not in other regions.

The expected buyer surplus generated by a LoLA with threshold price p_L is:

$$\begin{aligned} V(p_L) &= 2\int_{p_L}^{1}\left(\int_{0}^{c_2}[v(c_1)-c_2]\,dc_1\right)dc_2 + \int_0^{p_L}\int_0^{p_L}\left[\frac{1}{2}v(c_1)+\frac{1}{2}v(c_2)-p_L\right]dc_1dc_2. \\ &= \frac{1}{3}+\frac{1}{3}\cdot(p_L)^3\cdot(1-p_L). \end{aligned} \tag{2}$$

The first double integral covers the right-trapezoid region in which bidder 2 bids more than her opponent and above the "floor price" p_L. In this case, the LoLA prescribes that the lowest bidder supplies the good and is paid the second-lowest bid c_2. This term is doubled to account for the specular case in which bidder 1 bids the most. The second

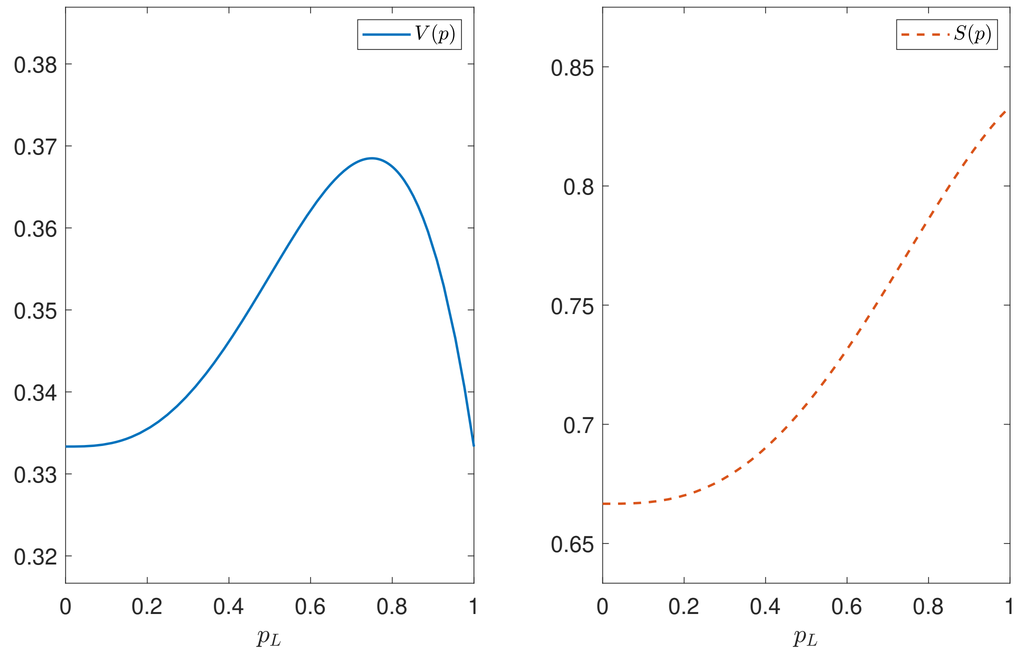

Figure 2: Expected buyer surplus V and social surplus S. V is maximal at ${p_L}^* = 3/4$ under a LoLA with floor price p_L.

double integral covers the inner-square region where two or more bidders bid below p_L. In this case, the LoLA prescribes that one of these bidders is randomly selected to supply the good and is paid p_L. The last equality follows from substituting for $v(\cdot)$ from (1) and solving the integrals.

The expected social surplus generated by a LoLA with threshold price p_L is:

$$\begin{aligned} S(p_L) &= 2\int_{p_L}^{1}\int_{0}^{c_2} [v(c_1) - c_1]\, dc_1 dc_2 + \int_{0}^{p_L}\int_{0}^{p_L} \left[\frac{1}{2}(v(c_1) - c_1) + \frac{1}{2}(v(c_2) - c_2)\right] dc_1 dc_2 \\ &= \frac{2}{3} + \frac{1}{3}\cdot\left(\frac{3}{2} - p_L\right)\cdot (p_L)^3. \end{aligned} \tag{3}$$

Figure 2 graphs the expected buyer surplus V and expected social surplus S as a function of p_L. The function V attains a maximum of about 0.37. By comparison, the first price auction and the random assignment mechanism, which correspond to LoLAs with $p_L = 0$ and $p_L = 1$, respectively, achieve a buyer's surplus of roughly 0.33 each.

Therefore, in this example the buyer-optimal LoLA is seen to improve the buyer's surplus by more than 10% relative to either the first price auction or the random assignment mechanism. The fact that the buyer-optimal p_L is interior indicates that the lemons problem is severe enough that the first price auction is not optimal; but not so severe that random allocation is optimal (i.e., Manelli and Vincent 1995 does not apply here).

By contrast, the expected social surplus $S(\cdot)$ is monotonically increasing in p_L, which implies that the socially optimal LoLA has $p_L = 1$. Therefore, in this example, random allocation is socially optimal but not buyer-optimal. That $V(\cdot)$ peaks earlier than $S(\cdot)$ is a general property: the buyer prefers a lower p_L than the social planner (see Proposition 3). This is intuitive: a small decrease from p_L to $p_L - \epsilon$ has first-order benefits for the buyer (the decrease in payment is realized whenever the winning bidder's type is smaller than p_L) but only second-order benefits for the social planner (the allocation improves only when the winning bidder's type is between p_L and $p_L - \epsilon$).

3 Model

A buyer with known type ξ seeks to procure an indivisible good from one of N potential suppliers. The suppliers' costs $c_1, ..., c_N$ are elements of the interval $[c_L, c_H]$. These costs are privately known, and they are independently drawn from the same distribution with density f. If a supplier with cost c is selected and paid m, the supplier's profit is

$$m - c,$$

and the buyer's surplus is

$$v(c, \xi) - m.$$

The function v represents the buyer's value from procuring the good from a buyer with cost c. If v is independent of c, we have the classic setting of Myerson (1981). If v is increasing in c there are quality concerns. The scalar ξ parameterizes the severity of the buyer's quality concerns: we assume that $v_{c\xi}(c, \xi) \geq 0$, meaning that when ξ is larger, intuitively, the quality concerns are more severe. For analytical convenience, we also assume $v(c_L, \xi) \geq c_L$, meaning that there are gains from trade at the lowest supplier type. This assumption does not imply that there are gains from trade for all types.

The virtual valuation function is defined as:

$$w(c;\xi,\beta) \equiv v(c;\xi) - c - \beta\frac{F(c)}{f(c)}. \tag{4}$$

The ratio $\frac{F(c)}{f(c)}$ represents the information rent earned by a supplier with type c. As we will show later, the scaling parameter $\beta \in [0,1]$ encodes the designer's concern for the buyer's share of the social surplus. When $\beta = 1$ the designer is solely focused on maximizing the buyer's surplus, as in Myerson (1981). When $\beta = 0$ the designer focuses entirely on social surplus. Interior values of β capture intermediate degrees of concern for buyer vs. social surplus. When $\beta = 0$ (resp., 1) we refer to (4) as the buyer's virtual valuation (resp., gains from trade).

From now on, we maintain the following regularity assumption.

Assumption 1 (Regularity of the virtual valuation function)**.** *The virtual valuation function $w(c;\xi,\beta)$ is quasiconcave in c.*

If w is decreasing in c, the lemons problem is mild or absent. In this special case of Assumption 1, Myerson (1981) proved that a second price auction is optimal. Assumption 1 allows for w to increase, because it only requires w to be single-peaked. The slope of w is partly determined by the slope of v. If v is sharply increasing there is a severe lemons problem, and w may be increasing in c.

A sufficient (but far from necessary) condition for Assumption 1 to hold is that w be concave in c. If v is concave and $\frac{F}{f}$ is convex, then w is concave. The ratio $\frac{F}{f}$ is convex if F is a Power distribution (of which the Uniform distribution is special case), a Pareto distribution, or an Exponential distribution.[6] Assumption 1 will be used to establish the optimality of a LoLA (Theorem 1).

The buyer can commit to any trading mechanism. By the revelation principle, any equilibrium outcome of any trading procedure is also the truth-telling equilibrium outcome of a direct mechanism. A direct mechanism is a set of $2N$ functions

$$q_i(c_i,\, c_{-i}),\; m_i(c_i,\, c_{-i})$$

[6] If F is a Power distribution then $\frac{F}{f}$ is linear. If $F(c) = 1 - x^{-\alpha}$ is a Pareto distribution $\frac{F}{f}(x)$ is proportional to $x^{\alpha+1} - x$ which is convex in x. If $F(x) = 1 - e^{-\lambda x}$ is an Exponential distribution $\frac{F}{f}(x)$ is proportional to $e^{\lambda x} - 1$ which is convex in x.

that, for each each i and any reported type profile c, specifies the probability that supplier i sells the object, and the expected payment that it receives from the buyer.

4 Results

We are interested in direct mechanisms that maximize any weighted average of the expected buyer surplus and the expected social surplus, with respective weights β and $1-\beta$, for any $\beta \in [0,1]$. Formally, we solve the following maximization problem:

Weighted welfare maximization problem

$$\max_{q,m} \int_{[c_L,c_H]^N} \left[\sum_{i=1}^{N} [v(c_i,\xi) - (1-\beta)\cdot c_i] \cdot q_i(c_i,c_{-i}) - \beta \cdot m_i(c_i,c_{-i}) \right] \prod_{i=1}^{N} f(c_i)\, dc_i \tag{5}$$

subject to, for all i, $c \in [c_L,c_H]^N$, $c_i, c_i' \in [c_L,c_H]$:

$$\sum_{i=1}^{N} q_i(c_i,c_{-i}) \leq 1 \tag{6}$$

$$q_i(c_i,c_{-i}) \geq 0 \tag{7}$$

$$\begin{aligned} &\int_{[c_L,c_H]^{N-1}} [m_i(c_i,c_{-i}) - c_i \cdot q_i(c_i,c_{-i})] \prod_{j\neq i} f(c_j)\, dc_j \\ &\quad \geq \int_{[c_L,c_H]^{N-1}} [m_i(c_i',c_{-i}) - c_i \cdot q_i(c_i',c_{-i})] \prod_{j\neq i} f(c_j)\, dc_j \end{aligned} \tag{8}$$

$$\int_{[c_L,c_H]^{N-1}} [m_i(c_i,c_{-i}) - c_i \cdot q_i(c_i,c_{-i})] \prod_{j\neq i} f(c_j)\, dc_j \;\geq\; 0. \tag{9}$$

In this section we prove that, for any $\xi \geq 0$ and $\beta \in [0,1]$, the above optimization problem is solved by a LoLA with suitably chosen "minimum price" p_L and reserve price p_H. In the optimal LoLA, it is an equilibrium for all suppliers to bid their cost ("sincere bidding"), and this equilibrium generates probabilities $q_i(c_i,\, c_{-i})$ and payments $m_i(c_i,\, c_{-i})$ that solve the above optimization problem. The LoLA format is formally defined next.

Lowball lottery auction (LoLA): formal definition

A LoLA with floor price p_L and reserve price $p_H \geq p_L$ is the following sealed-bid auction.

Let k denote the number of suppliers that bid at or below p_L:

- if $k \leq 1$, the outcome is the same as in the second-price auction with reserve price p_H;
- if $k > 1$, each of these k suppliers sells at price p_L and is paid p_L, with probability $1/k$

Individual rationality is guaranteed in a LoLA because any bidder can opt out by bidding above p_H. The next proposition is the main result of the paper.

Theorem 1 (Optimality of the LoLA). *In a LoLA with any $p_L \leq p_H$, it is a (weakly) dominant strategy for all suppliers to bid their cost. The resulting equilibrium implements the solution to optimization problem (5-9) if p_L and p_H are chosen such that:*

$$p_H{}^* = \sup\{c \in [c_L, c_H] \;\; s.t. \; w(c; \xi, \beta) > 0\}, \tag{10}$$

and

$$p_L{}^* = \sup\left\{p \in [c_L, c_H] \;\; s.t. \; \int_{c_L}^{p} w_c(c; \xi, \beta) \cdot F(c) \cdot dc > 0\right\}. \tag{11}$$

Proof. See Appendix A. ■

The challenge in proving Theorem 1 is that the monotonicity of the allocation function, i.e., the property that lower-cost bidders must win with weakly higher expected probability, can be binding (unless the optimal floor price equals c_L). Hence the standard proof technique, which hinges on side-stepping all monotonicity constraints, cannot be applied in our setting. Our approach relies on finding explicit expressions for the shadow values of violating these constraints, for all types. This is the most innovative part of our proof, and it is done in Lemma 4.

The reserve price $p_H{}^*$ defined in (10) is the same as the reserve price in standard auctions: it is the type at which the virtual valuation w becomes negative. The floor price $p_L{}^*$ defined in (11) identifies the threshold such that, if two or more suppliers have a cost below the threshold, the designer would prefer to source from the *higher-cost* type.

However, given the mononotonicity constraints mentioned above, the best feasible option is to randomize among them. A number of comparative static results about ${p_H}^*$ and ${p_L}^*$ follow immediately from conditions (10) and (11).

Proposition 1 (Comparative statics on ${p_H}^*$ and ${p_L}^*$)**.**

1. *Floor and reserve prices ${p_L}^*$ and ${p_H}^*$ are independent of the number of bidders.*
2. *The floor price is increasing in the severity of the lemons problem, i.e., ${p_L}^*$ is nondecreasing in ξ for any β.*
3. *If F is log-concave, the floor price is increasing in the degree to which the designer takes social welfare into account, i.e., ${p_L}^*$ is nonincreasing in β for any ξ.*
4. *The reserve price is increasing in the degree to which the designer takes social welfare into account, i.e., ${p_H}^*$ is decreasing in β for any ξ.*

Proof. **Part 1** Conditions (10) and (11) do not depend on N.

Part 2 Because $v_{c\xi} \geq 0$ by assumption, increasing ξ shifts the function w_c (at least weakly) upward (see eq. 4), and then condition (11) yields the result.

Part 3 Log-concavity of F implies that the ratio $\frac{F(c)}{f(c)}$ is increasing in c, therefore increasing β shifts the function w_c down (see eq. 4), whence the result follows.

Part 4 Increasing β shifts the function w_c downward (see eq. 4), and then condition (10) yields the result. ■

The property in Part 1 is shared by the reserve price in a standard auction (Myerson 1981). Part 2 says that the floor price is increasing in the parameter ξ that encodes the severity of the lemons problem. This is intuitive, because the only reason to have a floor price is to guard against lowball bidders. It is interesting that this effect obtains even if $\beta = 0$, i.e., when the designer maximizes social welfare. Part 3 requires log-concavity. Since most commonly-used F's are log-concave,[7] "typically," ${p_L}^*$ will be nonincreasing in β. The economic intuition for this results was provided earlier at the end of Section 2: the buyer prefers a lower p_L than the social planner because a small decrease from p_L to $p_L - \epsilon$ has first-order benefits for the buyer (the decrease in payment is realized whenever

[7]See Tables 1 and 3 in Bagnoli and Bergstrom (2005). Log-concavity of F obtains not only whenever f is log-concave (Bagnoli and Bergstrom 2005, Theorem 1) but also, often, when f is not log-concave.

the winning bidder's type is smaller than p_L) but only second-order benefits for the social planner (the allocation improves only when the winning bidder's type is between p_L and $p_L - \epsilon$).

The next result concerns uniqueness.

Proposition 2 (Sincere bidding is the unique equilibrium)**.** *Consider any LoLA with reserve price $p_H < c_H$ and three or more bidders. If the density f is positive on $[c_L, c_H]$ then the equilibrium outcome is unique almost surely. Up to changes of the bid functions on a set of measure zero, any equilibrium strategy profile entails sincere bidding for types with cost above p_L, and bidding any number less than or equal to p_L for all other types.*

Proof. The proof follows almost verbatim that of Proposition 1 in Blume and Heidhues (2004). ■

This result is a direct consequence of Corollary 1 in Blume and Heidhues (2004), who study uniqueness in Vickrey auctions. The reserve price is needed to rule out equilibria of the following form. Fix some $\hat{c} \in (p_L, c_H)$. Bidder 1 bids sincerely if her cost is below $\hat{c}$, and bids $\hat{c}$ otherwise. All other bidders bid sincerely if their cost is below $\hat{c}$, and bid c_H otherwise. In the absence of a reserve price, these strategies constitute an equilibrium. With a reserve price $p_H < c_H$, however, if bidder 1's cost exceeds the reserve price then bidder 1 prefers not to bid at all rather than to follow the recommended strategy.

5 Illustrative application: optimal procurement mechanisms for Italian public sector

This section illustrates the benefits of running the optimal auction in an adverse selection environment. Using information that was generously provided by Francesco Decarolis,[8] we perform a counterfactual experiment on Italian government procurement auctions. By making some stark assumptions about how quality enters the government's objective function (expression 12), we are able to compute the gain (buyer surplus) that the government could have made, had it used the optimal mechanism – which, conveniently, happens to be a LoLA – relative to a first-price auction, which is the format the government actually used.

[8] This information relates to Decarolis' (2014, 2018) structural analysis of Italian procurement firms.

The goal of this section is not to give policy recommendations, but merely to sketch out how real-world data can be used to find the optimal mechanism. Therefore, we forego the battery of robustness checks that would be essential if our goal was to give policy recommendations.

5.1 The available data

The available data is depicted in Figure 3. Panel A shows the estimated distribution of bidder costs $\hat{f}$, which was structurally estimated by Decarolis (2018) and corresponds to our $f(c)$.[9] Panels B and C show the empirical distributions of two measures of the auction winner's quality: the delivery delay ratio D, and the cost overrun ratio O.[10] The figure indicates that, in most cases, the government suffers a delay, a cost overrrun, or both.[11]

[9]In Decarolis' (2018) structural model, supplier i's cost in a given auction is given by:

$$c_i = y + z_i,$$

where the z_i's are idiosyncratic and privately-known cost components, and y is an auction-specific and commonly-known scalar. Decarolis (2018) estimates that $z_1, ..., z_N$ are i.i.d. draws from a random variable Z whose density is depicted in Figure 3, panel A. In what follows we assume, without loss of generality, that $y = 0$, which allows us to interpret z_i's as c_i's.

[10]Delay ratios D are measured as the difference between contractually-stipulated and actual delivery dates, divided by the former. Cost overrun ratios O are measured as the difference between the money eventually paid by the government and the winning bid, divided by the auction's reserve price.

[11]Note, for future reference, that panels B and C display the quality supplied by the *winner* in a first-price auction, which is not representative of the quality that would have been supplied by a *random bidder*.

Distributions of cost and quality measures

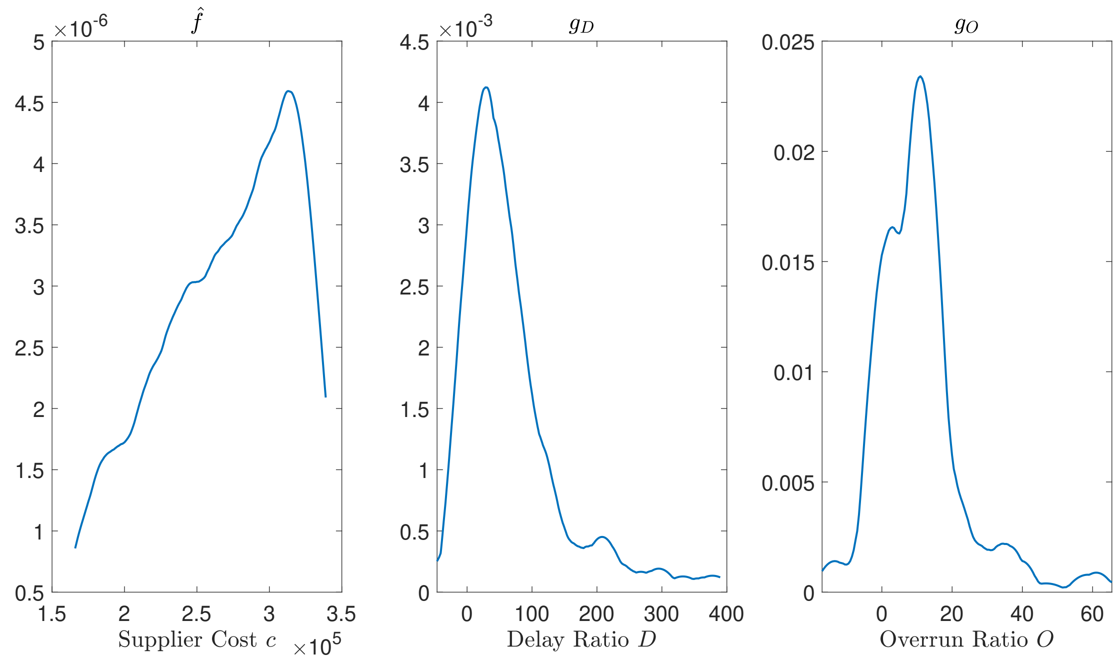

Figure 3: The left-hand panel depicts the estimated p.d.f. $\hat{f}_Z$ of the idiosyncratic cost component Z (unit is 10^5 euros) from Decarolis' (2018) assumed cost structure $c_i = y + z_i$, where the z_i's are iid draws from Z, and y is an auction-specific scalar. Without loss of generality we normalize $y = 0$, which allows us to replace z_i with c_i in the left-hand panel. The middle and right-hand panels display the empirical marginal distributions g_D and g_O of, respectively: the *delay ratio* D, which is the difference between the actual and the contractual time, as a percentage of the contractual time; and of the *overrun ratio* O which is the difference between the final payment and the winning bid as a percentage of the reserve price. See Decarolis (2014, p. 117). Kernel (Epanechnikov) smoothed distributions, the bandwidth used are 11000, 18.15 and 3.0071 respectively. Data generously provided by Francesco Decarolis.

5.2 Calibrating the buyer's payoff function $v(c, \xi)$

Based on these three distributions, we seek to obtain a calibrated counterpart for our theoretical construct $v(c, \xi)$. To cut down on expositional complexity, we assume the starkest possible functional form:

$$v(c, \xi) = const - K\mathbb{E}\left[D(c, \xi) + O(c, \xi)\right], \tag{12}$$

where $D(c, \xi)$ and $O(c, \xi)$ are unobserved random variables that represent the delays and cost overrruns, respectively, that are stochastically delivered by a supplier with cost c, conditional on the parameter ξ. The rationale for the minus sign is that delays and cost overruns *decrease* the buyer's value. K is a positive scaling parameter whose value will be calibrated later.[12]

The parameter ξ in expression (12) moderates the correlation between a supplier's cost c, and the qualities D and O stochastically provided by that supplier. This role appears to be conceptually different from the interpretation given to ξ in our theoretical model: in the theory, ξ is conceptualized as a buyer *type*; in(12), ξ is conceptualized as a feature of the supply-delivery technology. This conceptual distinction does not make a difference here because, operationally, what matters is that ξ determines the slope of the buyer's valuation, as it does in expression (13) below.

The distributions of the random variables $D(c, \xi)$ and $O(c, \xi)$ are as yet unspecified. We calibrate them semi-parametrically by requiring that, given that $c \sim \hat{f}$, their distributions for any given ξ coincide with the empirical marginal distributions g_D and g_O depicted in Figure 3.[13] Definition 2 in Appendix B provides formulae for constructing calibrated $\hat{D}(c, \xi)$ and $\hat{O}(c, \xi)$ with the desired marginals, *for any value of the parameter* ξ. Using these formulae allows us not to take a stand on the value of ξ. Plugging these formulae into expression (12) yields the following expression for the calibrated buyer payoff function:

$$\begin{aligned} \hat{v}(c, \xi) &= const - K\mathbb{E}\left[\hat{D}(c, \xi) + \hat{O}(c, \xi)\right] \\ &= const(\xi) - \xi K \left[\delta(c) + \omega(c)\right], \end{aligned} \tag{13}$$

where $const(\xi)$ is independent of c and, from Definition 2, we have:

$$\begin{aligned} \delta\left(c\right) &= G_D^{-1}\left(\left[1 - \hat{F}\left(c\right)\right]^N\right), \\ \omega\left(c\right) &= G_O^{-1}\left(\left[1 - \hat{F}\left(c\right)\right]^N\right), \end{aligned}$$

(refer to Appendix B.2 for the computations). Expression (13) is the calibrated buyer's

[12] There is no difficulty in making expression (12) more complex. For example, one could pre-multiply $D(c, \xi)$ and $O(c, \xi)$ by positive constants, and the analysis would be essentially unchanged. We don't do this, in order to minimize expositional complexity.

[13] Formally this means that, denoting the winning bidder's cost by $C_{(1)} = \min\{C_1, ..., C_N\}$, the random variable $D(C_{(1)}, \xi)$ has density g_D, and $O(C_{(1)}, \xi)$ has density g_O.

payoff. This expression is a fully specified function of (c, ξ) up to a constant. Indeed, the three quantities $\hat{F}$, G_D, and G_O are given in Figure 3; and the parameters N, K are assigned numerical values as described in Appendix B.2.

The parameter ξ will be treated as a free parameter. This parameter determines the sensitivity of the buyer's payoff to the quality concerns. If $\xi = 0$ the function $\hat{v}(c, \xi)$ does not depend on c and, therefore, there buyer has no quality concerns. If $\xi > 0$ the function $\hat{v}(c, \xi)$ is increasing in c (this is because $\delta(c)$ and $\omega(c)$ are decreasing functions of c). Intuitively, the parameter ξ modulates the buyer's quality concerns because, in the construction of $\hat{D}(c, \xi)$ and $\hat{O}(c, \xi)$, this parameter governs the correlation between supplier cost and quality.

The function $\hat{v}$ satisfies the two theoretical assumption imposed on page 8. Indeed, it can be checked from expression (13) that $\hat{v}_{c\xi} \geq 0$. Furthermore, we can (and will) make $const(\xi)$ in expression (13) large enough that $\hat{v}(c_L, \xi) \geq c_L$ for all $\xi \in [0, 1]$.

5.3 Buyer-optimal and socially optimal mechanisms are LoLAs

We compute the calibrated virtual valuation function:

$$\hat{w}(c; \xi, \beta) \equiv \hat{v}(c; \xi) - c - \beta \frac{\hat{F}(c)}{\hat{f}(c)}, \tag{14}$$

by substituting $\hat{v}$ from (13) and $\hat{F}$ from Figure 3 into the expression for the virtual valuation (4). We set $const(\xi)$ large enough that the virtual valuation (14) is positive for all values of c and β, which implies that it is optimal not to set any reserve price p_H in the LoLA.[14]

Each of the left-hand panels in Figure 4 displays $\hat{w}$ as a function of c, for $\beta = 0$ (gains from trade, dashed red line) and $\beta = 1$ (buyer's virtual valuation, solid blue line). These functions are shown for $\xi = 0$, 0.33, 0.67, and 1, respectively, in panels A-D. In all four left-hand panels, the buyer's virtual valuation and the gains from trade happen to be quasi-concave functions of c, so Assumption 1 is satisfied. Therefore, by Theorem 1 the LoLA is the buyer-optimal and the socially-optimal auction for all displayed values of ξ.

[14] $const(\xi)$ can be made arbitrarily large by setting $const$ large enough in expression (12): refer to Appendix B.2 for information about the calibration.

The right-hand panels of Figure 4 are calibrated counterparts to Figure 2. Each right-hand panel displays the expected buyer (solid blue line) and social (dashed red line) surplus in a LoLA with floor price p_L. The optimal floor prices are determined by equation (11) after setting β equal to one or zero: accordingly, they maximize the expected (buyer or social) surplus, as shown in Figure 4. Within each right-hand panel, the socially optimal floor price always exceeds the buyer-optimal one. This is a consequence of Proposition 1 part 3, because the estimated cost distribution $\hat{F}$ happens to be log-concave (see Figure 6).

As we move down from panel A to panel D, the parameter ξ (correlation between cost and quality) increases. Therefore, the buyer's quality concerns also increase, causing more-costly suppliers to become more socially valuable (as we move down the left-hand panels, the gains-from-trade dashed red line becomes increasing). Consistent with Proposition 1 part 2, the buyer-optimal and socially optimal floor prices increase with ξ: see the righ-hand panels. For low values of ξ, the buyer- and socially optimal auctions coincide with a first (or equivalently, second) price auction because the optimal floor prices coincide with c_L. As ξ increases, the optimal floor prices increase until, for sufficiently high values of ξ, the supplier is randomly selected in the socially optimal auction.

Optimal mechanisms with varying degrees of quality concerns

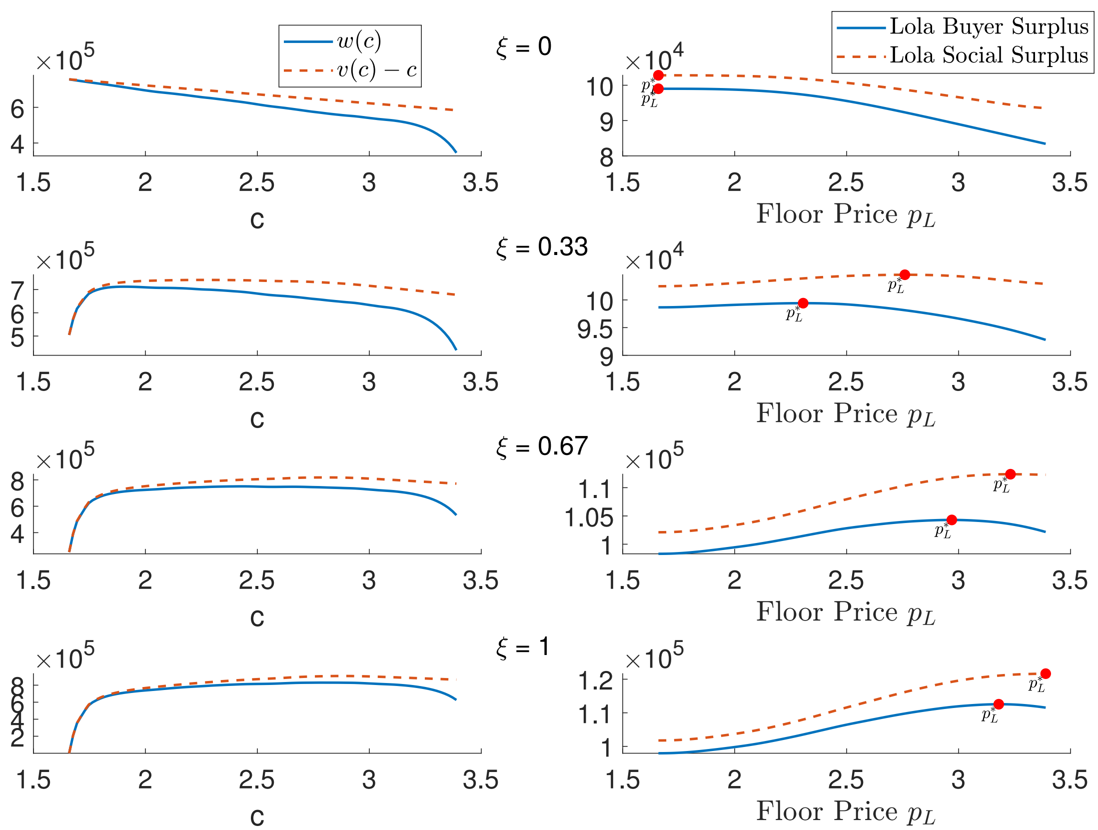

Figure 4: Virtual valuation functions $w(c)$ and gains from trade $v(c) - c$ for different values of ξ (left-hand column); expected buyer and social surplus in a LoLA with floor price p_L and no reserve price for different values of ξ (right-hand column). Recall that in our calibration it is optimal not to have a reserve price. Units of c are 10^5. As quality concerns increase (i.e., ξ increases), more-costly suppliers become more socially valuable (left panel, dashed red line). With minimal quality concerns, the optimal LoLAs reduce to standard auctions, i.e., first- or second-price auctions ($\xi = 0$, top right panel). With maximal quality concerns, the socially optimal LoLA reduces to the random allocation mechanism ($\xi = 1$, bottom-right panel).

5.4 Performance of the buyer-optimal mechanism vs. first-price auction

Figure 5 shows the performance gain of the buyer-optimal mechanism, which in our case is a LoLA with optimal floor price p_L^* and no reserve price, over a first-price (or, which is the same in our case, a second-price) auction, as ξ varies.[15] We analyze three performance metrics: expected buyer surplus (top panel), expected supplier profit (middle panel), and expected social surplus (bottom panel). In all three metrics, the buyer-optimal LoLA outperforms a conventional auction: for example, when $\xi = 1$ buyer surplus is 15% higher in the optimal LoLA than in a first price auction. The performance gain is increasing in the level of ξ, as one would expect. Even at relatively lower levels of $\xi \approx 0.5$, that is, when the quality concerns are relatively mild, a LoLA affords gains in the 2.5% range, which are nontrivial from a policy perspective.

Performance of buyer-optimal mechanisms relative to first-price auction

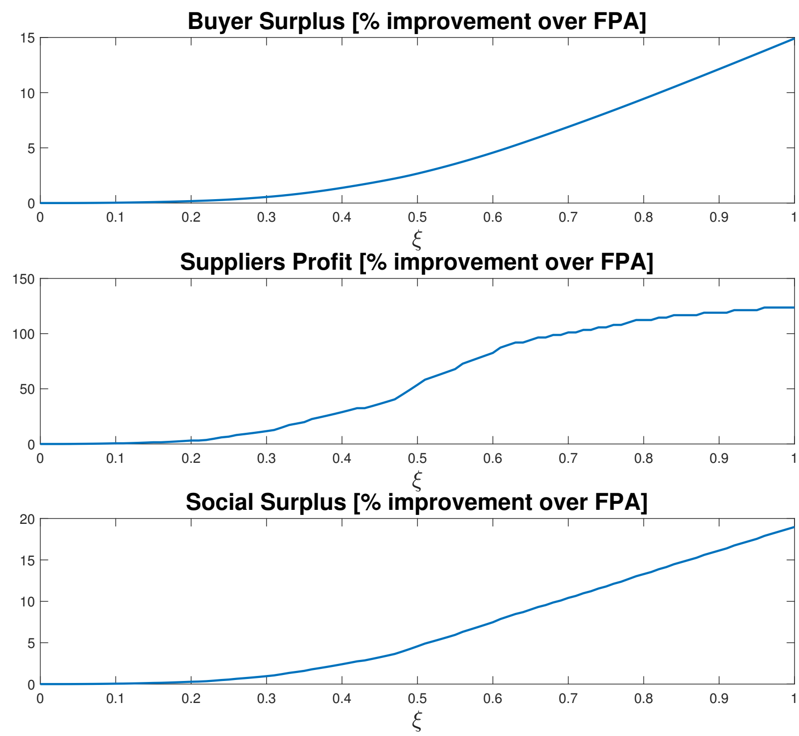

Figure 5: Performance improvement of optimal LoLA over first-price (or second-price) auction.

[15]The optimal floor p_L^* (not shown in the figure) changes as ξ varies.

6 Software for computing optimal procurement mechanisms

This section describes two software applications that we have created and made publicly available.[16] These applications compute the buyer-optimal procurement mechanisms in the presence of quality concerns. The purpose of disseminating these applications is twofold. First, we wish to allow business practitioners to assess whether they can benefit from a buyer-optimal LoLA and, if so, with what floor and reserve prices. Second, for pedagogical purposes, we want to facilitate the teaching of this paper in an engaging way.

6.1 App 1

An Excel-based visual interface asks the user to input a probability distribution of costs (corresponding to $f(c)$ in our theoretical model), a function $v(c)$ (corresponding to $v(c, \xi)$ for some fixed ξ), and the number of bidders N. The application assumes, as we did, that bidder costs are drawn independently from the cost distribution, and requires that $v(c_L) > c_L$.

The program then calls on Matlab to compute the virtual valuation function $w(c)$ using equation (4), and displays it. A prompt asks the reader to verify that $w(c)$ is quasi-concave, i.e., to check that Assumption 1 holds.[17] After this user check, if a reserve price is buyer-optimal then the program alerts the user and visually identifies the optimal reserve price in the graph of $w(c)$. Finally, the program displays the buyer and social surplus functions as a function of the LoLA floor price p_L, and displays the optimal floor and reserve prices (analogous to the right-hand panel of Figure 4). The program also displays the ratio between the social (or buyer) surplus under a LoLA with reserve price p_L, over a first price auction.

Screenshots for App number 1, and further details, are available in Appendix C.1.

[16] Downloadable from https://www.alessandrotenzinvilla.com/research.html.

[17] This check could have been done automatically, but prompting the reader to visually verify Assumption 1 has pedagogical value.

6.2 App 2

The second application is not visual, but it is more powerful. This application is realized in Matlab and IBM ILOG CPLEX. The application requires the same inputs as Application 1, and it computes the optimal mechanism *even when that mechanism is not a LoLA*. Therefore, Application 2 dispenses with Assumption 1 and with the requirement that $v(c_L) > c_L$. The application yields the buyer-optimal direct revelation mechanism, expressed through the interim probability $Q(c)$ that a generic bidder with cost c wins the auction. This application is helpful to deal with settings where assumptions made in this paper are violated, and so Theorem 1 does not apply.

7 Conclusions

Adverse selection is a major concern in procurement. In this paper we have presented a mechanism called LoLA which, under some regularity conditions, is the best incentive compatible mechanism for maximizing either the seller's or the social surplus (or any combination thereof). The mechanism features a floor (or minimum) price, and a reserve (or maximum) price. The sincere-bidding equilibrium of the LoLA is in dominant strategies, implements the surplus-maximizing allocation, and is unique under mild regularity conditions.

To illustrate the gains from the optimal mechanism, we performed a counterfactual experiment on Italian government procurement auctions. We computed the gain that the government could have made, had it used the optimal mechanism (which happens to be a LoLA), relative to a first-price auction, which is the format the government actually used. We find that, in a reasonably calibrated model, these savings can be nontrivial.

Our analysis has sidestepped the issues of repeated interaction and collusion. In the presence of collusion, it is possible that the presence of a floor price might help, as has been suggested in the literature. However, finding the optimal mechanism in the presence of collusion is beyond the scope of this paper.

We hope that our analysis can lead procurement agencies to consider experimenting with the LoLA.

References

[1] Albano, Gian Luigi, Milo Bianchi, and Giancarlo Spagnolo, 2006. "Bid Average Methods in Procurement." *Rivista di Politica Economica*, SIPI Spa, vol. 96(1), pages 41-62, January-F.

[2] Bagnoli, Mark, and Ted Bergstrom (2005). "Log-concave probability and its applications." Economic theory 26.2 (2005): 445-469.

[3] Blume, Andreas, and Paul Heidhues (2004). "All equilibria of the Vickrey auction." *Journal of economic Theory*, 114(1), 170-177.

[4] Chassang, Sylvain, and Juan Ortner. "Collusion in auctions with constrained bids: Theory and evidence from public procurement." *Journal of Political Economy* 127.5 (2019): 2269-2300.

[5] Dini, Federico, Riccardo Pacini, and Tommaso Valletti. "Scoring rules." *Handbook of procurement* (2006): 293-321.

[6] Calzolari, Giacomo and Giancarlo Spagnolo, (2006). "Reputation and collusion in procurement." Manuscript, University of Bologna and Stockholm School of Economics.

[7] Conley, Timothy G. and Francesco Decarolis (2016) "Detecting bidders groups in collusive auctions" *American Economic Journal: Microeconomics*, 2016

[8] Decarolis, Francesco (2014) "Awarding price, contract performance, and bids screening: Evidence from procurement auctions" *American Economic Journal: Applied Economics*, 2014

[9] Decarolis, Francesco (2018) "Comparing public procurement auctions" *International Economic Review*, 2018.

[10] Decarolis, Francesco and Michael Klein (2012) "Too Good to Be True Auctions," World Bank – Public Private Infrastructure Advisory Facility (PPIAF) Note, (2012). Part of Body of Knowledge on Infrastructure Regulation (BoKIR). Previoulsy circulated as Working Paper 186, Frankfurt am Main: Frankfurt School of Finance & Management

[11] Lopomo, Giuseppe (2001). "Optimality and robustness of the english auction." *Games and Economic Behavior* 36.2 (2001): 219-240.

[12] Manelli, Alejandro M., and Daniel R. Vincent (1995). "Optimal procurement mechanisms." *Econometrica* (1995): 591-620.

[13] Manelli, Alejandro M., and Daniel R. Vincent (2004) "Duality in Procurement Design," *Journal of Mathematical Economics*, 40: 411–428, 2004.

[14] Maskin. Eric and John Riley (1984) "Optimal Auctions with Risk Averse Buyers", *Econometrica*, Vol. 52, No. 6 (Nov., 1984), 1473-1518.

[15] Myerson, Roger B. "Optimal auction design." *Mathematics of Operations Research* 6.1 (1981): 58-73.

Appendices

A Proofs

For any pair p_L and p_H such that $c_L \leq p_L \leq p_H \leq c_H$, consider the LoLA with threshold prices p_L and p_H. Sincere bidding in the LoLA induces the following outcome:

$$q_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) \equiv \begin{cases} \frac{1}{\kappa+1}, & \text{if } c_i \leq p_L \;\; \& \; c_{-i}^{(\kappa)} \leq p_L < c_{-i}^{(\kappa+1)} \\ 1, & \text{if } p_L < c_i \;\; \& \; c_i < c_{-i}^{(1)} \end{cases} \tag{15}$$

and

$$m_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) \equiv \begin{cases} \frac{1}{\kappa+1} \cdot p_L, & \text{if } c_i \leq p_L \;\; \& \; c_{-i}^{(\kappa)} \leq p_L < c_{-i}^{(\kappa+1)} \\ c_{-i}^{(1)}, & \text{if } p_L < c_i \;\; \& \; c_i < c_{-i}^{(1)} \end{cases} \tag{16}$$

where $c_{-i}^{(\kappa)}$ denotes the κ-th lowest cost among all supplier i's opponents.

The functions $(q^{\mathrm{L}}, m^{\mathrm{L}})$ may also be interpreted as a direct revelation mechanism. We now show that, in this direct revelation mechanism, truthful reporting is a (weakly) dominant strategy.

Lemma 1. *$(q^L,\ m^L)$ satisfies, $\forall i = 1, ..., N$,*

$$\forall c_i, c_i', c_{-i}, \qquad m_i(c_i, c_{-i}) - c_i \cdot q_i(c_i, c_{-i}) \geq m_i(c_i', c_{-i}) - c_i \cdot q_i(c_i', c_{-i}) \tag{17}$$

and

$$\forall c_i, c_{-i}, \qquad m_i(c_i, c_{-i}) - c_i \cdot q_i(c_i, c_{-i}) \geq 0. \tag{18}$$

Proof. It is well known in mechanism design that conditions (17-18) hold if and only if the following conditions hold jointly: $\forall c_{-i} \in [c_L, c_H]^{N-1}$

$$m_i^{\mathrm{L}}(c_H, c_{-i}; p_L, p_H) \geq c_H \cdot q_i^{\mathrm{L}}(c_H, c_{-i}; p_L, p_H) \tag{19}$$

$$q_i^{\mathrm{L}}(\cdot, c_{-i}; p_L, p_H) \;\text{ is nonincreasing,} \tag{20}$$

and

$$\forall\, c_i \in [c_L, c_H] \qquad m_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) = c_i \cdot q_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) + \int_{c_i}^{c_H} q_i^{\mathrm{L}}(t, c_{-i}; p_L, p_H)\, dt. \tag{21}$$

Therefore, it suffices to show that (19-21) hold. To this end, observe that the inequalities in (19) and the monotonicity in (20) are immediate. The envelope condition in (21) holds because both m^{L} and q^{L} are constant in c_i on $[c_L, p_L) \cup (p_L, c_H]$, and

$$p_L \cdot \left[\lim_{x \uparrow p_L} q_i^{\mathrm{L}}(x, c_{-i}; p_L, p_H) - \lim q_{x \downarrow p_L} q_i^{\mathrm{L}}(x, c_{-i}; p_L, p_H) \right] = \lim_{x \uparrow p_L} m_i^{\mathrm{L}}(x, c_{-i}; p_L, p_H) - \lim_{x \downarrow p_L} m_i^{\mathrm{L}}(x, c_{-i}; p_L, p_H)$$

■

Our strategy of proof will involve restricting attention to candidate mechanisms that are symmetric, and this will be without loss of generality. Next, we introduce a formal definition of symmetric mechanism.

Definition 1. *A mechanism $(q_i, m_i)_{i=1,...,N}$ is symmetric if, for all i,*

$$q_i(c_{\pi(1)}, c_{\pi(2)}, ..., c_{\pi(N)}) = q_{\pi(i)}(c_1, c_2, ..., c_N),$$

and

$$m_i(c_{\pi(1)}, c_{\pi(2)}, ..., c_{\pi(N)}) = m_{\pi(i)}(c_1, c_2, ..., c_N),$$

for every permutation π of $\{1, 2, ..., N\}$. A symmetric mechanism is given by two functions

$$q \equiv q_1 : [c_L, c_H]^N \to [0, 1] \quad \text{and} \quad m \equiv m_1 : [c_L, c_H]^N \to [0, 1]$$

which are invariant to permutations of the last $N-1$ variables, i.e., letting $\mathcal{N}$ be the set of numbers $\{1, ..., N\}$, $\forall i \in \mathcal{N}$, $\forall$ permutation π of $\mathcal{N}$ we have:

$$q_i(c_1, c_2, ..., c_n) = q(c_i, c_2, ..., c_{i-1}, c_1, c_{i+1}, ..., c_N),$$

and

$$m_i(c_1, c_2, ..., c_n) = m(c_i, c_2, ..., c_{i-1}, c_1, c_{i+1}, ..., c_N).$$

If we restrict attention to symmetric mechanisms, the original weighted welfare maximization problem (5- 9) can be written more simply. We write down the simplified problem next and then, in Lemma 2, we show that the two maximization problems are equivalent.

First reformulation of the weighted welfare maximization problem

$$\max_{Q,M}\; N\int_{[c_L,c_H]} \left[\,[v(c_i,\xi)-(1-\beta)\cdot c_i]\cdot Q(c_i)-\beta\cdot M(c_i)\right] f(c_i)\,dc_i \tag{22}$$

subject to:

for all $c_i, c_i' \in [c_L, c_H]$:

$$M(c_i)-c_i\cdot Q(c_i)\geq M(c_i')-c_i\cdot Q(c_i'), \tag{23}$$

for all $c_i \in [c_L, c_H]$:

$$M(c_i)-c_i\cdot Q(c_i)\geq 0, \tag{24}$$

for all $c_i \in [c_L, c_H]$:

$$Q(c_i)\geq 0, \tag{25}$$

and

$$N\int_{c_L}^{c_1} Q(y)\,f(y)\,dy \;\leq\; 1-\left[1-F(c_1)\right]^N. \tag{26}$$

Lemma 2. *Restrict attention to symmetric mechanism. The value of the weighted welfare maximization problem (5- 9) is the same as the value of problem (22- 26).*

Proof. Define:

$$\begin{aligned} Q(c_1) &\equiv \int_{[c_L,c_H]^{N-1}} q(c_1,\,c_{-1})\cdot\prod_{j>1} dF(c_j) \\ M(c_1) &\equiv \int_{[c_L,c_H]^{N-1}} m(c_1,\,c_{-1})\cdot\prod_{j>1} dF(c_j). \end{aligned} \tag{27}$$

Because in solving problem (5- 9) we are restricting attention to mechanisms $(q_i, m_i)_{i=1,\ldots,N}$ that are symmetric, the objective function (5) can be re-written as (22). Similarly, the constraints (8) and (9) can be re-written as: (23) and (24). Furthermore, Border (1991) proves that, if the function q is symmetric in the sense of Definition 1, the demand constraints (6) and nonnegativity constraints (7) hold if and only if (25) and (26) are satisfied.

■

Problem (22- 26) can be further simplified, as follows.

Second reformulation of the weighted welfare maximization problem

$$\max_{Q} \; N \int_{c_L}^{c_H} w(c;\xi,\beta) \cdot Q(c) \cdot f(c)\, dc \tag{28}$$

where $w(c;\xi,\beta)$ is defined in (4), subject to:

$$Q \text{ is nonincreasing,} \tag{29}$$

and, for all $c \in [c_L, c_H]$:

$$Q(c) \geq 0, \tag{30}$$

and

$$N \int_{c_L}^{c} Q(y)\, f(y)\, dy \;\leq\; 1 - [1 - F(c)]^{N}. \tag{31}$$

Lemma 3. *The weighted welfare maximization problem (22- 26) can be reformulated as (28- 31).*

Proof. The incentive constraints (23) and (24) can be replaced without loss of generality by (29) and the envelope condition:

$$\forall\, c \in [c_L, c_H] \qquad M(c) = c \cdot Q(c) + \int_{c}^{c_H} Q(t)dt. \tag{32}$$

(This result is standard: see, e.g., Proposition 5.2 at p. 66 of Krishna 2010). Next, we use (32) to eliminate M from the problem. Substituting it into (22) and simplifying yields (28). Finally, (30) and (31) are identical to (25) and (26). ■

Next is the final reformulation of the problem.

Final (relaxed) formulation of the weighted welfare maximization problem

$$\max_{Q} \; N\int_{c_L}^{c_H} w(c;\xi,\beta)\cdot Q(c)\cdot f(c)\,dc \tag{33}$$

where $w(c;\xi,\beta)$ is defined in (4), subject to:

$$N\int_{c_L}^{c_H} w(c;\xi,\beta)\cdot Q(c)\cdot f(c)\,dc \;\leq\; N\int_{c_L}^{c_H} w(c;\xi,\beta)\cdot Q^L(c,p_L^*,p_H^*)\cdot f(c)\,d(c), \tag{34}$$

where $Q^L(c,p_L^*,p_H^*)$ is given by expression (27) with q being replaced by $q_i^{\mathrm{L}}(c_i,c_{-i};p_L^*,p_H^*)$ from expression (15).

Problem (33-34) below is actually a relaxation of (28- 31). Aggregating constraints (29-31) into the single inequality (34) is the most innovative part of the proof. This aggregation is proved in the next lemma.

Lemma 4. *Any allocation function Q that satisfies (29-31) also satisfies (34).*

Proof. The proof consists of multiplying both sides of each inequality (29-31) by a non-negative multiplier (which does not change the constraint), and then integrating over c on both sides of each constraint, and finally summing the three resulting inequalities. The resulting inequality identifies a superset of the original feasible set, and happens to equal (34).

The multipliers equal zero except:

$$\begin{cases} \forall c\in(p_H^*,c_H]: & \eta(c)\equiv -w(c;\xi,\beta)\cdot f(c) \\[2ex] \forall c\in(p_L^*,p_H^*): & \delta(c)\equiv -w'(c;\xi,\beta) \\[2ex] \forall c\in[c_L,p_L^*): & \mu(c)\equiv \dfrac{F(c)}{F(p_L^*)}\displaystyle\int_{c_L}^{p_L^*} w(t;\xi,\beta)\,dF(t)-\int_{c_L}^{c} w(t;\xi,\beta)\,dF(t) \end{cases} \tag{35}$$

To save on notation, in the rest of this proof we omit the dependence of w on (ξ,β).

Let us first show that the multipliers are nonnegative. We have $\eta(c)\geq 0\;\;\forall c\in(p_H^*,c_H]$, because w is negative on the interval $(p_H^*,c_H]$. We have $\delta(c)\geq 0\;\;\forall c\in(p_L^*,p_H^*)$, because

w is decreasing on the interval $[p_L^*, p_H^*]$. Finally, consider μ on $[c_L, p_L^*)$. First note that

$$\mu(c_L) = \mu(p_L^*) = 0 \tag{36}$$

If $c_L < p_L^*$, then the definition of p_L^* in (11) implies $w(c_L) < w(p_L^*)$. Since w is quasiconcave, there exists a point p_0 such that $w(p_0) = w(p_L^*)$ and

$$\begin{aligned} &\forall c \in [c_L, p_0) \quad w(p_L^*) - w(c) \geq 0, \text{ and} \\ &\forall c \in (p_0, p_L^*] \quad w(p_L^*) - w(c) \leq 0, \end{aligned}$$

Thus the derivative

$$\mu'(c) = f(c)\,[w(p_L^*) - w(c)]$$

is positive for $c < p_0$, and negative for $c > p_0$, that is μ is single-peaked on $[c_L, p_L^*]$. This, together with (36), implies that μ is nonnegative on $[c_L, p_L^*]$. Thus nonnegativity is established.

Now, we multiply both sides of: (29) by $\mu(c)$, (30) by $\eta(c)$, (31) by $\delta(c)$. We then integrate over c. Finally, we sum the three resulting inequalities. We arrive at:

$$\int_{c_L}^{p_L^*} \mu(y)\, dQ(y) + \int_{p_L^*}^{p_H^*} \delta(t) \int_{c_L}^{t} Q(y) f(y)\, dy\, dt - \int_{p_H^*}^{c_H} \eta(y)\, Q(y)\, dy \leq \int_{p_L^*}^{p_H^*} \delta(c)\, B(c)\, dc. \tag{37}$$

where

$$B(c) \equiv \frac{1}{N} \cdot \left(1 - [1 - F(c)]^N\right) \quad c \in [c_L, c_H]. \tag{38}$$

Let's focus first on the LHS of (37). The first integral can be rewritten as:

$$\begin{aligned} &\overbrace{\mu(p_L^*)}^{=0} \cdot Q(p_L^*) - \overbrace{\mu(c_L)}^{=0} \cdot Q(c_L) - \int_{c_L}^{p_L^*} Q(y)\, \mu'(y)\, dy \\ = -&\int_{c_L}^{p_L^*} \mu'(y)\, Q(y)\, dy. \end{aligned} \tag{39}$$

The second integral on the LHS of (37) can be rewritten as:

$$\begin{aligned}&\int_{c_L}^{p_L^*}\left(\int_{p_L^*}^{p_H^*}\delta(t)dt\right)Q(y)f(y)dy+\int_{p_L^*}^{p_H^*}\left(\int_{y}^{p_H^*}\delta(t)dt\right)Q(y)f(y)dy\\ =&\int_{c_L}^{p_L^*}\left(\int_{p_L^*}^{p_H^*}\delta(t)dt\right)Q(y)f(y)dy+\int_{p_L^*}^{p_H^*}w(y)Q(y)f(y)dy,\end{aligned} \tag{40}$$

where equality holds because:

$$\int_{y}^{p_H^*}\delta(t)dt=w(p_H^*)-\int_{y}^{p_H^*}w'(c)dc=w(y).$$

Adding (39) and (40) yields:

$$\begin{aligned}&\int_{c_L}^{p_L^*}\left(\int_{p_L^*}^{p_H^*}\delta(t)dt-\frac{\mu'(y)}{f(y)}\right)Q(y)\,f(y)\,dy\\ =&\int_{c_L}^{p_L^*}w(y)\,Q(y)\,f(y)\,dy,\end{aligned} \tag{41}$$

where equality holds because:

$$\begin{aligned}\int_{p_L^*}^{p_H^*}\delta(t)\,dt-\frac{\mu'(y)}{f(y)} &= \overbrace{w(p_H^*)-\int_{p_L^*}^{p_H^*}w'(t)dt}^{=w(p_L^*)}-\frac{\mu'(y)}{f(y)}\\ &= w(p_L^*)-\left(\frac{1}{F(p_L^*)}\cdot\int_{c_L}^{p_L^*}w(t)f(t)dt-w(y)\right)\\ &= \overbrace{w(p_L^*)-\frac{1}{F(p_L^*)}\cdot\int_{c_L}^{p_L^*}w(t)f(t)dt}^{=0}+w(y)\\ &= w(y)\end{aligned}$$

The third integral on the LHS of (37) can be rewritten as:

$$-\int_{p_H^*}^{c_H}\eta(y)Q(y)dy=\int_{p_H^*}^{c_H}w(y)f(y)Q(y)dy. \tag{42}$$

Combining (41) and (42) we can rewrite the LHS of (37) as

$$\int_{c_L}^{c_H} w(y)Q(y)f(y)dy.$$

Let's now focus on the RHS of (37). Plugging in the expression for δ and simplifying yields:

$$\begin{aligned}
&-\int_{p_L^*}^{p_H^*} w'(c)\, B(c)\, dc \\
&= w(p_L^*)\, B(p_L^*) + \int_{p_L^*}^{p_H^*} w(c)\, B'(c)\, dc \\
&= \int_{p_L^*}^{p_H^*} w(c)\, (1 - F(c))^{N-1}\, f(c)\, dc + \int_{c_L}^{p_L^*} \frac{1 - [1 - F(p_L)]^N}{N \cdot F(p_L)} w(c)\, f(c)\, dc,
\end{aligned} \tag{43}$$

where the second equality follows from the definition of B in (38). Now observe that:

$$\begin{aligned}
Q^L(c_1; p_L^*, p_H^*) &\equiv \int_{[c_L, c_H]^{N-1}} q(c_1,\, c_{-1}; p_L^*, p_H^*) \cdot \prod_{j>1} dF(c_j) \\
&= \begin{cases} 0, & c_1 \in (p_H^*, c_H]\,; \\ [1 - F(c_1)]^{N-1}, & c_1 \in (p_L^*, p_H^*]\,; \\ \frac{1-[1-F(p_L)]^N}{N \cdot F(p_L^*)}, & c_1 \in [c_L, p_L^*]\,. \end{cases}
\end{aligned} \tag{44}$$

Hence (43) boils down to:

$$\begin{aligned}
&\int_{p_L^*}^{p_H^*} w(c)\, Q^L(c, p_L^*, p_H^*)\, f(c)\, dc + \int_{c_L}^{p_L^*} w(c)\, Q^L(c, p_L^*, p_H^*)\, f(c)\, dc \\
&= \int_{c_L}^{p_H^*} w(c)\, Q^L(c, p_L^*, p_H^*)\, f(c)\, dc \\
&= \int_{c_L}^{c_H} w(c)\, Q^L(c, p_L^*, p_H^*)\, f(c)\, dc.
\end{aligned}$$

This completes the proof.

■

We are now ready to prove Theorem 1.

Proof of Theorem 1

Proof. Lemma 1 shows that the direct mechanism $(q^{\mathrm{L}}, m^{\mathrm{L}})$ satisfies both IC and IR *ex post*. Therefore, sincere bidding in the LoLA is a (weakly) dominant strategy equilibrium.

Moreover, $(q^{\mathrm{L}}, m^{\mathrm{L}})$ is a feasible mechanism, i.e., it satisfies constraints (6 - 9). Indeed, unit demand (6) and nonnegativity (7) can be checked directly from the definition (15), and the fact that $(q^{\mathrm{L}}, m^{\mathrm{L}})$ satisfy the ex-post incentive constraints, as proved in Lemma 1, immediately implies that it also satisfies their interim counterparts (8) and (9).

It remains to show that the mechanism $(q^{\mathrm{L}}, m^{\mathrm{L}})$ defined in (15) and (16) solves the weighted welfare problem. We proceed in two steps.

Maskin and Riley (1986, footnote 11) show that, in our setting, given any optimal mechanism for the weighted welfare problem, there is a symmetric mechanism that attains the same (maximal) value. Therefore, we can restrict the search for an optimal mechanism to the set of symmetric mechanisms (of which $(q^{\mathrm{L}}, m^{\mathrm{L}})$ is one) without loss of generality.

After restricting to symmetric mechanisms, Lemmas 2 -4 yield a relaxed problem with a set of feasible mechanisms (34) that contains the original feasible set. If a LoLA solves this relaxed problem, then a fortiori the LoLA solves the original problem. The LoLA defined by (44) solves this relaxed problem because $Q^L(c, p_L^*, p_H^*)$ satisfies (34) with equality.

■

B Material for Section 5

B.1 Semi-parametric identification of $\hat{D}$ and $\hat{O}$

We seek to recover the unobserved distribution of supplier quality conditional on cost c, that gives rise to the empirical distributions g_D and g_O in Figure 3. We take a guess-and-verify approach. In the next definition we guess a semi-parametric form of the distribution of supplier quality conditional on c; then we verify that the guess gives rise to the empirical distributions g_D and g_O, as it should.

Definition 2. ***(guess: distribution of supplier quality conditional on supplier cost)*** *For any $\xi \in [0,1]$ define:*

$$\hat{D}(c,\xi) = \begin{cases} \delta(c) & w.p.\ \xi \\ D & w.p.\ 1-\xi \end{cases}$$

$$\hat{O}(c,\xi) = \begin{cases} \omega(c) & w.p.\ \xi \\ O & w.p.\ 1-\xi, \end{cases}$$

where $\delta(c) = G_D^{-1}\left(\left[1-\hat{F}(c)\right]^N\right)$ *and* $\omega(c) = G_O^{-1}\left(\left[1-\hat{F}(c)\right]^N\right)$, *and D and O are the random variables with distributions depicted in Figure 3.*

Intuitively, $\hat{D}(c,\xi)$ is a random variable that represents the delay associated with a generic supplier with cost c. With probability ξ this delay is identically equal to the number $\delta(c)$; with complementary probability this delay is a random draw from the random variable D whose distribution is depicted in Figure 3, panel B. The same intuition holds for $\hat{O}(c,\xi)$. The functions $\delta(c)$ and $\omega(c)$ are specifically constructed so that the random variables D and O give rise to the "empirically correct marginals," in the following sense.

Lemma 5. ***(verify: $\hat{D}$ and $\hat{O}$ have the correct marginals)*** *Denote: $C_{(1)} = \min\{C_1, ..., C_N\}$. Then for any $\xi \in [0,1]$ we have: $\hat{D}\left(C_{(1)},\xi\right) \sim D$ and $\hat{O}\left(C_{(1)},\xi\right) \sim O$.*

Proof. We show the proof for the random variable D.

$$\begin{aligned}
\Pr\left(\delta\left(C_{(1)}\right) \le d\right) &= \Pr\left[G_D^{-1}\left(\left[1-\hat{F}\left(C_{(1)}\right)\right]^N\right) \le d\right] \\
&= \Pr\left[\left[1-\hat{F}\left(C_{(1)}\right)\right]^N \le G_D\left(d\right)\right] \\
&= \Pr\left[1-\left[G_D\left(d\right)\right]^{1/N} \le \hat{F}\left(C_{(1)}\right)\right] \\
&= \Pr\left[\hat{F}^{-1}\left(1-\left[G_D\left(d\right)\right]^{1/N}\right) \le C_{(1)}\right]
\end{aligned}$$

Since

$$\Pr\left(x \le C_{(1)}\right) = \left[1-\hat{F}\left(x\right)\right]^N,$$

then:

$$\begin{aligned}
\Pr\left(\delta\left(C_{(1)}\right) \le d\right) &= \left\{1-\hat{F}\left(\hat{F}^{-1}\left(1-\left[G_D\left(d\right)\right]^{1/N}\right)\right)\right\}^N \\
&= \left\{1-\left(1-\left[G_D\left(d\right)\right]^{1/N}\right)\right\}^N \\
&= \left\{G_D\left(d\right)^{1/N}\right\}^N \\
&= G_D\left(d\right).
\end{aligned}$$

The proof for the random variable O is virtually identical. ∎

This lemma proves that, if C is distributed according to $\hat{f}$, the delays and overruns of a bidder with cost c are drawn from $\hat{D}\left(c,\xi\right)$ and $\hat{O}\left(c,\xi\right)$, and there are N bidders, then the lowest bidder's marginal distributions of delays and overruns equals the observed marginal distributions of D and O from Figure 3. This property holds *for any value of the parameter* ξ. The parameter ξ encodes the correlation between cost and quality.

The calibrated buyer surplus function reads:

$$\begin{aligned}
\hat{v}(c,\xi) &= const - K\mathbb{E}\left[\hat{D}(c,\xi)+\hat{O}(c,\xi)\right] \\
&= const - (1-\xi)K\mathbb{E}\left[D+O\right] - \xi K\left[\delta(c)+\omega(c)\right] \\
&= const(\xi) - \xi K\left[\delta(c)+\omega(c)\right]. \qquad (45)
\end{aligned}$$

B.2 Calibration of $\hat{v}$

From expression (13), the calibrated buyer's payoff reads:

$$\hat{v}(c,\xi) = const(\xi) - \xi K \left[G_D^{-1}\left(\left[1-\hat{F}\left(c\right) \right]^N \right) + G_O^{-1}\left(\left[1-\hat{F}\left(c\right) \right]^N \right) \right]. \tag{46}$$

Our goal is to fully calibrate this function of (c,ξ). The constant $const(\xi)$ reads, from (45):

$$const(\xi) = const - (1-\xi)K\mathbb{E}\left[D+O\right]. \tag{47}$$

We set *const* large enough that the virtual valuation $\hat{w}$ is everywhere positive,[18] and K large enough that, as ξ varies between 0 and 1, the slope of the social welfare function (dashed red line in Figure 4) changes from positive to negative, while keeping at a magnitude that is reasonable. Specifically, we set $const = 1.0775 \times 10^6$ and $K = 2 \times 10^3$. With these values $\hat{w}$ is always positive (albeit barely so when c is small and ξ is large). Furthermore, the variation of the social surplus caused by a variation in supplier cost is reasonable. Indeed, given that the standard deviation of the distribution $\hat{f}$ (Figure 3, left-hand panel) equals 4.76×10^4, increasing the supplier's cost by one standard deviation around the mean (about one tick on the c-axis in Figure 4) yields variations in social surplus (dashed red line in Figure 4) that are plausible in magnitude, that is, not too large relative to average cost. With this choice of *const* and K, the social welfare evaluated at mean cost is of the same magnitude as the average cost for any ξ, which we view as a reassuring sanity check.

The three quantities $\hat{F}$, G_D, and G_O are given in Figure 3.

The number of bidders N is set equal to 7, the average number of bidders in the (first price) auctions studied by Decarolis (2014, 2016).

[18] This guarantees that the optimal LoLA does not involve a reserve price.

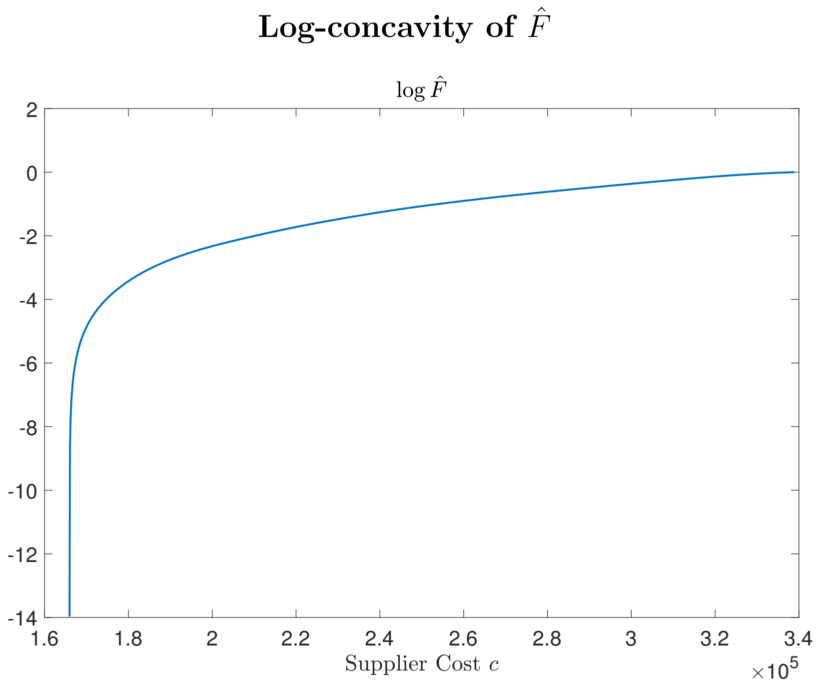

Figure 6: $\log(\hat{F})$ is concave.

C Software applications

C.1 Software 1

This software is a visually handy procedure realized in Matlab that does not require IBM ILOG CPLEX. The user specifies three inputs in an excel spreadsheet called “Input.xlsx”, as shown in Figure 7 (where input cells are colored in orange). There are four inputs: (i.) the minimum cost c_L (cell D21) and the maximum cost c_H (cell M21) used by the spreadsheet to automatically generate a linear cost grid with 10 nodes, (ii.) the 10 relative weights used to infer the cost distribution $f(c)$ (cells D20:M20), (iii.) the 10 values that represent the willingness to pay $v(c)$ (cells R20:AA20), and (iv.) the number of bidders N (cell R25).

Figure 7: The figure shows the inputs of the visual program that solves for the optimal LoLA among all LoLAs.

The Matlab script “FindOptimalLola.m” (which needs to be located in the same folder of the input file “Input.xlsx”) reads the 4 aforementioned inputs and calculates the virtual valuation function w. The script also re-samples all inputs on a grid with $T = 100$ nodes to increase the precision of the calculation. Given a grid $\{c_i\}_{i=1}^{T}$, the virtual valuation w is calculated as

$$\begin{cases} w_i = v_i - c_i, & i = 1 \\ w_i = v_i - c_i - (c_i - c_{i-1}) \cdot \frac{F_i}{f_i}, & \forall i > 1 \end{cases} \tag{48}$$

The result for w is showed to the user as in figure 9. The user is asked to check whether w is single-peaked in accordance to assumption 1.

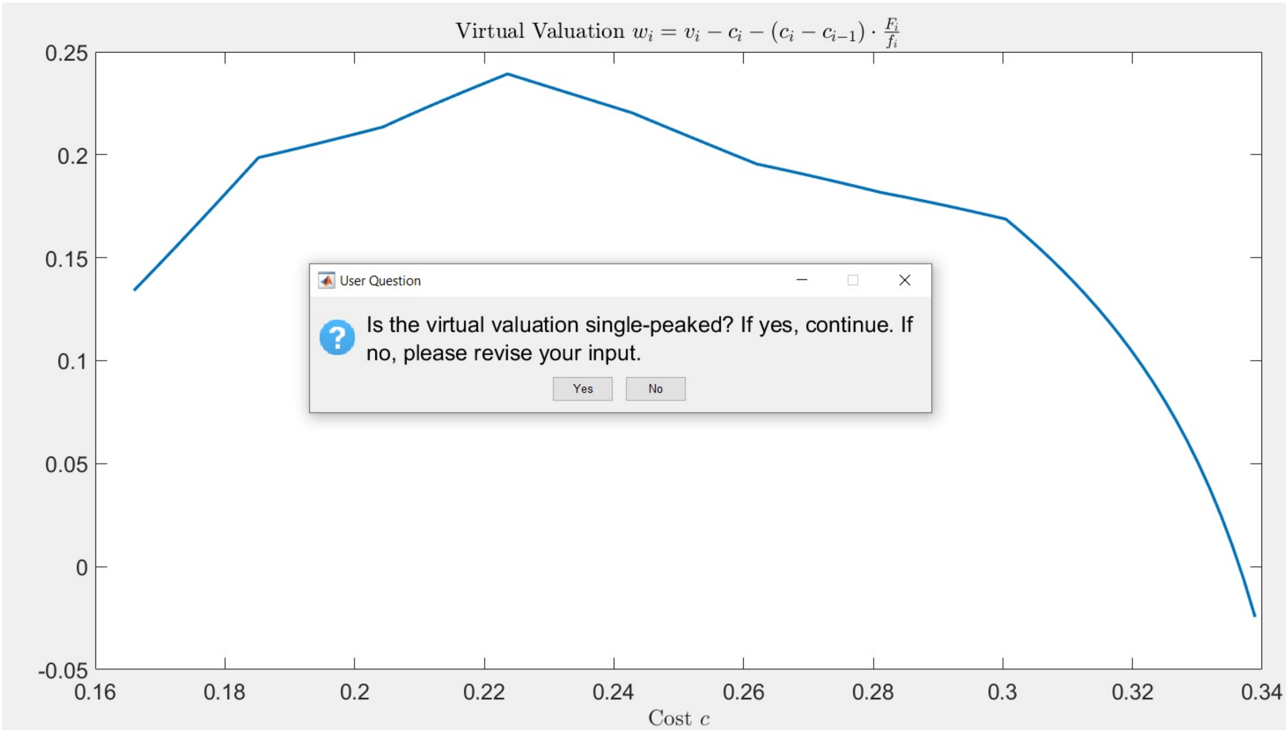

Figure 8: The figure shows w and it involves the user's participation by asking whether or not w is single-peaked.

If the user clicks "yes" the procedure continues, otherwise it stops as assumption 1 is violated. If "yes" is clicked, the procedure checks whether w has a root. If it does have a root, the software shows it in a new pop-up window as shown by figure 10. Hence, the software asks for the user's confirmation to set the root of w as a reservation price p_H.

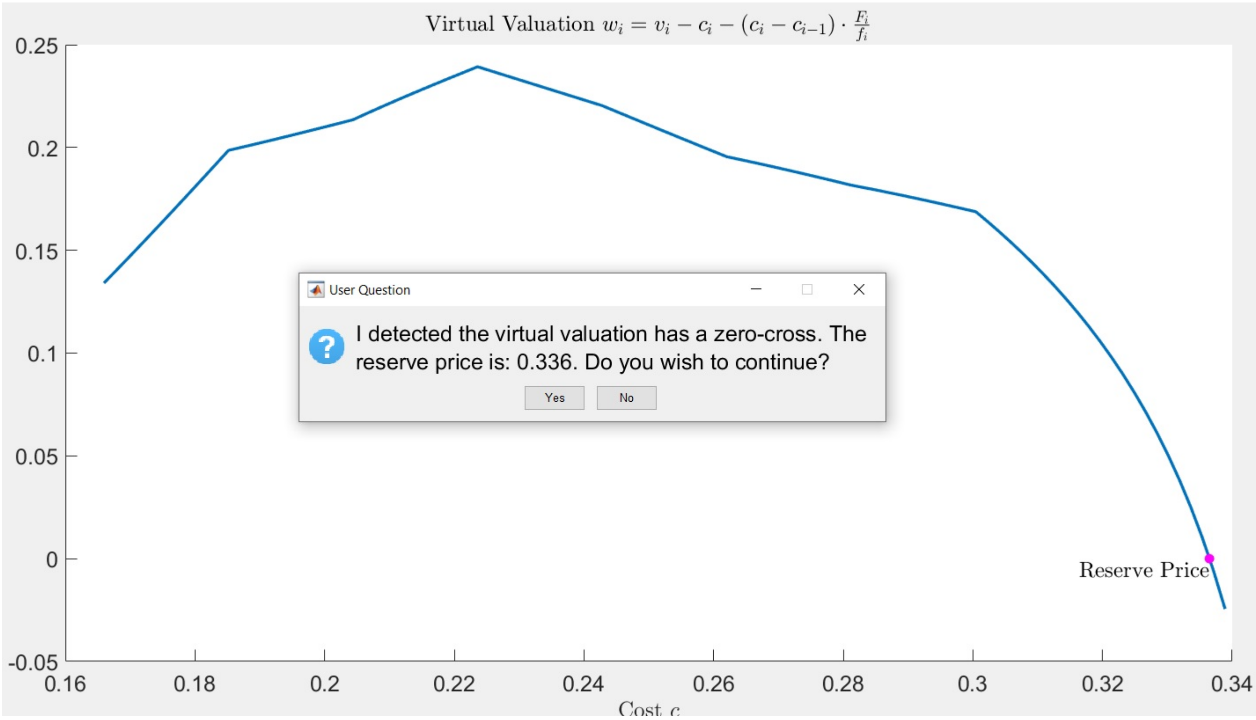

Figure 9: The figure shows the root of w calculated with a solver and using piece-wise linear interpolation on w. It involves the user's participation by acknowledging the root will be used as the reservation price.

Hence, the procedure iterates on all possible floor prices $\{p_{L,j}\}_{j=1}^{T}$ between c_L and c_H. For each floor price $p_{L,j}$, it calculates the associated buyer surplus $\sum_{i=1}^{T} w_i \cdot f_i \cdot Q_{i,j}$ and social surplus $\sum_{i=1}^{T}(v_i - c_i) \cdot f_i \cdot Q_{i,j}$. Note that $Q_{i,j} = Q(c_i, p_{L,j}, p_H)$ is calculated piece-wise as in equation 27 and it is function of the number of bidders N. The script terminates by showing the two resulting surpluses, optimal prices and benchmarks against the associated First Price Auction (FPA). The program shows results as reported in figure 10 and 11.

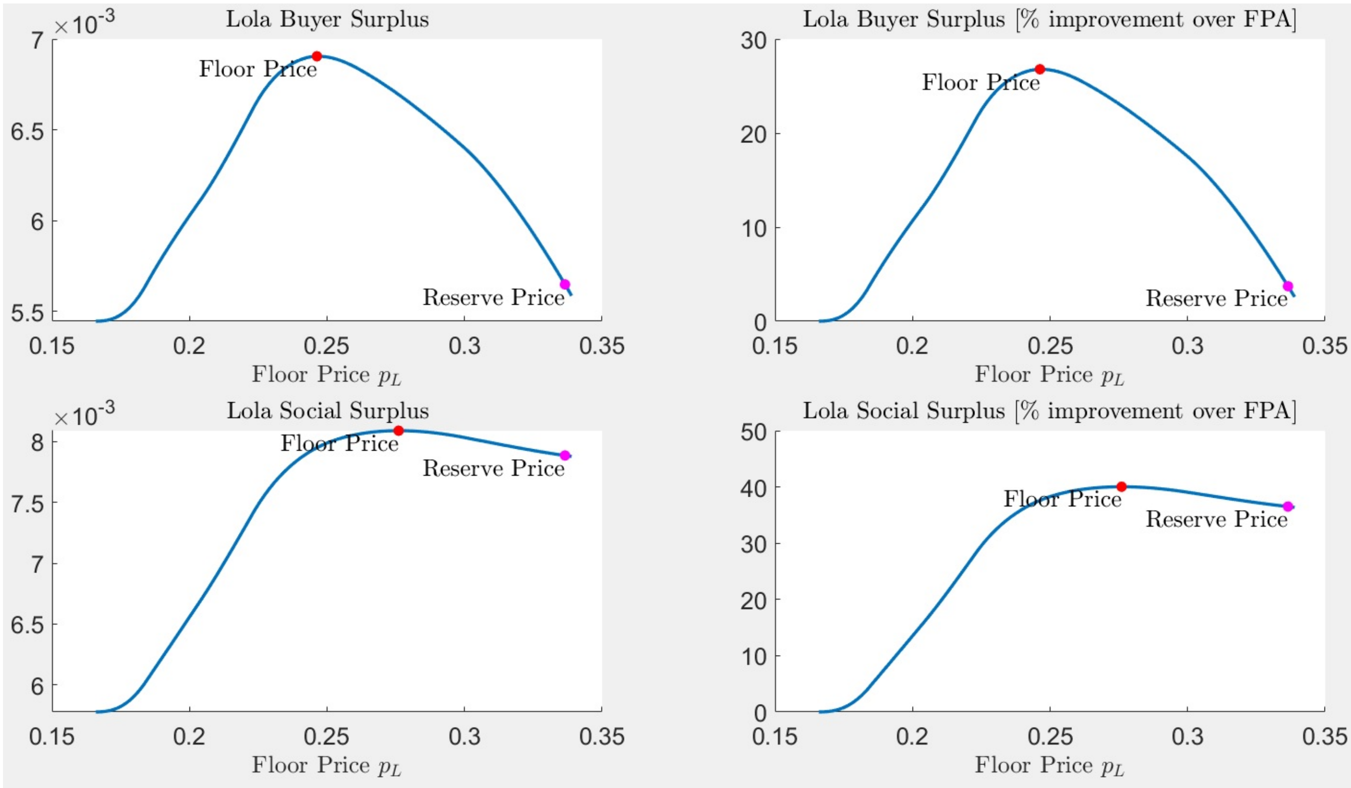

Figure 10: The figure shows the buyer surplus and social surplus in function of the floor price $p_{L,j}$. The points at which these functions are maximized correspond to the respective optimal LoLAs. In addition, the reservation price is also reported.

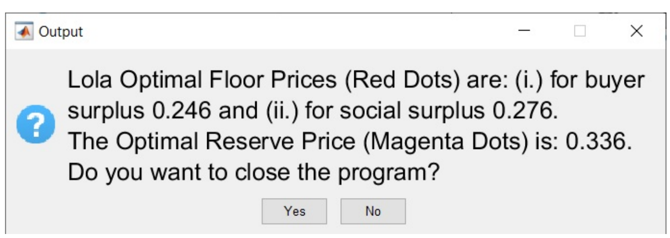

Figure 11: The figure shows the final report with the optimal floor and reservation prices.

C.2 Software 2

This software is realized in Matlab and IBM ILOG CPLEX. The entry point is "main.m". There are 5 inputs: (i.) the number of nodes T of the cost grid, (ii.) the minimum cost c_L, (iii.) the maximum cost c_H, (iv.) a vector of the willingness to pay $[v_1, \cdots, v_T]$, (v.) a vector of the cost distribution $[f_1, \cdots, f_T]$.

Given a distribution f, the virtual valuation is calculated as in (48). Then, the software passes all inputs to the script "CallCPLEX.m" in order to solve the linear program. This script generates two files: (i.) AMPL and (ii.) DAT.

The AMPL's file tells CPLEX how to generate the objective function and all constraints. In particular, it embeds the logic to generate: (i.) the demand constraints, (ii.) the non-negativity constraints, and (iii.) the monotonicity constraints.[19] The DAT's file specifies all numerical inputs.

Then, the program calls CPLEX to perform the high-scale optimization.

[19]CPLEX is preferable to Matlab because the optimization problem is large.