American Economic Review 2023, 113(6): 1505–1529
https://doi.org/10.1257/aer.20211437

Optimal Procurement with Quality Concerns†

By Giuseppe Lopomo, Nicola Persico, and Alessandro T. Villa*

Adverse selection in procurement arises when low-cost bidders are also low-quality suppliers. We propose a mechanism called LoLA (lowball lottery auction) which, under some conditions, maximizes any combination of buyer's and social surplus, subject to incentive compatibility, in the presence of adverse selection. The LoLA features a floor price, and a reserve price. The LoLA has a dominant strategy equilibrium that, under mild conditions, is unique. In a counterfactual analysis of Italian government auctions, we compute the gain that the government could have made, had it used the optimal procurement mechanism (a LoLA), relative to a first-price auction (the adopted format). (*JEL* D44, D82, H57, L14)

When the quality of a good or service is noncontractible, 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, it is common practice to reject abnormally low bids.[1] Some procurement rules deem bids to be "abnormally low" if they fall much below an engineering estimate of the work's cost.[2] Other rules, such as the "average bid auction" (ABA), disqualify bids that fall in extremely low (as well as extremely high) quantiles of the bid distribution. The rationale for disqualifying low bids is to weed out low-quality bidders.[3]

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).

* Lopomo: The Fuqua School of Business, Duke University (email: glopomo@duke.edu); Persico: Kellogg School of Management, Northwestern University, and NBER (email: nicola@nicolapersico.com); Villa: Federal Reserve Bank of Chicago (email: alessandro.villa@chi.frb.org). Sylvain Chassang was the coeditor for this article. We thank Jim Anton, and seminar participants at the EIEF and Duke University. Excellent research assistance by Ruofei Irene Guo, Zhaoyang Helen Xu, and Chenxi Xu is gratefully acknowledged. Special thanks to Francesco Decarolis for generously sharing the data on which Section V is based, and to Giancarlo Spagnolo for helpful feedback.

† Go to https://doi.org/10.1257/aer.20211437 to visit the article page for additional materials and author disclosure statements.

[1] The World Bank provides guidance for identifying abnormally low bids and deciding whether to accept or reject them. See World Bank (2016).

[2] Such is the case, for example, in the Korean procurement mechanism studied in Eun (2018).

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

A LoLA with floor price p_L and reserve price p_H is a (reverse) second-price sealed-bid auction in which bids below p_L and above p_H are not allowed, and ties are broken uniformly. When two or more bidders bid p_L, one of these bidders is randomly selected to supply the good and is paid p_L. In a LoLA, no bid is ever rejected for being too low: cheap suppliers are allowed to compete, but they are not allowed to bid too aggressively, and so they are not preferentially selected.

In a LoLA, the buyer effectively commits to pay no less than a (publicly announced) 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. Interestingly, floor prices are a feature of certain Medicare auctions[4] and of some Japanese procurement auctions.[5]

We show that, under mild regularity assumptions, the buyer's expected surplus is maximized by a LoLA among all interim incentive compatible (IC) and individually rational (IR) mechanisms. To our knowledge, this is the first time that a floor price emerges as part of an optimal selling mechanism.

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 surplus, 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 are 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. The theoretical and practical concerns with the ABA are documented by Albano, Bianchi, and Spagnolo (2006); Decarolis (2014, 2018); and Conley and Decarolis (2016).

Due to the adverse selection problem, in a standard first- or second-price auction both buyer surplus and social welfare may well decrease as the number of potential bidders increases. In the optimal LoLA, however, increasing the number of potential bidders improves both the buyer surplus and the social welfare. This difference

[4] Bids to supply the government with durable medical equipment, prosthetics, and orthotics, are limited by both ceilings and floors. See https://www.cms.gov/dmeposfeesched/downloads/dme10_c_summary.pdf. We thank a referee for pointing this out.

[5] See Chassang and Ortner (2019).

highlights the role that the floor price p_L^* plays in protecting the auctioneer from adverse selection.[6]

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 (Decarolis 2019), 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.[7] These applications compute the buyer-optimal procurement mechanisms in the presence of quality concerns, whether or not the optimal mechanism is a LoLA.

The two closest papers in the literature are Myerson (1981) and Manelli and Vincent (1995). When there is no lemons problem, first- and second-price auctions are both socially optimal and maximize the buyer's surplus (Myerson 1981). 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.[8]

The formal literature on (nonoptimal) procurement in the presence of quality concerns goes back to, at least, Dini, Pacini, and Valletti (2006) and Albano, Bianchi, and Spagnolo (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 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 the winner is chosen "randomly," suggesting that these auctions suffer from adverse selection.[9]

A sizable theoretical literature looks at settings where adverse selection arises endogenously through the winning bidder's strategic choice of performance (performing may mean paying one's bid or, in a procurement context, providing a suitable good

[6] Calzolari and Spagnolo (2006) show that, in a dynamic model where the provision of noncontractible quality is sustained by the threat of exclusion, the auctioneer may want to limit the number of bidders.

[7] See https://github.com/forket86/Software-1-Optimal-LoLA and https://github.com/forket86/Software-2-Optimal-Mechanism.

[8] 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.

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

or service). In this literature, after the winning bidder is selected, some uncertainty is realized that may lead the winning bidder to declare bankruptcy rather than perform. Because the option to declare bankruptcy is valuable, bidders who are more likely to take advantage of the option will bid more aggressively. Since more-aggressive bidders are less likely to perform, the auctioneer is exposed to adverse selection. This "strategic performance" paradigm blends moral hazard and adverse selection; our model, in contrast, may be regarded as a pure adverse selection model in the spirit of Manelli and Vincent (1995).

Within the "strategic performance" literature, Waehrer (1995) compares efficiency and revenue of first- and second-price auctions under different specifications for what happens after a default. Spulber (1990) analyzes first-price auctions, and shows that damages for nonperformance can play a key role in achieving allocational efficiency. Rhodes-Kropf and Viswanathan (2005) and Zheng (2001) study a setting where budget-constrained bidders borrow money in order to place their bid, and may later default on their loan; both papers study the efficiency of different contractual arrangements between bidders and lenders. Board (2007) compares first- and second-price auctions and finds that, depending on what happens to the assets of a bankrupt winner, one or the other auction format is preferred by the auctioneer. None of these papers seeks to identify the optimal auction mechanism.

Within this "strategic performance" paradigm, two papers adopt a mechanism design approach. Chillemi and Mezzetti (2014) study a complex design problem in which the mechanism determines not only the winning bidder, but—also—the type of damages to be paid in case of nonperformance. Closer to our approach, Burguet, Ganuza, and Hauk (2012) take as given what happens in case of nonperformance. In both papers, the optimal mechanism features pooling (the random choice of winner) only among types that underperform with probability zero—who are also the least-aggressive bidders. In a procurement auction, this type of pooling can be implemented with a price cap but not with a price floor: hence, as stated by Burguet, Ganuza, and Hauk (2012, fn. 25), "a price floor … is never optimal." By contrast, our mechanism leverages price floors to manage adverse selection.

Technically, our model differs from "strategic performance" models in the role that the winning bid plays in determining ex post performance. In the "strategic performance" literature, equilibrium performance depends on the winning bid's level: a higher winning bid is less likely to force the winner to declare bankruptcy. Thus, conditional on the winning bidder's type, reducing competition among bidders improves ex post performance. In our paper, by contrast, conditional on the winning bidder's type, there is no correlation between the winning bid's level and ex post performance. This lack of conditional correlation reflects the "pure adverse selection" nature of the model and is, admittedly, a stark feature. However, this feature does not preclude using our framework to model quality concerns arising from ex post performance. Indeed, in Section IVB we extend our framework to model ex post performance.[10]

Finally, Che and Kim (2010) compare auction formats that differ in the kind of legal tender that is allowed in the auction. The value of some legal tenders can

[10] In our extension the winning bid is, effectively, a "sunk cost" that does not affect performance.

depend on the bidder's unobservable type (e.g., if the tender is shares in entities that are managed by the bidder), which can create an adverse selection risk for the auctioneer. The value of cash is independent of the bidder's type. Che and Kim (2010) prove that the revenue-maximizing auction format uses cash, thereby completely eliminating adverse selection. Our setting is different in that bidders are restricted to bidding with cash, and yet an adverse selection problem exists. Furthermore, we do not allow mechanisms that eliminate adverse selection entirely, except for those that also eliminate competition entirely (random allocation).

This paper abstracts from both collusion and endogenous supplier entry. 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.[11] Their evidence suggests that introducing minimum prices causes potential suppliers to enter the auction, which helps destabilize cartels.

In sum, our first and main contribution relative to the literature is that we characterize the *optimal* procurement mechanism in the presence of pure adverse selection (i.e., abstracting from strategic performance considerations). 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 enough to the existing procurement formats 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, ancillary 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).[12]

The paper proceeds as follows. The next section contains a simple illustrative example. Section II lays out the model. Section III derives the optimal mechanism and some comparative static results. Section IV features several extensions. Section V analyzes the Italian procurement auctions. Section VI concludes.

I. 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

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

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

[11] Calzolari and Spagnolo (2006) also study repeated procurement in the presence of quality concerns.
[12] Reference to this software is provided in footnote 24.

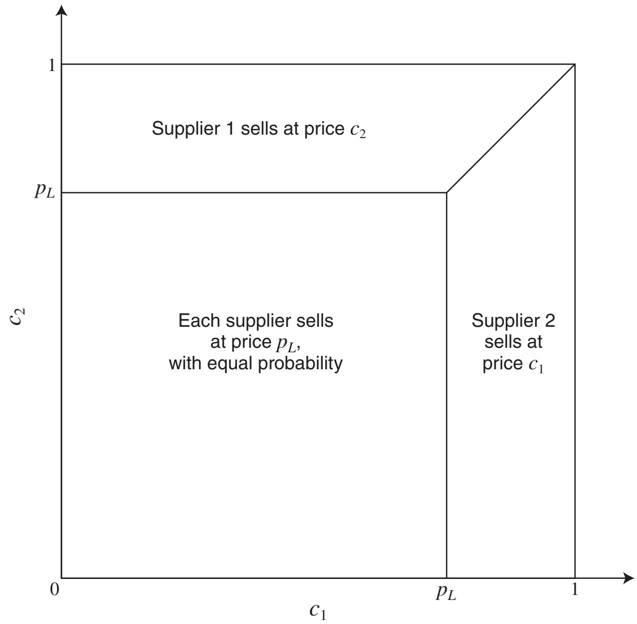

Figure 1. Outcome of the LoLA with Floor Price p_L

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 is paid 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.

Note that setting $p_L = 0$ yields the second-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

$$(2) \quad V(p_L) = \int_{p_L}^{1}\left\{\int_0^{c_2}\left[v(c_1) - c_2\right]dc_1\right\}dc_2 + \int_{p_L}^{1}\left\{\int_0^{c_1}\left[v(c_2) - c_1\right]dc_2\right\}dc_1$$
$$+ \int_0^{p_L}\int_0^{p_L}\left[\frac{1}{2}v(c_1) + \frac{1}{2}v(c_2) - p_L\right]dc_1 dc_2$$
$$= \frac{1}{3} + \frac{1}{3}\cdot(p_L)^3\cdot(1 - p_L).$$

The first two double integrals cover the upper- and right-trapezoid regions respectively, where bidder 2 (respectively, 1) bids more than its opponent and above the "floor price" p_L. In this case, the LoLA prescribes that the lowest bidder supplies

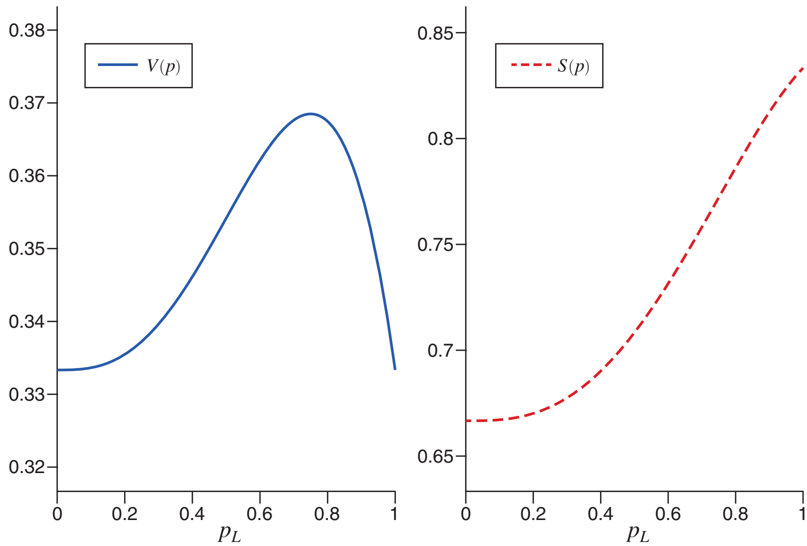

Figure 2. Expected Buyer Surplus V and Social Surplus S

Note: Expected buyer surplus V is maximal at $p_L^* = 3/4$ under a LoLA with floor price p_L.

the good and is paid the second-lowest bid c_2. The third double integral covers the inner-square region where both 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

$$(3) \qquad S(p_L) = 2\int_{p_L}^{1}\int_{0}^{c_2}\left[v(c_1) - c_1\right]dc_1dc_2 + \int_{0}^{p_L}\int_{0}^{p_L}\left\{\frac{1}{2}\left[v(c_1) - c_1\right] + \frac{1}{2}\left[v(c_2) - c_2\right]\right\}dc_1dc_2 = \frac{2}{3} + \frac{1}{3}\cdot\left(\frac{3}{2} - p_L\right)\cdot(p_L)^3.$$

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 second-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 percent 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, the 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). Intuitively, this is because a benevolent designer does not internalize the buyer's monetary savings from lowering p_L.

How would an ABA perform in this scenario? Decarolis (2014, 2018) and Conley and Decarolis (2016) have shown that the ABA is vulnerable to multiple coordination equilibria, some of which can be very unfavorable for the auctioneer. To illustrate their argument, allow for $N > 2$ bidders with the same uniform cost distribution. Define an (admittedly stylized) ABA as an auction where the lowest bidder wins, all bidders are paid their bid, but bids in the lowest or highest one-*N*th quantile of the bid distribution are discarded. Then, the strategy profile in which all bidders bid b is an equilibrium. To see this, observe that if bidder i deviates from b, its bid belongs either to the one-*N*th highest, or to the one-*N*th lowest quantile, and thus is automatically discarded. In this equilibrium, the buyer's expected surplus equals $E\left[v(c_i)\right] - b$, which can be made arbitrarily small by making b arbitrarily large. If, for example, $b = 1$ then the buyer's surplus equals 0.33, compared with about 0.37 that is attainable with the optimal LoLA with $p_L^* = 3/4$.

II. Model

A buyer with known type ξ seeks to procure an indivisible good from one of $N > 1$ potential suppliers. The suppliers' costs $c_1, \ldots, c_N$ 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 standard setting of Myerson (1981) in reverse, because the auctioneer buys rather than sells. 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 cost. This assumption does not imply that there are gains from trade for all cost realizations.

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 $F(c)/f(c)$ 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 versus social surplus.

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.

Assumption 1 will be used to establish the optimality of a LoLA (Theorem 1). A sufficient (but far from necessary) condition for this assumption to hold is that w be concave in c. If v is concave and F/f is convex, then w is concave. The ratio 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.[13]

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

$$q_i(c_i, c_{-i}), \quad m_i(c_i, c_{-i}) \tag{5}$$

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

III. 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.—

$$\max_{q,m} \int_{[c_L,c_H]^N} \left\{ \sum_{i=1}^{N} \left[\left(v(c_i,\xi) - (1-\beta) \cdot c_i \right) \cdot q_i(c_i,c_{-i}) - \beta \cdot m_i(c_i,c_{-i}) \right] \right\} \times \prod_{j=1}^{N} f(c_j) dc_j \tag{6}$$

[13] 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.

subject to, for all $i, c_i, c_i' \in [c_L, c_H], c_{-i} \in [c_L, c_H]^{N-1}$:

$$(7) \quad \sum_{i=1}^{N} q_i(c_i, c_{-i}) \leq 1,$$

$$(8) \quad q_i(c_i, c_{-i}) \geq 0,$$

$$(9) \quad \int_{[c_L,c_H]^{N-1}} \left[m_i(c_i, c_{-i}) - c_i \cdot q_i(c_i, c_{-i})\right] \prod_{j \neq i} f(c_j) dc_j \geq \int_{[c_L,c_H]^{N-1}} \left[m_i(c_i', c_{-i}) - c_i \cdot q_i(c_i', c_{-i})\right] \prod_{j \neq i} f(c_j) dc_j,$$

$$(10) \quad \int_{[c_L,c_H]^{N-1}} \left[m_i(c_i, c_{-i}) - c_i \cdot q_i(c_i, c_{-i})\right] \prod_{j \neq i} f(c_j) dc_j \geq 0.$$

The inequalities in (9) are the standard (interim) incentive-compatibility constraints. The inequalities in (10) are (interim) individual rationality constraints; these constraints capture the idea that suppliers are free not to bid.

In this section we prove that, for any ξ and any $\beta \in [0,1]$, the optimization problem above 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 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 a (reverse) second-price sealed-bid auction in which bids below p_L and above p_H are not allowed, and ties are broken uniformly.*

The next proposition is the main result of the paper.

THEOREM 1 (Optimality of the LoLA): *In a LoLA, it is a weakly dominant strategy for any supplier to not bid if its cost exceeds p_H, bid its cost when its cost is between p_H and p_L, and bid p_L otherwise. Furthermore, if Assumption 1 holds, the resulting equilibrium implements the solution to the optimization problem (6–10), provided that the reserve and floor prices are set to*

$$(11) \quad p_H^* = \sup\left\{c \in [c_L, c_H] \text{ such that } w(c; \xi, \beta) > 0\right\},$$

and

$$(12) \quad p_L^* = \max\left\{p \in [c_L, c_H] \text{ such that } w(p; \xi, \beta) \geq E\left[w(c; \xi, \beta) \mid c \leq p\right]\right\}.$$

PROOF.

See online Appendix A.

The reserve price p_H^* defined in (11) is the same as the reserve price in standard auctions: it is the type at which the virtual valuation w becomes negative. The inequality within curly brackets in equation (12) captures the trade-off that determines the optimal floor price p_L. If p_L is increased marginally, types slightly above p_L win with positive probability. These "marginal" types generate virtual surplus close to $w(p_L;\xi,\beta)$, which is the left-hand side of the inequality in (12). If, instead, p_L is not increased, then the marginal types are excluded and the virtual surplus generated is the average among all types below p_L, which is the right-hand side of the inequality in (12). The optimal floor price, if interior, equates the two: the equality reflects the optimal way to offer the same interim allocation to an interval of types below p_L.[14]

Equation (12) covers three different scenarios: the one in which standard auctions are optimal (Myerson 1981), the scenario in which random mechanisms are optimal (Manelli and Vincent 1995), and our intermediate scenario where a LoLA is optimal. If w is strictly decreasing in c, the inequality in equation (12) holds only for $p = c_L$, hence $p_L^* = c_L$ is the optimal floor price. This is the standard Myerson case in which the optimal mechanism is a standard first- or second-price auction. If, instead, w is strictly increasing in c, this inequality holds for all p in $[c_L, c_H]$. Then, the max operator in (12) uniquely selects $p_L^* = c_H$ as the optimal floor price: this is the random mechanism identified by Manelli and Vincent (1995). Finally, in the intermediate scenario where w peaks in the interior of $[c_L, c_H]$, the optimal floor price can be in the interior of $[c_L, c_H]$. To build intuition for this case, focus first on the case where w is negative in a neighborhood of c_H. In this case the optimal reserve price is interior, and strictly larger than the optimal floor price which is also interior. The first claim follows directly from equation (11). The second claim holds because the inequality in equation (12) must fail at any $p \geq p_H^*$, and must hold strictly at $p = c_L$; therefore, by continuity, the inequality must hold with equality at some point in the interior of $[c_L, p_H^*]$. This implies that the optimal floor price is interior and strictly lower than the optimal reserve price. This logic extends to the case where w is positive over its entire domain: in this case it is optimal not to use a reserve price; however, the optimal floor price may still exceed c_L.

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 sidestepping 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.

A number of comparative static results about p_H^* and p_L^* follow immediately from conditions (11) and (12).

PROPOSITION 1 (Comparative Statics on p_H^* and p_L^*):

(*i*) *Floor and reserve prices p_L^* and p_H^* are independent of the number of bidders.*

(*ii*) *The floor price is increasing in the severity of the lemons problem; i.e., p_L^* is nondecreasing in ξ for any β.*

[14] See, e.g., Section 6 in Bulow and Roberts (1989).

(*iii*) *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.

(*iv*) *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 (11) and (12) do not depend on N.

Part 2: Condition (12) is equivalent to

$$(13)\quad p_L^* = \max\left\{p \in [c_L, c_H] \text{ such that } \int_{c_L}^{p} w_c(c;\xi,\beta)\cdot F(c)\cdot dc \geq 0\right\}.$$

(To check this, integrate by parts the inequality in 13). Because $v_{c\xi} \geq 0$ by assumption, increasing ξ shifts the function w_c (at least weakly) upward (see equation (4)), and then condition (13) yields the result.

Part 3: Log concavity of F implies that the ratio $F(c)/f(c)$ is increasing in c; therefore increasing β shifts the function w_c down (see equation (4)), and then condition (13) yields the result.

Part 4: Increasing β shifts the function w downward (see equation (4)), and then condition (11) 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 Fs are log concave,[15] "typically," p_L^* will be nonincreasing in β. The economic intuition for this result was provided earlier at the end of Section I: the buyer prefers a lower p_L than the social planner because a benevolent designer does not internalize the buyer's monetary savings from lowering p_L.

Next, we show that increasing the number of potential suppliers N increases the weighted welfare generated by the optimal LoLA.

PROPOSITION 2 (Effect of the Number of Suppliers on Weighted Welfare): *Increasing the number of potential suppliers N increases the weighted welfare generated by the optimal LoLA.*

PROOF:

See online Appendix A.

[15] 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.

This result is not immediate because, as N increases, the adverse selection problem worsens. Indeed, if a naïve auctioneer used a standard first- or second-price auction rather than a LoLA, weighted welfare would *decrease* with N, at least for large N. To see this, assume that the optimal LoLA has an interior floor price. Then the function $w(\cdot)$ must be strictly increasing near c_L. In a standard first- or second-price auction, expected weighted welfare equals $E\left[w\left(c^{(1)}\right)\right]$, where $c^{(1)}$ denotes the lowest cost among all N suppliers. As N increases, the distribution of $c^{(1)}$ shifts toward the left and thus, eventually, $E\left[w\left(c^{(1)}\right)\right]$ must decrease with N. This observation highlights the role of the optimal floor price p_L^* in protecting the auctioneer from an adverse selection problem that worsens as N grows.

The next result concerns uniqueness. In what follows, "sincere bidding" means that all types between p_L and p_H bid their cost, and all types below p_L bid p_L.

PROPOSITION 3 (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 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 its 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.

IV. Extensions

A. *Reinterpreting v as Willingness to Pay for Expected Quality*

So far, we have assumed that the auctioneer's willingness to pay $v(c, \xi)$ is an increasing function of cost. This model can be thought of as the "reduced form" of a more complex model where a second dimension is present: the quality x_i provided by each supplier. We now spell out this model.

Assume that the auctioneer only cares about quality and, as before, each supplier cares only about its cost. Each supplier draws its quality and cost from a joint distribution $\Psi(c, x; \xi)$. Adverse selection arises when cost c and quality x are positively correlated. Quality, like cost, is noncontractible.[16]

[16] If quality, cost, or a combination of the two, were contractible, it would be beneficial for the auctioneer to use scoring rules.

In this setting, types are two-dimensional vectors (c,x). However, it turns out that there is no loss of generality in restricting attention to mechanisms $q_i(c_i,c_{-i})$, $m_i(c_i,c_{-i})$ which, as in (5), depend on c but not on x (see online Appendix B.1 for a proof of this statement). This implies that quality x only shows up in the objective function of the weighted welfare maximization problem (6), but not in any of the constraints (7)–(10). After integrating out x in the objective function, the buyer's willingness to pay becomes

$$v(c,\xi) = \int x d\Psi(x|c,\xi). \tag{14}$$

If x and c are stochastically affiliated, the expectation $v(c,\xi)$ is nondecreasing in c. Thus, the function $v(c,\xi)$, which is a primitive of the baseline model of Section II, can be interpreted in the present two-dimensional setting as the auctioneer's willingness to pay for the expected quality supplied by a bidder with cost c. In this interpretation, the parameter ξ modulates the correlation between cost and quality. Equation (14) will be used in Section VB to construct the auctioneer's willingness to pay function $v(c,\xi)$ based on the winning suppliers' performance in Italian procurement auctions.

B. *Reinterpreting Adverse Selection as Low Supplier Performance*

So far, the auctioneer's willingness to pay for supplier i's good $v(c_i)$ has been assumed to be an exogenous function of the supplier's cost c_i. In this section, we sketch out a setting in which the winning supplier's cost and the auctioneer's willingness to pay are determined endogenously by the winning bidder's performance. The positive correlation between the winning supplier's cost and the auctioneer's willingness to pay will emerge endogenously. For expositional simplicity, we restrict attention to a functional form example.

In what follows, there is no exogenously assigned cost c to each supplier before the auction. Rather, the cost of supplying the good is determined after the auction, by the winning supplier's choice of performance quality. Assume that, after the auction, the winning bidder must exert noncontractible effort $e \in [0,1]$ in order to fulfill its contractual obligations. The cost of effort is $\gamma(e,t) = 1 + t(e-1)$. The parameter $t \in [0,1]$ is supplier specific and captures heterogeneity across suppliers. Higher effort levels increase performance quality. The incentive to exert effort comes from a contractual specification that imposes a fine on the supplier if its performance is inadequate. The expected fine resulting from any effort level e is given by $\phi(e) = (e-1)^2$. The function $\phi(\cdot)$ is decreasing on $[0,1]$; i.e., higher effort results in a lower expected fine. Expected fines may be interpreted as a reduced-form proxy for any expected renegotiation costs incurred by the supplier, because renegotiation is more likely to occur when the winning supplier chooses a lower performance level.

The supplier seeks to minimize its overall cost inclusive of any fines, so the optimal effort level is given by

$$e^*(t) = \arg\min_e \gamma(e,t) + \phi(e) = 1 - \frac{t}{2}. \tag{15}$$

This expression represents the performance level of a supplier with type t. We assume that the auctioneer prefers higher effort, i.e., higher performance quality (and less renegotiation). Hence, expression (15) implies that suppliers with higher type t are worse from the auctioneer's viewpoint.

The resulting overall cost for the winning supplier is given by

$$\gamma\big(e^*(t),t\big) + \phi\big(e^*(t)\big) \;=\; 1 - \frac{t^2}{4}. \tag{16}$$

This expression is decreasing in t; i.e., a higher supplier type has a lower overall cost. Expected fines are given by

$$\phi\big(e^*(t)\big) \;=\; \frac{t^2}{4}. \tag{17}$$

If expected fines are interpreted as renegotiation costs, expression (17) indicates that renegotiation is more prevalent when a higher type wins the auction.

Expressions (15) and (16) imply that suppliers with a higher type t exert less effort (lower performance level) *and* incur a lower overall cost (inclusive of fines). Thus, as in our baseline model, the supplier's cost and the auctioneer's willingness to pay are positively correlated. However, unlike in the baseline model, here cost and quality are determined endogenously by the ex post behavior of the winning bidder.

C. *Descending LoLAs*

We have defined the LoLA as a sealed-bid auction. Alternatively, a LoLA can be implemented with a descending clock auction format with irrevocable exit. In this implementation, the price starts at p_H and is lowered continuously until either only one bidder is left, or the clock reaches p_L. In the first case the remaining bidder sells at the price where the clock stopped. In the second case, each remaining bidder sells at price p_L with equal probability.

D. *First-Price LoLAs*

In some procurement settings it may be desirable to use an auction format in which, unlike in the LoLA, the winner pays its bid. Next, we introduce an auction format with this property.

DEFINITION 1 (FPLoLA): *A first-price LoLA, or FPLoLA, with minimum bid b_L and reserve price $p_H \geq b_L$ is a (reverse) first-price sealed-bid auction in which bids below b_L and above p_H are not allowed, and ties are broken uniformly.*

In a first-price LoLA, the winning supplier always pays its bid. Individual rationality is guaranteed because suppliers are free not to bid.

The next proposition shows that the allocation induced by any LoLA, i.e., who wins the contract and how much each type expects to get paid, can be replicated by the symmetric equilibrium of a suitably designed FPLoLA.

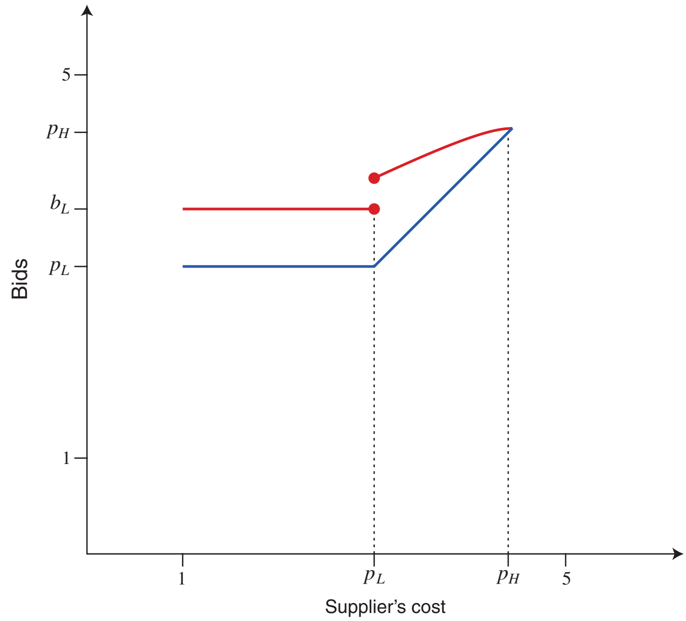

Figure 3. Equilibrium Strategies in a LoLA and Its Equivalent FPLoLA

Notes: The blue line is the equilibrium bidding strategy in a LoLA with two bidders, costs drawn from the uniform distribution on $[1,5]$, and $p_L = 3, p_H = 4.4$. The red line is the equilibrium bidding strategy in the equivalent first-price LoLA; bidders with cost larger than p_H choose not to bid.

PROPOSITION 4 (Implementation via an Equivalent FPLoLA): *The allocation induced by the sincere equilibrium in a LoLA with any reserve price and floor price p_L, can be implemented by the symmetric equilibrium of an "equivalent FPLoLA" with the same reserve price and a suitably chosen minimum bid b_L.*

PROOF:

See Lemma 6 in the online Appendix for a complete characterization of the equivalent FPLoLA and its equilibrium.

Figure 3 compares the equilibrium bidding strategies in a LoLA and its equivalent FPLoLA, in an environment with two bidders and costs drawn from the uniform distribution on $[1,5]$. Consider a LoLA with floor price p_L and reserve price p_H. The red curve in Figure 3 represents the equilibrium bidding function in its equivalent FPLoLA (this is the strategy $\beta^{f\mathbb{L}}$ given in Lemma 6). In this equilibrium, types $c_i > p_L$ bid as in a (reverse) first-price auction with no minimum bid, and types $c_i \leq p_L$ bid the minimum bid b_L. Type p_L is indifferent between bidding on the increasing portion of the red curve and bidding the minimum bid b_L. The minimum bid b_L is carefully chosen to ensure that the discontinuity in the bidding function arises precisely at type p_L: this property must hold for the FPLoLA to be equivalent to the LoLA.

Figure 3 also displays the equilibrium bidding strategy in the LoLA (blue line). All types between p_L and p_H bid their cost, and all types below p_L bid p_L. Per the

LoLA rules, any bidder who wins with a bid of p_L is paid at least p_L, and sometimes more; in expectation, such a bidder is paid an amount that equals exactly b_L, the minimum bid in the equivalent FPLoLA. The blue line is uniformly below the red line, meaning that bidders in a LoLA bid more aggressively than in the equivalent FPLoLA.

Implementing a given allocation via a LoLA is less informationally demanding than implementing it through an equivalent FPLoLA. Indeed, in a LoLA all suppliers have a dominant strategy and so they do not need to concern themselves with the behavior of others. Furthermore, the optimal floor price p_L^* is independent of the number of bidders N (see expression (12)). In contrast, the corresponding minimum bid in the FPLoLA depends on N (see online Appendix expression (53)).

E. *Asymmetric Bidders*

In our setting, bidders may be asymmetric in two dimensions: in the parameter ξ and in the cost distribution f. We were unable to obtain an analytic solution comparable to Theorem 1 for the asymmetric case.[17] However, we used our software applications to compute the optimal mechanism in asymmetric environments close to the symmetric one studied in Section I. The main insight from the numerical analysis is that a key feature of optimal LoLAs is robust to the introduction of asymmetries across bidders. This feature is that, when all bidders have relatively high cost, the auctioneer can afford to induce price competition because the adverse selection problem is mild. However, when multiple bidders have relatively low costs, the auctioneer prefers to suppress price competition in order to avoid buying from the lowest cost (hence, lowest quality) bidder.

In our numerical analysis, supplier 1's cost x is drawn from a distribution with density $f_1(x;a) = a \cdot (x - 1/2) + 1$ on $[0,1]$. Supplier 2's cost y is drawn independently from the uniform distribution on $[0,1]$. The buyer's willingness to pay for each supplier's good is, respectively,

$$v_1(x) = v_0 - 4 \cdot \left(\frac{1}{2}x^2 - x + \frac{1}{3}\right) \text{ and } v_2(y;\xi_2) = v_0 - \xi_2 \cdot \left(\frac{1}{2}y^2 - y + \frac{1}{3}\right).$$

The parameter ξ_2 modulates the severity of supplier 2's adverse selection: if ξ_2 equals zero, there is no adverse selection. When f is uniform, the functional form of $v_2(y;\xi_2)$ guarantees that the ex ante expected gains from trade with supplier 2 are independent of ξ_2.[18] Setting $a = 0$, $v_0 = 4/3$, and $\xi_2 = 4$ yields the symmetric example of Section I. Here we set $v_0 = 2$ to guarantee that, when we introduce asymmetries, both virtual valuations remain positive.

When a is fixed at zero and ξ_2 varies in the interval $(2.5, 4)$, the optimal mechanism is qualitatively illustrated in Figure 4 panel A. When both suppliers' costs are relatively high, supplier 1 wins more often than in the symmetric case depicted in Figure 1. Conversely, when both suppliers' costs are relatively low, supplier 2 wins

[17] The proof of Theorem 1 relies on identifying the analytic expression of the dual solution, i.e., the shadow prices of the weighted welfare problem. In the asymmetric case the dual solution is not unique, and this multiplicity makes it more difficult to identify the analytic expression of any dual solution.

[18] For the same reason, the gains from trade with supplier 1 are the same as with supplier 2.

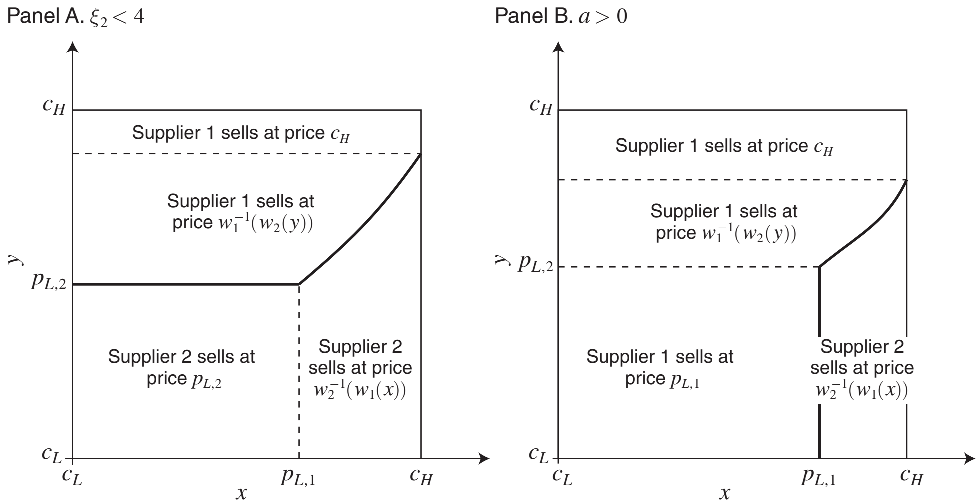

Figure 4. Optimal Auctions in Asymmetric Settings

Notes: In panel A, the ex ante expected gains from trading with either supplier are the same, but quality concerns are less severe for supplier 2. In panel B, supplier 1's cost distribution is higher than (i.e., stochastically dominates) its opponent's.

more often. This property reflects the fact that the optimal mechanism rewards suppliers with a relatively high virtual valuation. Because $\xi_2 < 4$, supplier 2's virtual valuation (refer to expression (4)) exceeds its opponent's if both costs are small; conversely, if both costs are high, supplier 1's virtual valuation is higher.[19] Note, also, that when both suppliers' costs are low, supplier 2 wins for sure. To understand this property, observe that in the symmetric case the auctioneer was indifferent between the two suppliers, and so was willing to randomize between the two; here, instead, supplier 2 is strictly preferable.

When ξ_2 is fixed at 4 and a varies in the interval $(0, 0.5)$, the optimal mechanism is qualitatively illustrated in Figure 4 panel B. Supplier 1 wins more often than in the symmetric case depicted in Figure 1. This property results from a standard property that does not depend on quality concerns: for any supplier, lower-cost types command more information rents. When $a > 0$ low-cost types are less likely for supplier 1 than for supplier 2, and thus it is better for the auctioneer to buy from supplier 1. This causes the optimal mechanism to favor bidder 1. Note, also, that when both suppliers have a relatively low cost, supplier 1 wins for sure. To understand this feature, observe that in the symmetric case the auctioneer was indifferent between the two suppliers, and so was willing to randomize between the two. Now, supplier 1 is strictly more attractive than supplier 2.

[19] This can be verified by plugging the expressions for v_1 and v_2 in (4) and setting x close to y.

V. Illustrative Application: Optimal Procurement Mechanisms for the 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 (Decarolis 2019),[20] 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 (18)), 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.[21]

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.

A. *The Available Data*

The available data are depicted in Figure 5. Panel A shows the estimated distribution of bidder costs $\hat{f}$, which was structurally estimated by Decarolis (2018) and corresponds to our $f(c)$.[22] 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.[23] The figure indicates that, in most cases, the government suffers a delay, a cost overrun, or both.[24]

B. *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,

$$v(c,\xi) = const - KE\big[D(c,\xi) + O(c,\xi)\big], \tag{18}$$

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

[21] To compute optimal mechanisms, this section leverages two software applications that we have created and made publicly available. Taking as input the bidders' cost distribution F and the auctioneer's valuation function $v(c,\xi)$, these applications yield the optimal procurement mechanism (5), even when Assumption 1 is violated and, so, the optimal mechanism may not be a LoLA. Applications downloadable from https://github.com/forket86/Software-1-Optimal-LoLA and https://github.com/forket86/Software-2-Optimal-Mechanism.

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

$$c_i = y + z_i,$$

where the z_is are idiosyncratic and privately known cost components, and y is an auction-specific and commonly known scalar. Decarolis (2018) estimates that $z_1, \ldots, z_N$ are i.i.d. draws from a random variable Z whose density is depicted in Figure 5 panel A. In what follows we assume, without loss of generality, that $y = 0$, which allows us to interpret z_is as c_is.

[23] 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.

[24] 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*.

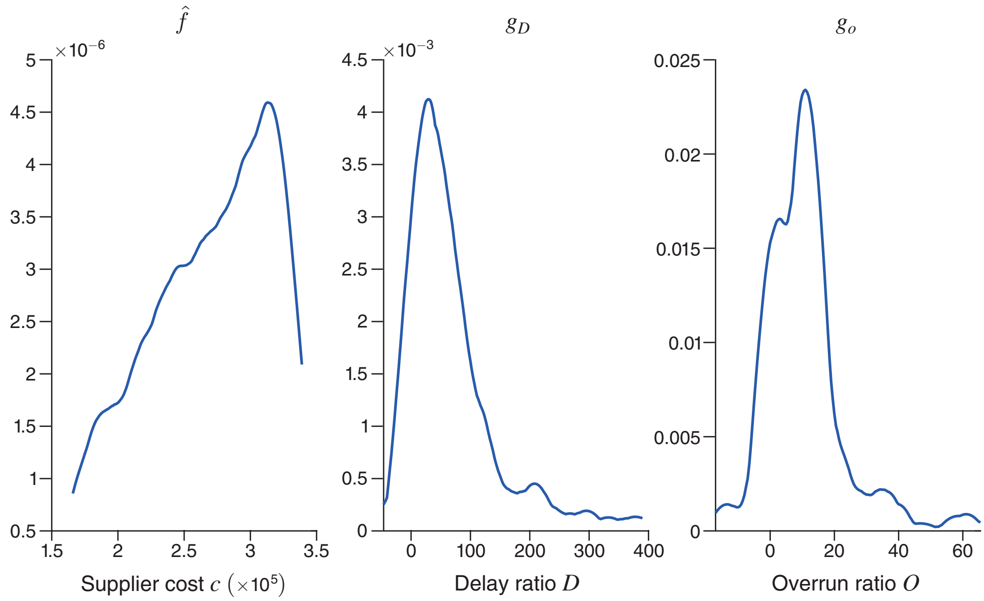

Figure 5. Distributions of Cost and Quality Measures

Notes: The left-hand panel depicts the estimated probability density function $\hat{f}_Z$ of the idiosyncratic cost component Z (unit is 10^5 euros) from Decarolis's (2018) assumed cost structure $c_i = y + z_i$, where the z_is are i.i.d. 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 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 bandwidths used are 11000, 18.15 and 3.0071 respectively. Data generously provided by Francesco Decarolis.

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. We use K as a positive scaling parameter whose value will be calibrated later.[25]

The parameter ξ in expression (18) 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(18), 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 (19) below.

The distributions of the random variables $D(c,\xi)$ and $O(c,\xi)$ are as yet unspecified. We calibrate them semiparametrically by requiring that, given that $c \sim \hat{f}$, their distributions for any given ξ coincide with the empirical marginal distributions g_D

[25] There is no difficulty in making expression (18) more complex. For example, one could premultiply $D(c,\xi)$ and $O(c,\xi)$ by positive constants, and the analysis would be essentially unchanged.

and g_O depicted in Figure 5.[26] Definition 3 in online Appendix C 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 (18) yields the following expression for the calibrated buyer payoff function:

$$\hat{v}(c,\xi) = const - KE\left[\hat{D}(c,\xi) + \hat{O}(c,\xi)\right] \tag{19}$$

$$= const(\xi) - \xi K\left[\delta(c) + \omega(c)\right],$$

where $const(\xi)$ is independent of c and, from Definition 3, we have

$$\delta(c) = G_D^{-1}\left(\left[1 - \hat{F}(c)\right]^N\right),$$

$$\omega(c) = G_O^{-1}\left(\left[1 - \hat{F}(c)\right]^N\right)$$

(refer to online Appendix C.2 for the computations). Expression (19) is the calibrated buyer's 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 5; and the parameters N, K are assigned numerical values as described in online Appendix C.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, the 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 assumptions imposed on page 3. Indeed, it can be checked from expression (19) that $\hat{v}_{c\xi} \geq 0$. Furthermore, we can (and will) make $const(\xi)$ in expression (19) large enough that $\hat{v}(c_L,\xi) \geq c_L$ for all $\xi \in [0,1]$.

C. *Buyer-Optimal and Socially Optimal Mechanisms Are LoLAs*

We compute the calibrated virtual valuation function

$$\hat{w}(c;\xi,\beta) \equiv \hat{v}(c;\xi) - c - \beta\frac{\hat{F}(c)}{\hat{f}(c)}, \tag{20}$$

by substituting $\hat{v}$ from (19) and $\hat{F}$ from Figure 5 into the expression for the virtual valuation (4). We set $const(\xi)$ large enough that the virtual valuation (20) 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.[27]

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

[27] By setting *const* large enough in expression (18), $const(\xi)$ can be made arbitrarily large: refer to online Appendix C.2 for information about the calibration.

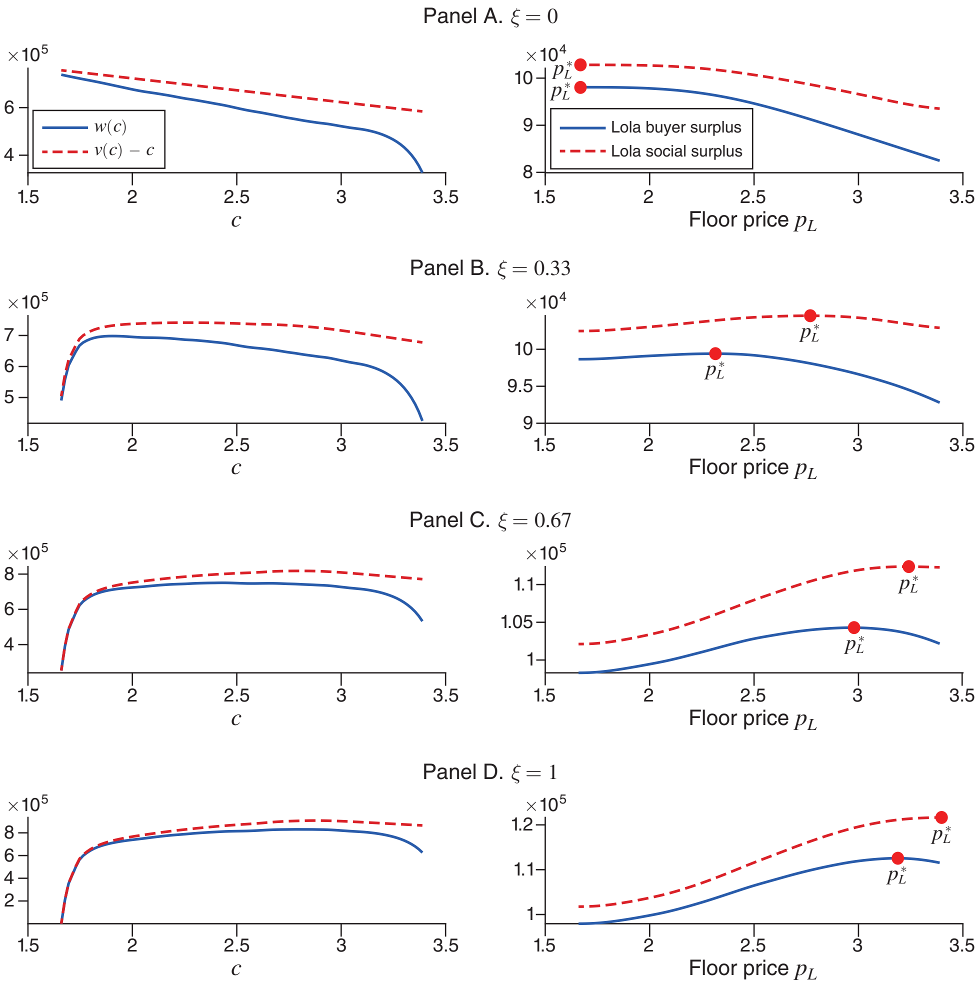

Figure 6. Optimal Mechanisms with Varying Degrees of Quality Concerns

Notes: 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 column, 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 graph). With maximal quality concerns, the socially optimal LoLA reduces to the random allocation mechanism ($\xi = 1$, bottom-right graph).

Each of the left-hand graphs in Figure 6 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 graphs, the buyer's virtual valuation and the gains from trade happen to be quasiconcave 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.

The right-hand graphs of Figure 6 are calibrated counterparts to Figure 2. Each right-hand graph displays the expected buyer (solid blue line) and social (dashed red

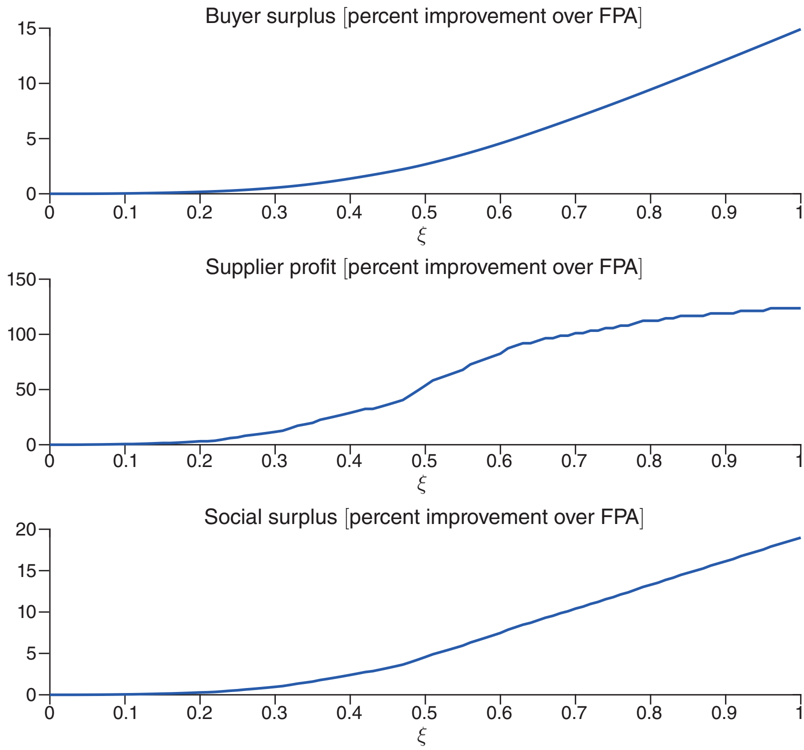

Figure 7. Performance Improvement of Optimal LoLA over First-Price (or Second-Price) Auction

line) surplus in a LoLA with floor price p_L. The optimal floor prices are determined by equation (12) after setting β equal to one or zero: accordingly, they maximize the expected (buyer or social) surplus, as shown in Figure 6. Within each right-hand graph, 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 8 in the online Appendix).

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 graphs, 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 right-hand graphs. For low values of ξ, the buyer-optimal 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.

D. *Performance of the Buyer-Optimal Mechanism versus First-Price Auction*

Figure 7 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.[28] 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 percent 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 percent range, which are nontrivial from a policy perspective.

In Section IVD we showed that the optimal LoLA can also be implemented via a first-price auction with an appropriately chosen minimum bid b_L. Within the parametric setting that gives rise to Figure 6, we computed the minimum bids b_L corresponding to the buyer-optimal floor prices p_L (these p_Ls are marked by the red dots on the blue curves in the figure). We know from the theory that $b_L \geq p_L$. Using online Appendix expression (53) we find that b_L is up to 24 percent higher than p_L when $\xi = 0$; the two thresholds both converge to c_H (and therefore to each other) as ξ increases toward 1.[29]

VI. 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 surplus 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

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

Bagnoli, Mark, and Ted Bergstrom. 2005. "Log-Concave Probability and Its Applications." *Economic Theory* 26 (2): 445–69.

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

[29] The code used to compute b_L is available at https://www.alessandrotenzinvilla.com/research.html.

Blume, Andreas, and Paul Heidhues. 2004. "All Equilibria of the Vickrey Auction." *Journal of Economic Theory* 114 (1): 170–77.

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

Bulow, Jeremy, and John Roberts. 1989. "The Simple Economics of Optimal Auctions." *Journal of Political Economy* 97 (5): 1060–90.

Burguet, Roberto, Juan-José Ganuza, and Esther Hauk. 2012. "Limited Liability and Mechanism Design in Procurement." *Games and Economic Behavior* 76 (1): 15–25.

Calzolari, Giacomo, and Giancarlo Spagnolo. 2006. "Reputation and Collusion in Procurement." Unpublished.

Chassang, Sylvain, and Juan Ortner. 2019. "Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement." *Journal of Political Economy* 127 (5): 2269–2300.

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

Chillemi, Ottorino, and Claudio Mezzetti. 2014. "Optimal Procurement Mechanisms: Bidding on Price and Damages for Breach." *Economic Theory* 55 (2): 335–55.

Conley, Timothy G., and Francesco Decarolis. 2016. "Detecting Bidders Groups in Collusive Auctions." *American Economic Journal: Microeconomics* 8 (2): 1–38.

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

Decarolis, Francesco. 2018. "Comparing Public Procurement Auctions." *International Economic Review* 59 (2): 391–419.

Decarolis, Francesco. 2019. "Dati Aste" Unpublished data. Obtained by email on May 10, 2019.

Decarolis, Francesco, and Michael Klein. 2012. "Too Good to Be True Auctions." Unpublished.

Dini, Federico, Riccardo Pacini, and Tommaso Valletti. 2006. "Scoring Rules." In *Handbook of Procurement*, edited by Nicola Dimitri, Gustavo Piga, and Giancarlo Spagnolo, 293–321. Cambridge, MA: Cambridge University Press.

Eun, Dong-Jae. 2018. "Reducing Ex-post Adjustment by Bid Screening Ex-ante: Evidence from Procurement Auctions for Paving Work." Unpublished.

Lopomo, Giuseppe. 2001. "Optimality and Robustness of the English Auction." *Games and Economic Behavior* 36 (2): 219–40.

Lopomo, Giuseppe, Nicola Persico, and Alessandro T. Villa. 2023. "Replication Data for: Optimal Procurement with Quality Concerns." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E182801V1.

Manelli, Alejandro M., and Daniel R. Vincent. 1995. "Optimal Procurement Mechanisms." *Econometrica* 63 (3): 591–620.

Manelli, Alejandro M., and Daniel R. Vincent. 2004. "Duality in Procurement Design." *Journal of Mathematical Economics* 40 (3-4): 411–28.

Maskin. Eric, and John Riley. 1984. "Optimal Auctions with Risk Averse Buyers." *Econometrica* 52 (6): 1473–1518.

Myerson, Roger B. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (1): 58–73.

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

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

Waehrer, Keith. 1995. "A Model of Auction Contracts with Liquidated Damages." *Journal of Economic Theory* 67 (2): 531–55.

World Bank. 2016. *Abnormally Low Bids and Proposals: Guide to the Identification and Treatment of Abnormally Low Bids and Proposals*. Washington, DC: World Bank.

Zheng, Charles Z. 2001. "High Bids and Broke Winners." *Journal of Economic Theory* 100 (1): 129-171.