NONCONFIDENTIAL // EXTERNAL

Federal Reserve Bank of Chicago

Optimal Procurement with Quality Concerns

Giuseppe Lopomo, Nicola Persico, and Alessandro T. Villa

REVISED
August 3, 2022
WP 2022-21

https://doi.org/10.21033/wp-2022-21

Optimal Procurement With Quality Concerns*

Giuseppe Lopomo
Duke[†]

Nicola Persico
Northwestern & NBER[‡]

Alessandro T. Villa
FRBC[§]

August 3, 2022

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 any combination of buyer's and 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.

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 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 6 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, 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). 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

[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.

[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.

procurement auctions.[5]

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 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 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. The theoretical and practical concerns with the ABA are documented by Albano et al. (2016), 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 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 ex-

[5] See Chassang and Ortner (2019).

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

periment 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.[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 (1982) 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 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.[8]

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 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

[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.

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 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 (2007) compares efficiency and revenue of first- and second-price auctions under different specifications for what happens after a default. Spulber (2000) analyzes first-price auctions, and shows that damages for non-performance can play a key role in achieving allocational efficiency. Rhodes-Kropf and Viswanathan (2000) 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 non-performance. Closer to our approach, Burguet et al. (2012) take as given what happens in case of non-performance. 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 et al. (2012, fn. 25), "a price floor ... is never optimal." By contrast, our mechanism leverages price floors to manage adverse selection.

[9]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).

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 5.2 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 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

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

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

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 3 lays out the model. Section 4 derives the optimal mechanism and some comparative static results. Section 5 features several extensions. Section 6 analyzes the Italian procurement auctions. 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 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.

[12]Reference to this software is provided in footnote 21.

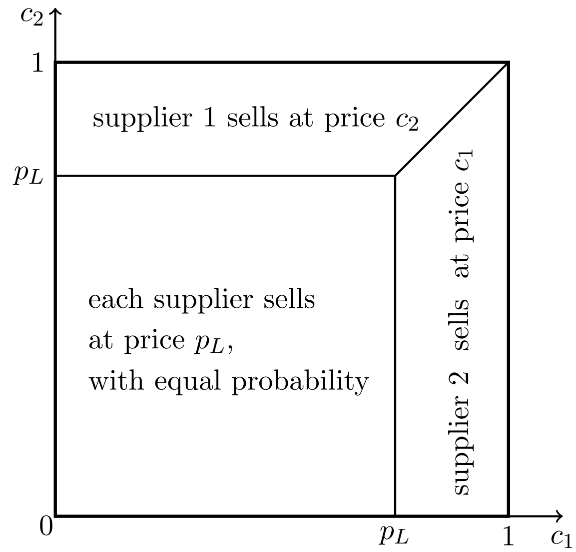

Figure 1: Outcome of the LoLA with floor price p_L

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:

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

The first two double integrals cover the upper- and right-trapezoid regions respectively, where bidder 2 (resp., 1) bids more than its 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. 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

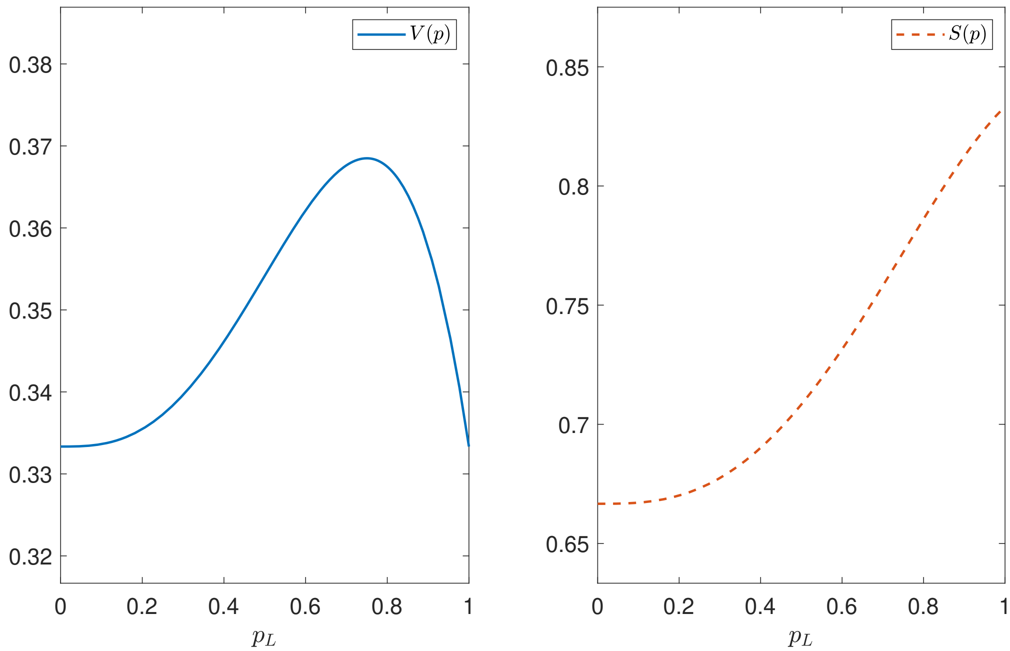

Figure 2: Expected buyer surplus V and social surplus S. V is maximal at ${p_L}^* = 3/4$ under a LoLA with floor price p_L.

substituting for $v(\cdot)$ from (1) and solving the integrals.

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

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

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 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% 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 $1/N$-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 $1/N$-th highest, or to the $1/N$-th lowest quantile, and thus is automatically discarded. In this equilibrium, the buyer's expected surplus equals $\mathbb{E}[v(c_i)] - b$, 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 $p_L^* = 3/4$.

3 Model

A buyer with known type ξ seeks to procure an indivisible good from one of $N > 1$ potential suppliers. The suppliers' costs $c_1, ..., c_N$ 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 $\frac{F(c)}{f(c)}$ 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.

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 $\frac{F}{f}$ 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.[13]

[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

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

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

that, for each 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.

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:

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.

Weighted welfare maximization problem

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

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

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

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

$$\begin{aligned} &\int_{[c_L,c_H]^{N-1}} [m_i(c_i,c_{-i}) - c_i \cdot q_i(c_i,c_{-i})] \prod_{j\neq i} f(c_j)\, dc_j \\ &\qquad \geq \int_{[c_L,c_H]^{N-1}} [m_i(c_i',c_{-i}) - c_i \cdot q_i(c_i',c_{-i})] \prod_{j\neq i} f(c_j)\, dc_j \end{aligned} \tag{9}$$

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

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 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 is formally defined next.

Lowball lottery auction (LoLA): formal definition

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.

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. Furthermore, if Assumption 1 holds, the resulting equilibrium implements the solution to optimization problem (6-10) provided that the reserve and floor prices are set to:*

$$p_H{}^* = \sup\{c \in [c_L, c_H] \;\; s.t. \; w(c; \xi, \beta) > 0\}, \tag{11}$$

and

$$p_L{}^* = \max\left\{p \in [c_L, c_H] \;\; s.t. \; w(p; \xi, \beta) \geq \mathbb{E}\left[w(c; \xi, \beta) \mid c \leq p\right]\right\}. \tag{12}$$

Proof. See 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 tradeoff 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 mechanims 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

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

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 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.

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 ${p_L}^*$)**.**

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 non-decreasing 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 (11) and (12) do not depend on N.

Part 2 Condition (12) is equivalent to:

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

(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 eq. 4), and then condition (13) 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), and then condition (13) yields the result.

Part 4 Increasing β shifts the function w downward (see eq. 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 F's are log-concave,[15] "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 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 Appendix A. ■

This result is not immediate because, as N increases, the adverse selection problem worsens. Indeed, if a naive 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

[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.

welfare equals $\mathbb{E}[w(c^{(1)})]$, 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, $\mathbb{E}[w(c^{(1)})]$ 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.

5 Extensions

5.1 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 non-contractible.[16]

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 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

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

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 3, 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 6.2 to construct the auctioneer's willingness to pay function $v(c, \xi)$ based on the winning suppliers' performance in Italian procurement auctions.

5.2 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 non-contractible effort $e \in [0, 1]$ in order to fulfill its contrac-

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

tual 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

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

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

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

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

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

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.

5.3 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.

5.4 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.

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 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.

Equilibrium strategies in a LoLA and its equivalent FPLoLA

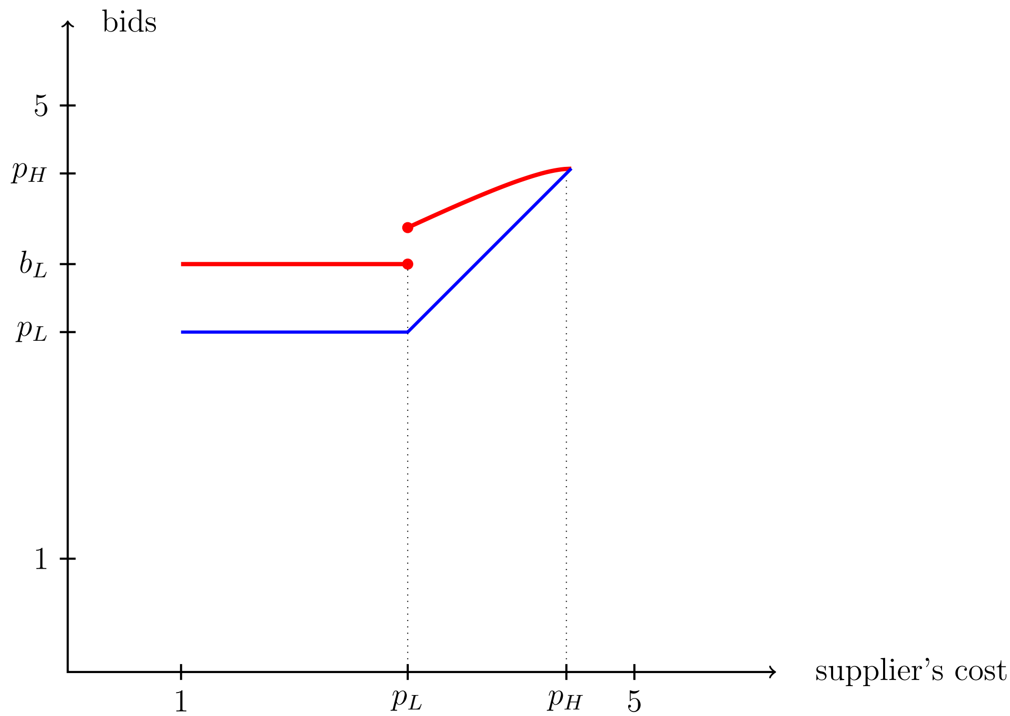

Figure 3: The blue line is the equilibrium bidding strategy in a LoLA with two bidders, costs drawn from the uniform distribution on $[1, 5]$, and $p_L = 3, p_H = 4.4$. The red line is the equilibrium bidding strategy in the equivalent first-price LoLA; bidders with cost larger than p_H choose not to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

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 expression 53).

5.5 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 2. 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 $f_1(x;a) = a \cdot \left(x - \frac{1}{2}\right) + 1$ 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:

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

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 $a = 0$, $v_0 = \frac{4}{3}$, and $\xi_2 = 4$ yields the symmetric example of Section 2. 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. (Refer to Appendix C for details about the computations.) 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 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

[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.

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

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 $(0, .5)$, 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.

Optimal auctions in asymmetric settings

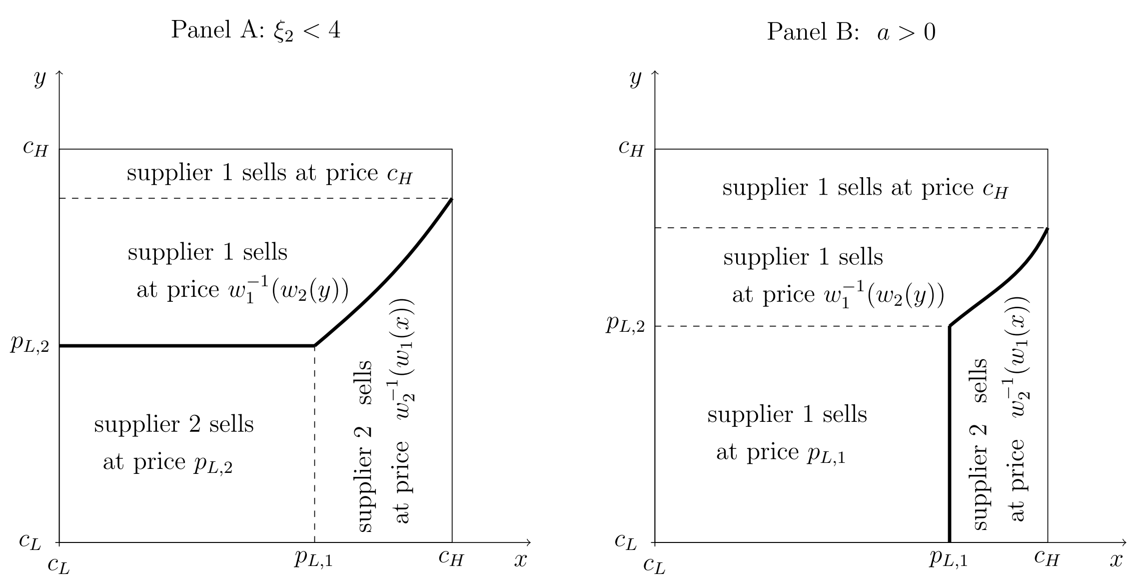

Figure 4: In panel A, the ex-ante expected gains from trading with either supplier are the same, but quality concerns are less severe for supplier 2. In panel B, supplier 1's cost distribution is higher than (i.e., stochastically dominates) its opponent's.

6 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,[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.

6.1 The available data

The available data is 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

[20]This information relates to Decarolis' (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' (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 5, 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.

[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

indicates that, in most cases, the government suffers a delay, a cost overrun, or both.[24]

Distributions of cost and quality measures

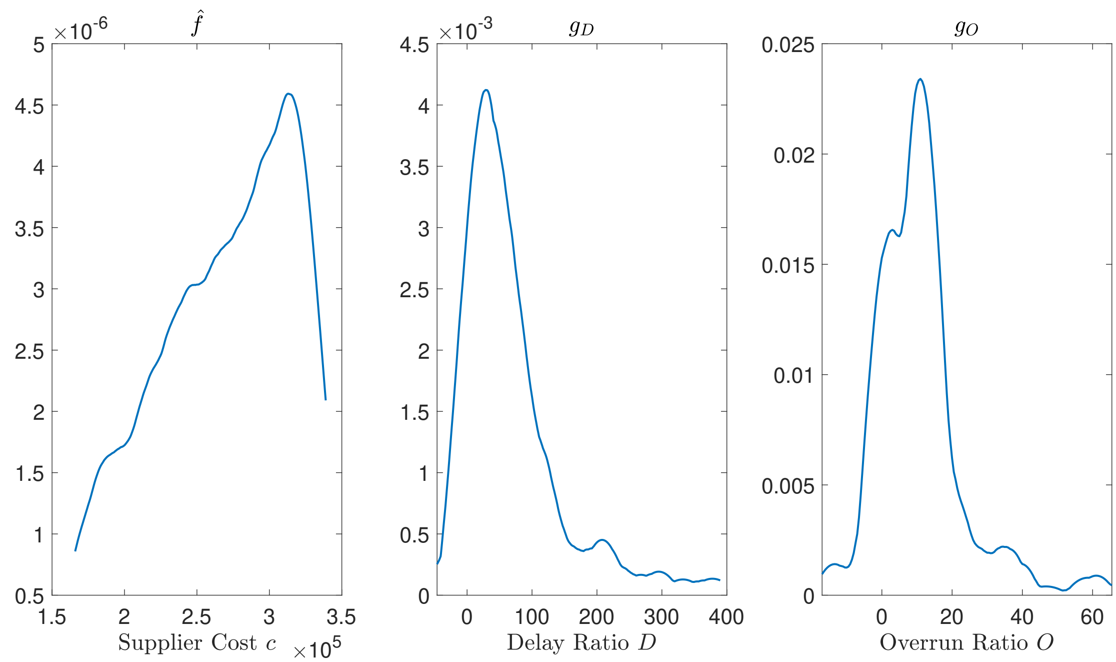

Figure 5: The left-hand panel depicts the estimated p.d.f. $\hat{f}_Z$ of the idiosyncratic cost component Z (unit is 10^5 euros) from Decarolis' (2018) assumed cost structure $c_i = y + z_i$, where the z_i's are iid draws from Z, and y is an auction-specific scalar. Without loss of generality we normalize $y = 0$, 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.

6.2 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:

$$v(c, \xi) = const - K\mathbb{E}\left[D(c, \xi) + O(c, \xi)\right], \tag{18}$$

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*.

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.[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 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 5.[26] Definition 3 in Appendix D 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:

$$\begin{aligned}\hat{v}(c,\xi) &= const - K\mathbb{E}\left[\hat{D}(c,\xi)+\hat{O}(c,\xi)\right] \\ &= const(\xi) - \xi K\left[\delta(c)+\omega(c)\right],\end{aligned} \tag{19}$$

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

[25] There is no difficulty in making expression (18) 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.

[26] 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.

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 Appendix D.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 11. 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 $\xi \in [0, 1]$.

6.3 Buyer-optimal and socially optimal mechanisms are LoLAs

We compute the calibrated virtual valuation function:

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

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]

Each of the left-hand panels 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 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.

The right-hand panels of Figure 6 are calibrated counterparts to Figure 2. Each right-

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

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 (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 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 8).

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-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.

Optimal mechanisms with varying degrees of quality concerns

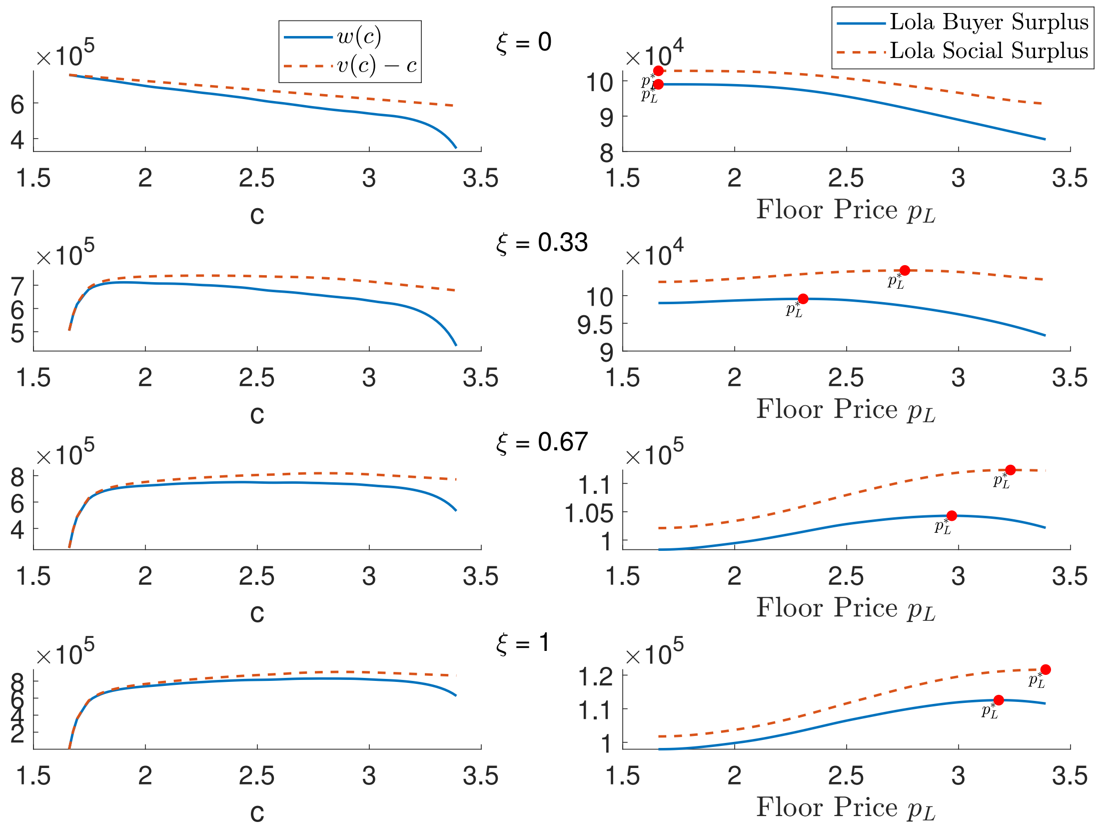

Figure 6: Virtual valuation functions $w(c)$ and gains from trade $v(c) - c$ for different values of ξ (left-hand column); expected buyer and social surplus in a LoLA with floor price p_L and no reserve price for different values of ξ (right-hand column). Recall that in our calibration it is optimal 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).

6.4 Performance of the buyer-optimal mechanism vs. 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

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

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

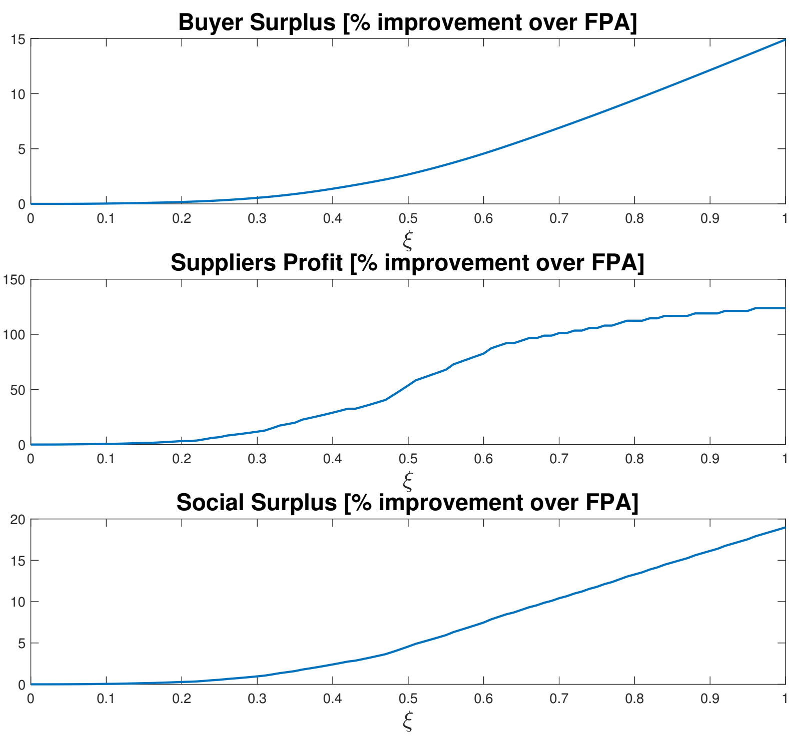

Figure 7: Performance improvement of optimal LoLA over first-price (or second-price) auction.

In Section 5.4 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_L's 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 expression 53 we find that b_L is up to 24% 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]

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 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.

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

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] Board, Simon (2007). "Bidding into the red: A model of post-auction bankruptcy." *The Journal of Finance* 62.6 (2007): 2695-2723.

[5] Bulow, Jeremy, and John Roberts (1989). "The simple economics of optimal auctions." *Journal of political economy* 97.5 (1989): 1060-1090.

[6] Burguet, Roberto, Juan-José Ganuza, and Esther Hauk (2012). "Limited liability and mechanism design in procurement." *Games and Economic Behavior* 76.1 (2012): 15-25.

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

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

[9] Che, Yeon-Koo, and Jinwoo Kim (2010). "Bidding with securities: Comment." American Economic Review 100.4 (2010): 1929-35.

[10] Chillemi, Ottorino, and Claudio Mezzetti (2014). "Optimal procurement mechanisms: bidding on price and damages for breach." Economic Theory 55.2 (2014): 335-355.

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

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

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

[14] 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

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

[16] Eun, Dong-Jae (2018) "Reducing ex-post adjustment by bid screening ex-ante: Evidence from procurement auctions for paving work." Manuscript, U. of Washington, August 7, 2018.

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

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

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

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

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

[22] Rhodes-Kropf, Matthew, and S. Viswanathan. (2005) "Financing auction bids." *RAND Journal of Economics* (2005): 789-815.

[23] Spulber, Daniel F. (1990) "Auctions and contract enforcement." *JL Econ & Org.* 6 (1990): 325.

[24] Waehrer, Keith. (1995) "A model of auction contracts with liquidated damages." *Journal of Economic Theory* 67.2 (1995): 531-555.

[25] World Bank (2016) "Abnormally Low Bids and Proposals: Guide to the identification and treatment of Abnormally Low Bids and Proposals." The World Bank, 1818 H Street NW Washington, DC 20433 USA. Published February 1, 2016 – Second Edition.

Appendices

A Proofs for Section 4

A.1 Proof of Theorem 1

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:

$$q_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) \equiv \begin{cases} 1 & \text{if } \ p_L \leq c_i < c_{-i}^{(1)} \\ 1 & \text{if } c_i \leq p_L < c_{-i}^{(1)} \\ \frac{1}{\kappa+1} & \text{if } \max\left\{c_{-i}^{(\kappa)}, c_i\right\} \leq p_L < c_{-i}^{(\kappa+1)} \\ 0 & \text{else} \end{cases} \tag{21}$$

and

$$m_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) \equiv \begin{cases} c_{-i}^{(1)} & \text{if } \ p_L \leq c_i < c_{-i}^{(1)} \\ c_{-i}^{(1)} & \text{if } c_i \leq p_L < c_{-i}^{(1)} \\ \frac{1}{\kappa+1} \cdot p_L & \text{if } \max\left\{c_{-i}^{(\kappa)}, c_i\right\} \leq p_L < c_{-i}^{(\kappa+1)} \\ 0 & \text{else,} \end{cases} \tag{22}$$

where $c_{-i}^{(\kappa)}$ denotes the κ-th lowest cost among all supplier i's opponents. For expositional simplicity, events where two or more bidders have the same cost are ignored in (21, 22) because they happen with probability zero.

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, $\forall i = 1, ..., N$,*

$$\forall c_i, c_i', c_{-i}, \qquad m_i(c_i, c_{-i}) - c_i \cdot q_i(c_i, c_{-i}) \geq m_i(c_i', c_{-i}) - c_i \cdot q_i(c_i', c_{-i}) \tag{23}$$

and

$$\forall c_i, c_{-i}, \qquad m_i(c_i, c_{-i}) - c_i \cdot q_i(c_i, c_{-i}) \geq 0. \tag{24}$$

Proof. It is well known in mechanism design that conditions (23-24) hold if and only if the following conditions hold jointly: $\forall c_{-i} \in [c_L, c_H]^{N-1}$

$$m_i^{\mathrm{L}}(c_H, c_{-i}; p_L, p_H) \geq c_H \cdot q_i^{\mathrm{L}}(c_H, c_{-i}; p_L, p_H) \tag{25}$$

$$q_i^{\mathrm{L}}(\cdot, c_{-i}; p_L, p_H) \text{ is nonincreasing,} \tag{26}$$

and

$$\forall\, c_i \in [c_L, c_H] \qquad m_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) = c_i \cdot q_i^{\mathrm{L}}(c_i, c_{-i}; p_L, p_H) + \int_{c_i}^{c_H} q_i^{\mathrm{L}}(t, c_{-i}; p_L, p_H)\, dt. \tag{27}$$

Therefore, it suffices to show that (25-27) hold. To this end, observe that the inequalities in (25) and the monotonicity in (26) are immediate. The envelope condition in (27) holds because both m^{L} and q^{L} are constant in c_i on $[c_L, p_L)$ and on $(p_H, c_H]$, and

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

■

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

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

and

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

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

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

which are invariant to permutations of the last $N-1$ *variables, i.e., letting* $\mathcal{N}$ *be the set* *of numbers* $\{1, ..., N\}$, $\forall i \in \mathcal{N}$, $\quad \forall$ *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 (6- 10) 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. Define:

$$\begin{aligned} Q(c_1) &\equiv \int_{[c_L, c_H]^{N-1}} q(c_1,\, c_{-1}) \cdot \prod_{j>1} dF(c_j) \\ M(c_1) &\equiv \int_{[c_L, c_H]^{N-1}} m(c_1,\, c_{-1}) \cdot \prod_{j>1} dF(c_j). \end{aligned} \tag{28}$$

First reformulation of the weighted welfare maximization problem

$$\max_{Q,M} \; N \int_{[c_L, c_H]} \left[\left[v(c_i, \xi) - (1-\beta) \cdot c_i \right] \cdot Q(c_i) - \beta \cdot M(c_i) \right] f(c_i)\, dc_i \tag{29}$$

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

$$M(c_i) - c_i \cdot Q(c_i) \geq M(c_i') - c_i \cdot Q(c_i'), \tag{30}$$

$$M(c_i) - c_i \cdot Q(c_i) \geq 0, \tag{31}$$

$$Q(c_i) \geq 0, \tag{32}$$

and

$$N \int_{c_L}^{c_i} Q(y)\, f(y)\, dy \;\leq\; 1 - \left[1 - F(c_i)\right]^N. \tag{33}$$

Lemma 2. *Restrict attention to symmetric mechanism. The value of the weighted welfare maximization problem (6- 10) is the same as the value of problem (29- 33).*

Proof. Because in solving problem (6- 10) we are restricting attention to mechanisms $(q_i, m_i)_{i=1,...,N}$ that are symmetric, the objective function (6) can be re-written as (29).

Similarly, the constraints (9) and (10) can be re-written as: (30) and (31). Furthermore, Border (1991) proves that, if the function q is symmetric in the sense of Definition 2, the demand constraints (7) and nonnegativity constraints (8) hold if and only if (32) and (33) are satisfied.

■

Problem (29- 33) can be further simplified, as follows.

Second reformulation of the weighted welfare maximization problem

$$\max_{Q} \; N \int_{c_L}^{c_H} w(c;\xi,\beta) \cdot Q(c) \cdot f(c)\, dc \tag{34}$$

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

$$Q \text{ is nonincreasing,} \tag{35}$$

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

$$Q(c) \geq 0, \tag{36}$$

and

$$N \int_{c_L}^{c} Q(y)\, f(y)\, dy \;\leq\; 1 - [1 - F(c)]^N . \tag{37}$$

Lemma 3. *The weighted welfare maximization problem (29- 33) can be reformulated as (34- 37).*

Proof. The incentive constraints (30) and (31) can be replaced without loss of generality by (35) and the envelope condition:

$$\forall\, c \in [c_L, c_H] \qquad M(c) = c \cdot Q(c) + \int_{c}^{c_H} Q(t) dt. \tag{38}$$

(This result is standard: see, e.g., Proposition 5.2 at p. 66 of Krishna 2010). Next, we use (38) to eliminate M from the problem. Substituting it into (29) and simplifying yields (34). Finally, (36) and (37) are identical to (32) and (33). ■

Next is the final reformulation of the problem.

Final (relaxed) formulation of the weighted welfare maximization problem

$$\max_{Q} \; N\int_{c_L}^{c_H} w(c;\xi,\beta)\cdot Q(c)\cdot f(c)\,dc \tag{39}$$

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

$$N\int_{c_L}^{c_H} w(c;\xi,\beta)\cdot Q(c)\cdot f(c)\,dc \;\leq\; N\int_{c_L}^{c_H} w(c;\xi,\beta)\cdot Q^L(c,p_L^*,p_H^*)\cdot f(c)\,dc, \tag{40}$$

where $Q^L(c,p_L^*,p_H^*)$ is given by expression (28) with q being replaced by $q_i^{\mathrm{L}}(c_i,c_{-i};p_L^*,p_H^*)$ from expression (21).

Problem (39-40) below is actually a relaxation of (34- 37). Aggregating constraints (35-37) into the single inequality (40) is the most innovative part of the proof. This aggregation is proved in the next lemma.

Lemma 4. *Any allocation function Q that satisfies (35-37) also satisfies (40).*

Proof. The proof consists in multiplying both sides of each inequality (35-37) by a nonnegative 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 (40).

The multipliers equal zero except:

$$\begin{cases} \forall c\in(p_H^*,c_H]: & \eta(c)\equiv -w(c;\xi,\beta)\cdot f(c) \\ \\ \forall c\in(p_L^*,p_H^*): & \delta(c)\equiv -w_c(c;\xi,\beta) \\ \\ \forall c\in[c_L,p_L^*): & \mu(c)\equiv \dfrac{F(c)}{F(p_L^*)}\displaystyle\int_{c_L}^{p_L^*} w(t;\xi,\beta)\,dF(t)-\int_{c_L}^{c} w(t;\xi,\beta)\,dF(t) \end{cases} \tag{41}$$

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

$$\mu(c_L) = \mu(p_L^*) = 0 \tag{42}$$

If $c_L < p_L^*$, then the definition of p_L^* in (12) 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 (42), implies that μ is nonnegative on $[c_L, p_L^*]$. Thus nonnegativity is established.

Now, we multiply both sides of: (35) by $\mu(c)$, (36) by $\eta(c)$, (37) by $\delta(c)$. We then integrate over c. Finally, we sum the three resulting inequalities. We arrive at:

$$\int_{c_L}^{p_L^*} \mu(y)\, dQ(y) \; + \int_{p_L^*}^{p_H^*} \delta(t) \int_{c_L}^{t} Q(y) f(y)\, dy\, dt - \int_{p_H^*}^{c_H} \eta(y)\, Q(y)\, dy \;\leq\; \int_{p_L^*}^{p_H^*} \delta(c)\, B(c)\, dc. \tag{43}$$

where

$$B(c) \equiv \frac{1}{N} \cdot \left(1 - [1 - F(c)]^N\right) \quad c \in [c_L, c_H]. \tag{44}$$

To see that (43) is equivalent to (40), let's focus first on the LHS of (43). The first integral can be rewritten as:

$$\begin{aligned} &\overbrace{\mu(p_L^*)}^{=0} \cdot Q(p_L^*) - \overbrace{\mu(c_L)}^{=0} \cdot Q(c_L) - \int_{c_L}^{p_L^*} Q(y)\, \mu'(y)\, dy \\ = &- \int_{c_L}^{p_L^*} \mu'(y)\, Q(y)\, dy. \end{aligned} \tag{45}$$

The second integral on the LHS of (43) can be rewritten as:

$$
\begin{aligned}
&\int_{c_L}^{p_L^*}\left(\int_{p_L^*}^{p_H^*}\delta(t)dt\right)Q(y)f(y)dy+\int_{p_L^*}^{p_H^*}\left(\int_{y}^{p_H^*}\delta(t)dt\right)Q(y)f(y)dy\\
=&\int_{c_L}^{p_L^*}\left(\int_{p_L^*}^{p_H^*}\delta(t)dt\right)Q(y)f(y)dy+\int_{p_L^*}^{p_H^*}w(y)Q(y)f(y)dy,
\end{aligned}
\tag{46}
$$

where equality holds because:

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

Adding (45) and (46) yields:

$$
\begin{aligned}
&\int_{c_L}^{p_L^*}\left(\int_{p_L^*}^{p_H^*}\delta(t)dt-\frac{\mu'(y)}{f(y)}\right)Q(y)\,f(y)\,dy\\
=&\int_{c_L}^{p_L^*}w(y)\,Q(y)\,f(y)\,dy,
\end{aligned}
\tag{47}
$$

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 (43) can be rewritten as:

$$
-\int_{p_H^*}^{c_H}\eta(y)Q(y)dy=\int_{p_H^*}^{c_H}w(y)f(y)Q(y)dy.
\tag{48}
$$

Combining (47) and (48) we can rewrite the LHS of (43) as

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

Let's now focus on the RHS of (43). Plugging in the expression for δ and simplifying yields:

$$\begin{aligned}
&-\int_{p_L^*}^{p_H^*} w'(c)\, B(c)\, dc \\
&= w(p_L^*)\, B(p_L^*) + \int_{p_L^*}^{p_H^*} w(c)\, B'(c)\, dc \\
&= \int_{p_L^*}^{p_H^*} w(c)\, (1-F(c))^{N-1}\, f(c)\, dc + \int_{c_L}^{p_L^*} \frac{1-[1-F(p_L)]^N}{N\cdot F(p_L)} w(c)\, f(c)\, dc,
\end{aligned} \tag{49}$$

where the second equality follows from the definition of B in (44). Now observe that:

$$\begin{aligned}
Q^L(c_1; p_L^*, p_H^*) &\equiv \int_{[c_L,c_H]^{N-1}} q(c_1,\, c_{-1}; p_L^*, p_H^*)\cdot \prod_{j>1} dF(c_j) \\
&= \begin{cases} 0, & c_1 \in (p_H^*, c_H]\,; \\ [1-F(c_1)]^{N-1}, & c_1 \in (p_L^*, p_H^*]\,; \\ \frac{1-\left[1-F(p_L^*)\right]^N}{N\cdot F(p_L^*)}, & c_1 \in [c_L, p_L^*]\,. \end{cases}
\end{aligned} \tag{50}$$

Hence (49) 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 (7 - 10). Indeed, unit demand (7) and nonnegativity (8) can be checked directly from the definition (21), 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 (9) and (10).

It remains to show that the mechanism $(q^{\mathrm{L}}, m^{\mathrm{L}})$ defined in (21) and (22) 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 (40) 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 (50) solves this relaxed problem because $Q^L(c, p_L^*, p_H^*)$ satisfies (40) with equality.

■

A.2 Proof of Proposition 2

Proof. The weighted welfare generated by the optimal LoLA is:

$$N \cdot \int_{c_L}^{p_H^*} w(c_i) \cdot Q_*^{\mathbb{L}}(c_i) \cdot dF(c_i)$$

$$= N \cdot \overbrace{\frac{1 - [1 - F(p_L^*)]^N}{N \cdot F(p_L^*)}}^{\bar{Q}} \cdot \int_{c_L}^{p_L^*} w(c_i) \cdot dF(c_i) + N \cdot \overbrace{\int_{p_L^*}^{p_H^*} w(c_i) \cdot [1 - F(c_i)]^{N-1} \cdot dF(c_i)}^{Q_*^{\mathbb{L}}(c_i) \text{ above } p_L}$$

$$= \frac{1 - [1 - F(p_L^*)]^N}{F(p_L^*)} \cdot \int_{c_L}^{p_L^*} w(c_i) \cdot dF(c_i) + N \cdot \int_{p_L^*}^{c_H} \max\{w(c_i), 0\} \cdot [1 - F(c_i)]^{N-1} \cdot dF(c_i)$$

$$= \Big(1 - [1 - F(p_L^*)]^N\Big) \cdot \overbrace{\int_{c_L}^{p_L^*} w(c_i) \cdot \frac{f(c_i)}{F(p_L^*)}\, dc_i}^{w(p_L^*)} + \int_{p_L^*}^{c_H} \max\{w(c_i), 0\}\, dG(c_i)$$

$$= \left(1 - \overbrace{[1 - F(p_L^*)]^N}^{\equiv P(N)}\right) \cdot w(p_L^*) + \overbrace{[1 - F(p_L^*)]^N}^{\equiv P(N)} \cdot \overbrace{\mathbb{E}\left[\max\left\{w(c^{(1)}), 0\right\} \mid p_L^* < c^{(1)}\right]}^{\equiv \mathcal{E}(N)}$$

$$= w(p_L^*) - P(N) \cdot \Big(w(p_L^*) - \mathcal{E}(N)\Big), \tag{51}$$

where $G(c_i) \equiv 1 - (1 - F(c_i))^N$ is the c.d.f. of the lowest cost, and $g(c_i) = G'(c_i) = N\,(1 - F(c_i))^{N-1}\, f(c_i)$ is its density.

The first equality follows from replacing $Q_*^{\mathbb{L}}$ with (56). The second equality follows from canceling N in the first term, and extending the second integral to c_H. The third equality follows from pulling $F(p_L^*)$ inside the integral, and using the definition of G. The fourth equality makes use of the fact that

$$w(p_L^*) = \frac{1}{F(p_L^*)} \int_{c_L}^{p^*} w(c)\, dF(c)\, dc, \tag{52}$$

which follows from integrating by parts the integral in(12).

Now note that, by definition (12), if p_L^* is greater than c_L it must lie in the region where the function w is decreasing. This region extends all the way to c_H by Assumption 1. Because $\max\{w(\cdot), 0\}$ is positive at p_L^* (see (52)) and nonincreasing on $[p_L, c_H]$, stochastic dominance implies that the conditional expectation $\mathcal{E}(N)$ in (51) is strictly increasing in N. This implies that the term in parenthesis in (51) is positive. Because $P(N)$ in (51) is decreasing in N, expression (51) is increasing in N. ■

B Proofs for Section 5

B.1 Proofs for Section 5.1

In this extension, each supplier $i = 1, ..., N$ draws its type $\theta_i = (c_i, x_i)$ independently from a distribution with density ϕ and support $\Theta_1 \equiv [c_L, c_H] \times [x_L, x_H]$. Let $\Theta \equiv \Theta_1^N$.

A direct mechanism $\mathcal{M}$ consists of $2N$ functions

$$\mathcal{M} \equiv \{q_i(c, x), m_i(c, x) \mid (c, x) \in \Theta\}_{i=1}^N$$

The restriction to symmetric mechanisms is wlog in the case as well.

Define supplier i's profit function as

$$\Pi(c_i, x_i) \equiv \sup \{M_i(c_i', x_i') - c_i \cdot Q_i(c_i', x_i') \mid (c_i', x_i') \in \Theta_1\}$$

Lemma 5. *If a mechanism is incentive compatible, then its reduced form Q must be independent of x_i, except possibly at for a zero measure set, i.e., for all i, all c_i, x_i and x_i'*

$$Q_i(c_i, x_i) = Q_i(c_i, x_i')$$

Proof. Standard mechanism design arguments imply that Π is convex and absolutely continuous.
The envelope theorem implies

$$\frac{\partial\,\Pi(c_i, x_i)}{\partial\, x_i} = 0 \quad \text{a.e.}$$

and

$$\frac{\partial\,\Pi(c_i, x_i)}{\partial\, c_i} = -Q(c_i, x_i) \quad \text{a.e.}$$

For any types (c_i, x_i) and (c_i', x_i'), the profit difference $\Pi(c_i', x_i') - \Pi(c_i, x_i)$ is equal to the line integral of the gradient of Π along any path. Therefore we have

$$\int_{c_i}^{c_i'} \overbrace{\frac{\partial\,\Pi(t, x_i)}{\partial\, c_i}}^{=-Q(t,x_i)} dt + \int_{x_i}^{x_i'} \overbrace{\frac{\partial\,\Pi(c_i', t)}{\partial\, x_i}}^{=0} dt \;=\; \int_{x_i}^{x_i'} \overbrace{\frac{\partial\,\Pi(c_i, t)}{\partial\, x_i}}^{=0} dt + \int_{c_i}^{c_i'} \overbrace{\frac{\partial\,\Pi(t, x_i')}{\partial\, c_i}}^{=-Q(t,x_i')} dt$$

If $(c_i, x_i) < (c_i', x_i')$, the path in the LHS is "first east and then north"; and the path in the RHS is "first north and then east".

Thus we have

$$\int_{c_i}^{c_i'} Q(t, x_i)dt \;=\; \int_{c_i}^{c_i'} Q(t, x_i')dt$$

Because the last equality must hold for any c_i and c_i', Q must be independent of x_i, except possibly for a zero measure set. ■

Lemma 5 implies that Π is also independent of x_i. and thus M must satisfy the envelope condition

$$M(c_i) = \Pi(c_L) + \int_{c_L}^{c_i} Q_i(t)\, dt,$$

and thus must be independent of x_i.

B.2 Proof for Section 5.4

Lemma 6. *Fix a LoLA with floor price p_L and reserve price p_H, and denote by $\bar{Q}$ the probability of winning for any type with cost below p_L. The FPLoLA with the same reserve price and minimum bid given by*

$$b_L = \frac{\bar{Q} - [1 - F(p_L)]^{N-1}}{\bar{Q}} \cdot p_L + \frac{[1 - F(p_L)]^{N-1}}{\bar{Q}} \cdot \beta(p_L; p_H) \tag{53}$$

is equivalent in the sense that it generates the same interim expected payments and profits for each supplier, and the same buyer's expected surplus.

The following strategy is a symmetric equilibrium in the equivalent first-price LoLA:

$$\beta^{f\mathbb{L}}(c_i; p_L, p_H) \equiv \begin{cases} b_L & \text{if } c_i \leq p_L \\ \beta(c_i; p_H) & \text{if } p_L < c_i \leq p_H \\ \text{no bid} & \text{if } c_i > p_H, \end{cases} \tag{54}$$

where

$$\beta(c_i; p_H) = \mathbb{E}\left[\min\left\{c_{-i}^{(1)}, p_H\right\} \mid c_i < c_{-i}^{(1)}\right]$$

is the equilibrium bidding strategy in the standard (reverse) first-price auction with reserve price p_H and no minimum bid, and $c_{-i}^{(1)}$ denotes the lowest cost among i's opponents.

Proof. The proof proceeds as follows. First, we show that when bidder i computes its best response in the FPLoLA, there is no loss of generality in ignoring the interval of "unused bids" $(b_L, \beta(p_L))$. Since all the remaining bids are used by some type, we can restate the best response problem as reporting a type in the direct revelation mechanism induced by $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$.

Next, we compute the interim probability of winning and expected payment for each type c_i in the direct revelation mechanism induced by $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$, and show that these functions coincide with their counterparts in the sincere equilibrium in the LoLA. This implies the equivalence of the two auction formats in terms of buyer expected surplus, interim expected payments, and expected profits.

Finally, because sincere bidding is an equilibrium in the LoLA, truthful reporting must also be an equilibrium in the direct revelation mechanism induced in the FPLoLA by $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$. Equivalently, $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$ is an equilibrium in the FPLoLA.

No loss of generality in ignoring "unused bids"

Suppose that all bidders in the FPLoLA except i follow the strategy $\beta^{f\mathbb{L}}$ given in (54). Then bidder i's expected payoff function is:

$$\Pi_i(b, c_i) = \begin{cases} 0 & \text{if no bid} \\ (b - c_i)\ [1 - F(\beta^{-1}(b))]^{N-1} & \text{if } \beta(b_L; p_H) \le b \le p_H \\ (b - c_i)\ [1 - F(b_L)]^{N-1} & \text{if } b_L < b < \beta(b_L; p_H) \\ (b_L - c_i)\ \bar{Q} & \text{if } b = b_L \end{cases}$$

For any $c_i \in [c_L, c_H]$, the payoff function $\Pi_i(\cdot, c_i)$ is linear and strictly increasing on the middle interval $(b_L, \beta(b_L; p_H)]$. Therefore all bids in this interval cannot be optimal for any type. Once all bids in $(b_L, \beta(b_L; p_H)]$ are removed from consideration, all remaining bids are in the range of $\beta^{f\mathbb{L}}$. Since all the remaining bids are used by some type, we can interpret choosing the best response in the FPLoLA as choosing a type report in the direct revelation mechanism induced by $\beta^{f\mathbb{L}}$.

The direct revelation mechanism induced by the LoLA

In the LoLA, for any type profile (c_i, c_{-i}) supplier i sells with probability

$$q_i^{\mathbb{L}}(c_i, c_{-i}) = \begin{cases} 1 & \text{if } c_i < \min\{c_{-i}^{(1)},\, p_H\} \text{ and } p_L < c_{-i}^{(1)} \\ \frac{1}{k+1} & \text{if } c_i \le p_L \text{ and } c_{-i}^{(k)} \le p_L < c_{-i}^{(k+1)} \\ 0 & \text{else} \end{cases} \tag{55}$$

The resulting *interim* probability of selling is

$$Q^{\mathbb{L}}(c_i) \equiv \int_{[c_L,c_H]^{N-1}} q_i^{\mathbb{L}}(c_i, c_{-i}) \prod_{j\neq i}^{N} f(c_j)\, dc_j$$

$$= \begin{cases} \bar{Q} & \text{if } c_i \in [c_L, p_L] \\ [1 - F(c_i)]^{N-1} & \text{if } c_i \in (p_L, p_H] \\ 0 & \text{if } c_i \in (p_H, c_H] \end{cases} \tag{56}$$

where

$$\begin{aligned} \bar{Q} &= \sum_{j=0}^{N-1} \overbrace{\binom{N-1}{j} \cdot F(p_L)^j \cdot [1 - F(p_L)]^{N-1-j}}^{\Pr[j \text{ opponents have cost below } p_L]} \cdot \overbrace{\frac{1}{j+1}}^{\text{uniform tie-breaking}} \\ &= \sum_{j=0}^{N-1} \frac{(N-1)!}{j!\,(N-1-j)!} F(p_L)^j (1 - F(p_L))^{N-1-j} \frac{1}{j+1} \\ &= \frac{1}{N\,F(p_L)} \sum_{j=0}^{N-1} \frac{N!}{(j+1)!\,(N-1-j)!} F(p_L)^{j+1} (1 - F(p_L))^{N-1-j} \\ &= \frac{1}{N\,F(p_L)} \sum_{t=1}^{N} \binom{N}{t} F(p_L)^t (1 - F(p_L))^{N-t} \\ &= \frac{1}{N\,F(p_L)} \left[1 - (1 - F(p_L))^N\right] \end{aligned}$$

In a LoLA, the *ex post* expected payment function is

$$m_i^{\mathbb{L}}(c_i, c_{-i}) = \begin{cases} \min\{c_{-i}^{(1)},\, p_H\} & \text{if } c_i < \min\{c_{-i}^{(1)},\, p_H\} \text{ and } p_L < c_{-i}^{(1)} \\ \frac{1}{k+1}\, p_L & \text{if } c_i \leq p_L \text{ and } c_{-i}^{(k)} \leq p_L < c_{-i}^{(k+1)} \\ 0 & \text{else} \end{cases} \tag{57}$$

and the resulting *interim* expected payment function is

$$
\begin{aligned}
M^{\mathbb{L}}(c_i) &\equiv \int_{[c_L,c_H]^{N-1}} m_i^{\mathbb{L}}(c_i, c_{-i}) \prod_{j\neq i}^{N} f(c_j)\, dc_j \\
&= \begin{cases} \bar{M} & \text{if } c_i \in [c_L, p_L] \\ \beta(c_i; p_H) \cdot [1 - F(c_i)]^{N-1} & \text{if } c_i \in (p_L, p_H] \\ 0 & \text{if } c_i \in (p_H, c_H] \end{cases}
\end{aligned} \tag{58}
$$

We have

$$
\bar{M} = \bar{Q} \cdot \ p_L + [1 - F(p_L)]^{N-1} \cdot [\beta(p_L; p_H) - p_L]
$$

The first term in the RHS captures the fact that any bidder with type below p_L wins with probability $\bar{Q}$ and is paid at least p_L. The second term captures the event in which the costs of all the bidder's opponents exceed p_L; in this case, which happens with probability $[1 - F(p_L)]^{N-1}$, the bidder is paid more.

The second line in (58) holds because any type $c_i \in [p_L, p_H]$ sells at price $\min\left\{c_{-i}^{(1)}, p_H\right\}$ when $c_i < c_{-i}^{(1)}$. Therefore its expected payment is

$$
\int_{c_i}^{c_H} \min\{y, p_H\}\, dF_{-i}^{(1)}(y) = \beta(c_i; p_H) \cdot [1 - F(c_i)]^{N-1}.
$$

The direct revelation mechanism induced by $\beta^{f\mathbb{L}}$ in the FPLoLA coincides with its LoLA counterpart.

Type c_i's interim probability of winning in the direct revelation mechanism induced by $\beta^{f\mathbb{L}}$ is the same as the probability of winning in the FPLoLA assuming that all other bidders follow the strategy $\beta^{f\mathbb{L}}$ given in (54) and i bids according to $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$. Because the strategy $\beta^{f\mathbb{L}}$ is strictly increasing in the region above p_L and flat below p_L, the regions of the type space in which the lowest type wins with probability 1 are the same as in the sincere equilibrium of the equivalent LoLA. Similarly, the regions in which multiple suppliers win with positive probability are the same as in the two auctions. Therefore, both the ex post and interim probability of winning in the FPLoLA are the same as in the sincere equilibrium of the equivalent LoLA.

To see that the interim expected payment function in the first-price LoLA is equal to

$M^{\mathbb{L}}$, note first that all types above p_H do not bid and thus are paid zero. Next, note that any type between p_L and p_H sells at price $\beta(c_i; p_H)$ when all other suppliers bid above. Therefore the expected payment of all types between p_L and p_H is $\beta(c_i; p_H) \cdot [1 - F(c_i)]^{N-1}$, as in (58). Finally, for all types below p_L, the *interim* expected payment

$$b_L \cdot \bar{Q} = \left[\bar{Q} - [1 - F(p_L)]^{N-1} \right] \cdot p_L + [1 - F(p_L)]^{N-1} \cdot \beta(p_L; p_H)$$

is equal to $\bar{M}$ in (58).

Equivalence between FPLoLA and LoLA

Because the direct revelation mechanism induced by $\beta^{f\mathbb{L}}$ in the FPLoLA coincides with its LoLA counterpart, the two auction formats are equivalent in terms of buyer expected surplus, interim expected payments, and expected profits.

$\beta^{f\mathbb{L}}(c_i; p_L, p_H)$ **is an equilibrium in the FPLoLA**

Because sincere bidding is an equilibrium in the LoLA, truthful reporting must also be an equilibrium in the direct revelation mechanism induced in the FPLoLA by $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$. Equivalently, $\beta^{f\mathbb{L}}(c_i; p_L, p_H)$ is an equilibrium in the FPLoLA.

■

The next figure compares the equilibrium outcomes in a LoLA and its equivalent

FPLoLA.

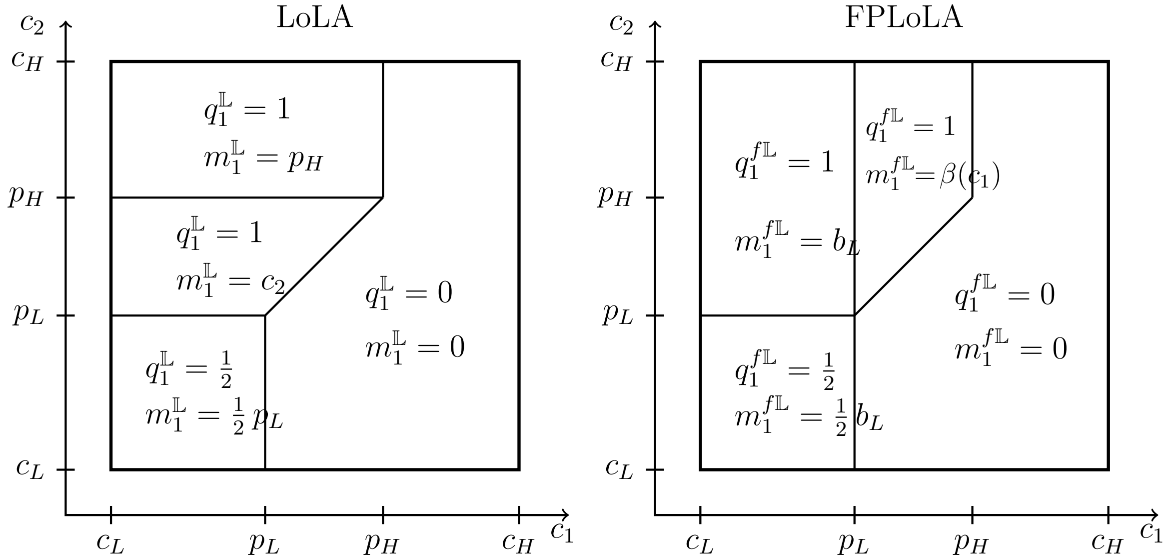
LoLA
c_2
c_H
p_H
p_L
c_L
$q_1^{\mathbb{L}} = 1$
$m_1^{\mathbb{L}} = p_H$
$q_1^{\mathbb{L}} = 1$
$m_1^{\mathbb{L}} = c_2$
$q_1^{\mathbb{L}} = 0$
$m_1^{\mathbb{L}} = 0$
$q_1^{\mathbb{L}} = \frac{1}{2}$
$m_1^{\mathbb{L}} = \frac{1}{2}\,p_L$
c_L p_L p_H c_H c_1
FPLoLA
c_2
c_H
p_H
p_L
c_L
$q_1^{f\mathbb{L}} = 1$
$m_1^{f\mathbb{L}} = b_L$
$q_1^{f\mathbb{L}} = 1$
$m_1^{f\mathbb{L}} = \beta(c_1)$
$q_1^{f\mathbb{L}} = 0$
$m_1^{f\mathbb{L}} = 0$
$q_1^{f\mathbb{L}} = \frac{1}{2}$
$m_1^{f\mathbb{L}} = \frac{1}{2}\,b_L$
c_L p_L p_H c_H c_1

C Numerical Analysis of Asymmetric Environments (NOT FOR PUBLICATION)

In this appendix we solve numerically for the optimal mechanism in the two-bidder environments described in Section 5.5. These environments are characterized by small asymmetries in ξ_2 and a around the parameter constellation $\xi_2 = 4$ and $a = 0$ that defines the symmetric case. The symmetric case is treated in Section 2.

The densities of x and y are:

$$f_1(x; a) = a \cdot \left(x - \tfrac{1}{2}\right) + 1 \quad \text{and } f_2(y) = 1.$$

The buyer's willingness to pay for suppliers 1's good and 2's good are

$$v_1(x; \xi_1) = v_0 - \xi_1 \cdot \left(\tfrac{1}{2}\, x^2 - x + \tfrac{1}{3}\right) \quad \text{and} \quad v_2(y; \xi_2) = v_0 - \xi_2 \cdot \left(\tfrac{1}{2}\, y^2 - y + \tfrac{1}{3}\right).$$

The virtual valuations are

$$w_1(x; \xi_1, a) = v_1(x; \xi_1) - x - \tfrac{F_1(x)}{f_1(x)} = v_0 - x - \xi_1 \left(\tfrac{x^2}{2} - x + \tfrac{1}{3}\right) - \frac{\frac{a\,x^2}{2} - x\left(\frac{a}{2} - 1\right)}{a\left(x - \frac{1}{2}\right) + 1}$$

and

$$w_2(y; \xi_2) = v_2(y; \xi_2) - y - \tfrac{F_2(y)}{f_2(y)} = v_0 - 2\,y - \xi_2 \left(\tfrac{y^2}{2} - y + \tfrac{1}{3}\right).$$

Throughout, ξ_1 is set to 4. Setting $\xi_2 = 4$ and $v_0 = \frac{4}{3}$ yields the example of Section 2. Here, we set $v_0 = 2$ to guarantee that, when we introduce asymmetries, both virtual valuations remain positive.

The next two sections display virtual valuations and cost distributions, and the corresponding optimal auctions computed by our software application. The software solves a discretized version of the following auctioneer's problem:

$$\begin{cases} \max\limits_{q_1,\, q_2} \displaystyle\int_0^1 \int_0^1 \; [w_1(x,y)\cdot q_1(x,y) + w_2(y,x)\cdot q_2(y,x)]\, f_1(x)\, f_2(y)\, dx\, dy \\ \text{subject to, for all } (x,y)\in[0,1]^2\text{:} \\ \quad q_1(x,y) + q_2(x,y) \le 1 \\ \quad q_1(\cdot,y) \;\; \text{nonincreasing} \\ \quad -q_1(1,y) \le 0 \\ \quad q_2(\cdot,x) \;\; \text{nonincreasing} \\ \quad -q_2(1,x) \le 0 \end{cases}$$

Section C.1 deals with parameter configurations where a is fixed at zero and ξ_2 varies in the interval $(2.5, 4)$. Section C.2 deals with parameter configurations where ξ_2 is fixed at 4 and a varies in the interval $(0, .5)$.

C.1 Asymmetry on ξ

$\xi_1 = 4, \quad \xi_2 = 2.5, \quad a = 0, \quad p_{L,1} = 0.747475, \quad p_{L,2} = 0.57$

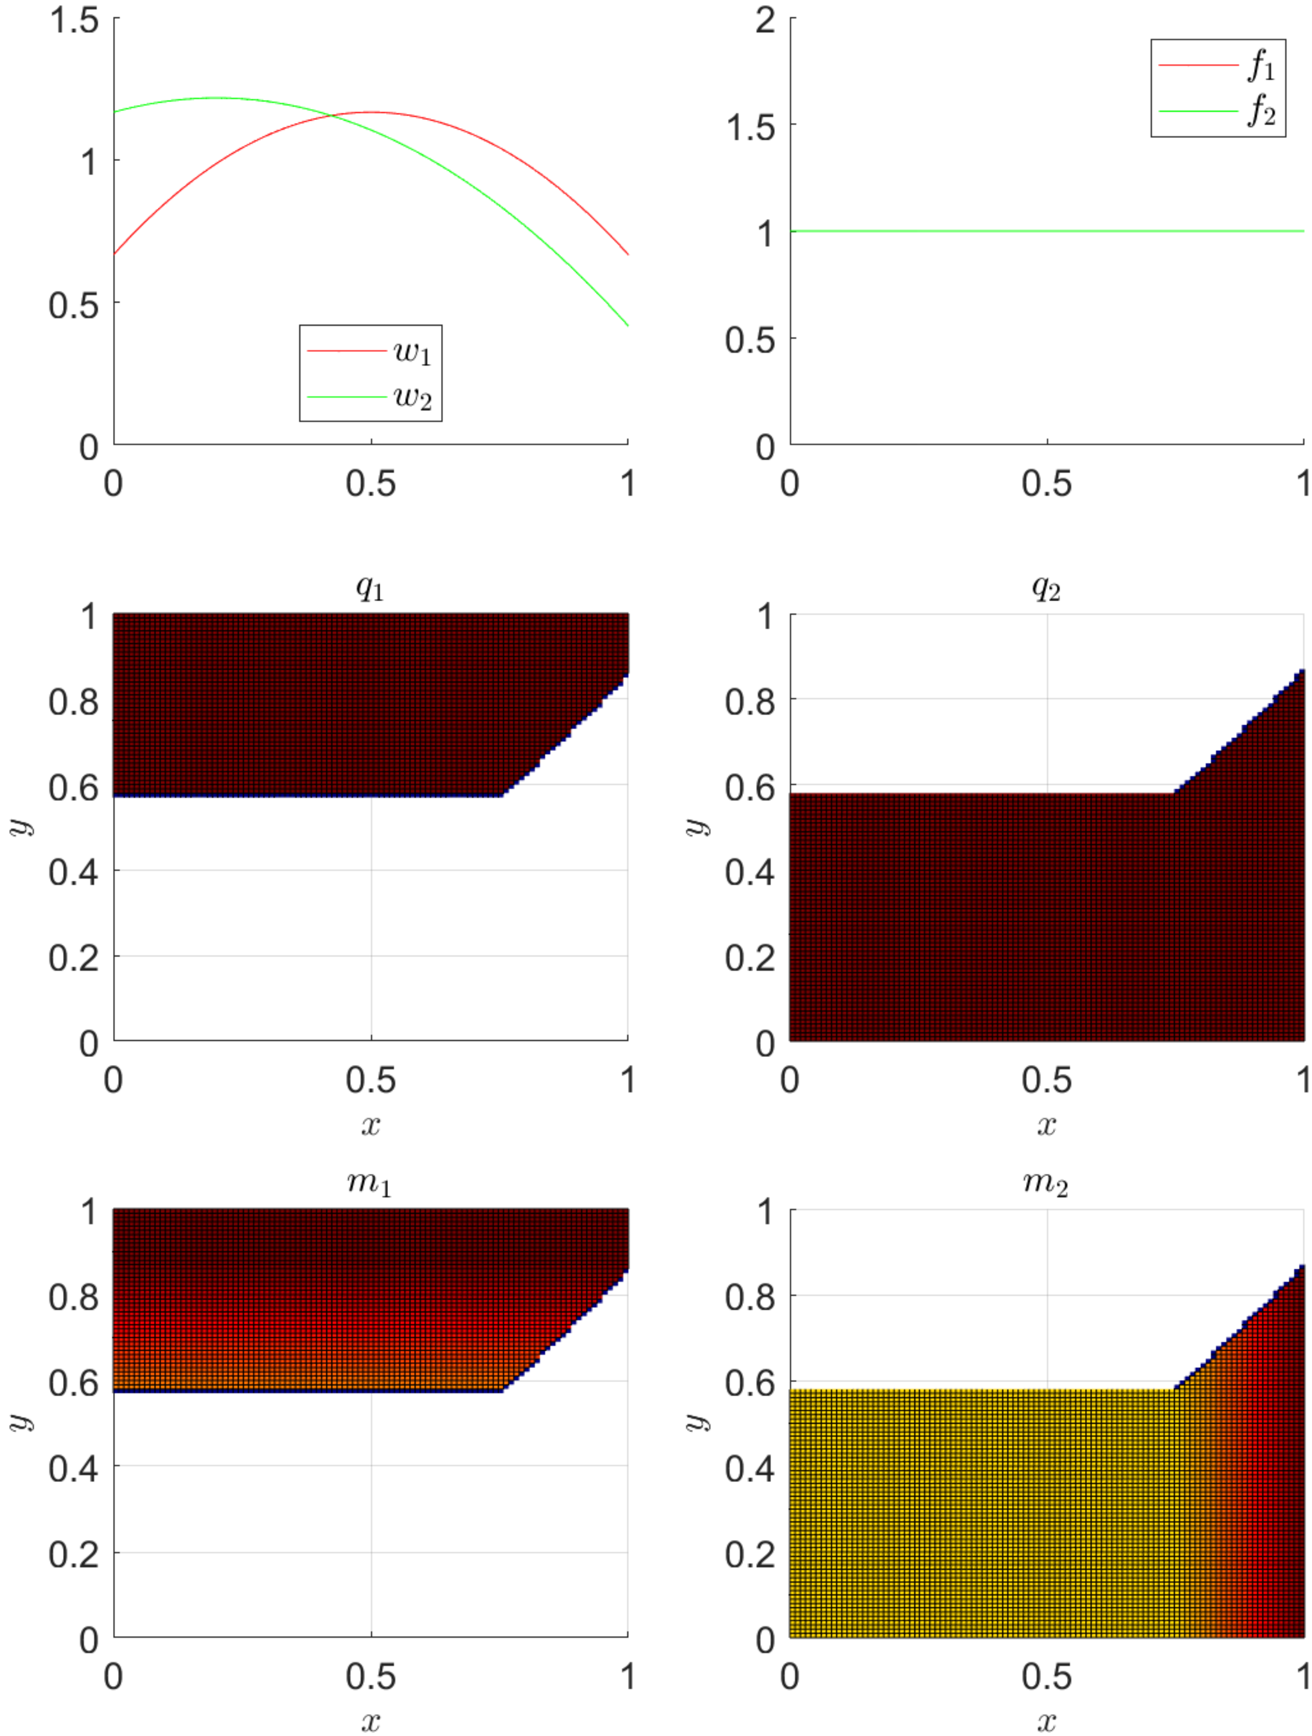

$$\xi_1 = 4, \quad \xi_2 = 2.8, \quad a = 0, \quad p_{L,1} = 0.757576, \quad p_{L,2} = 0.6$$

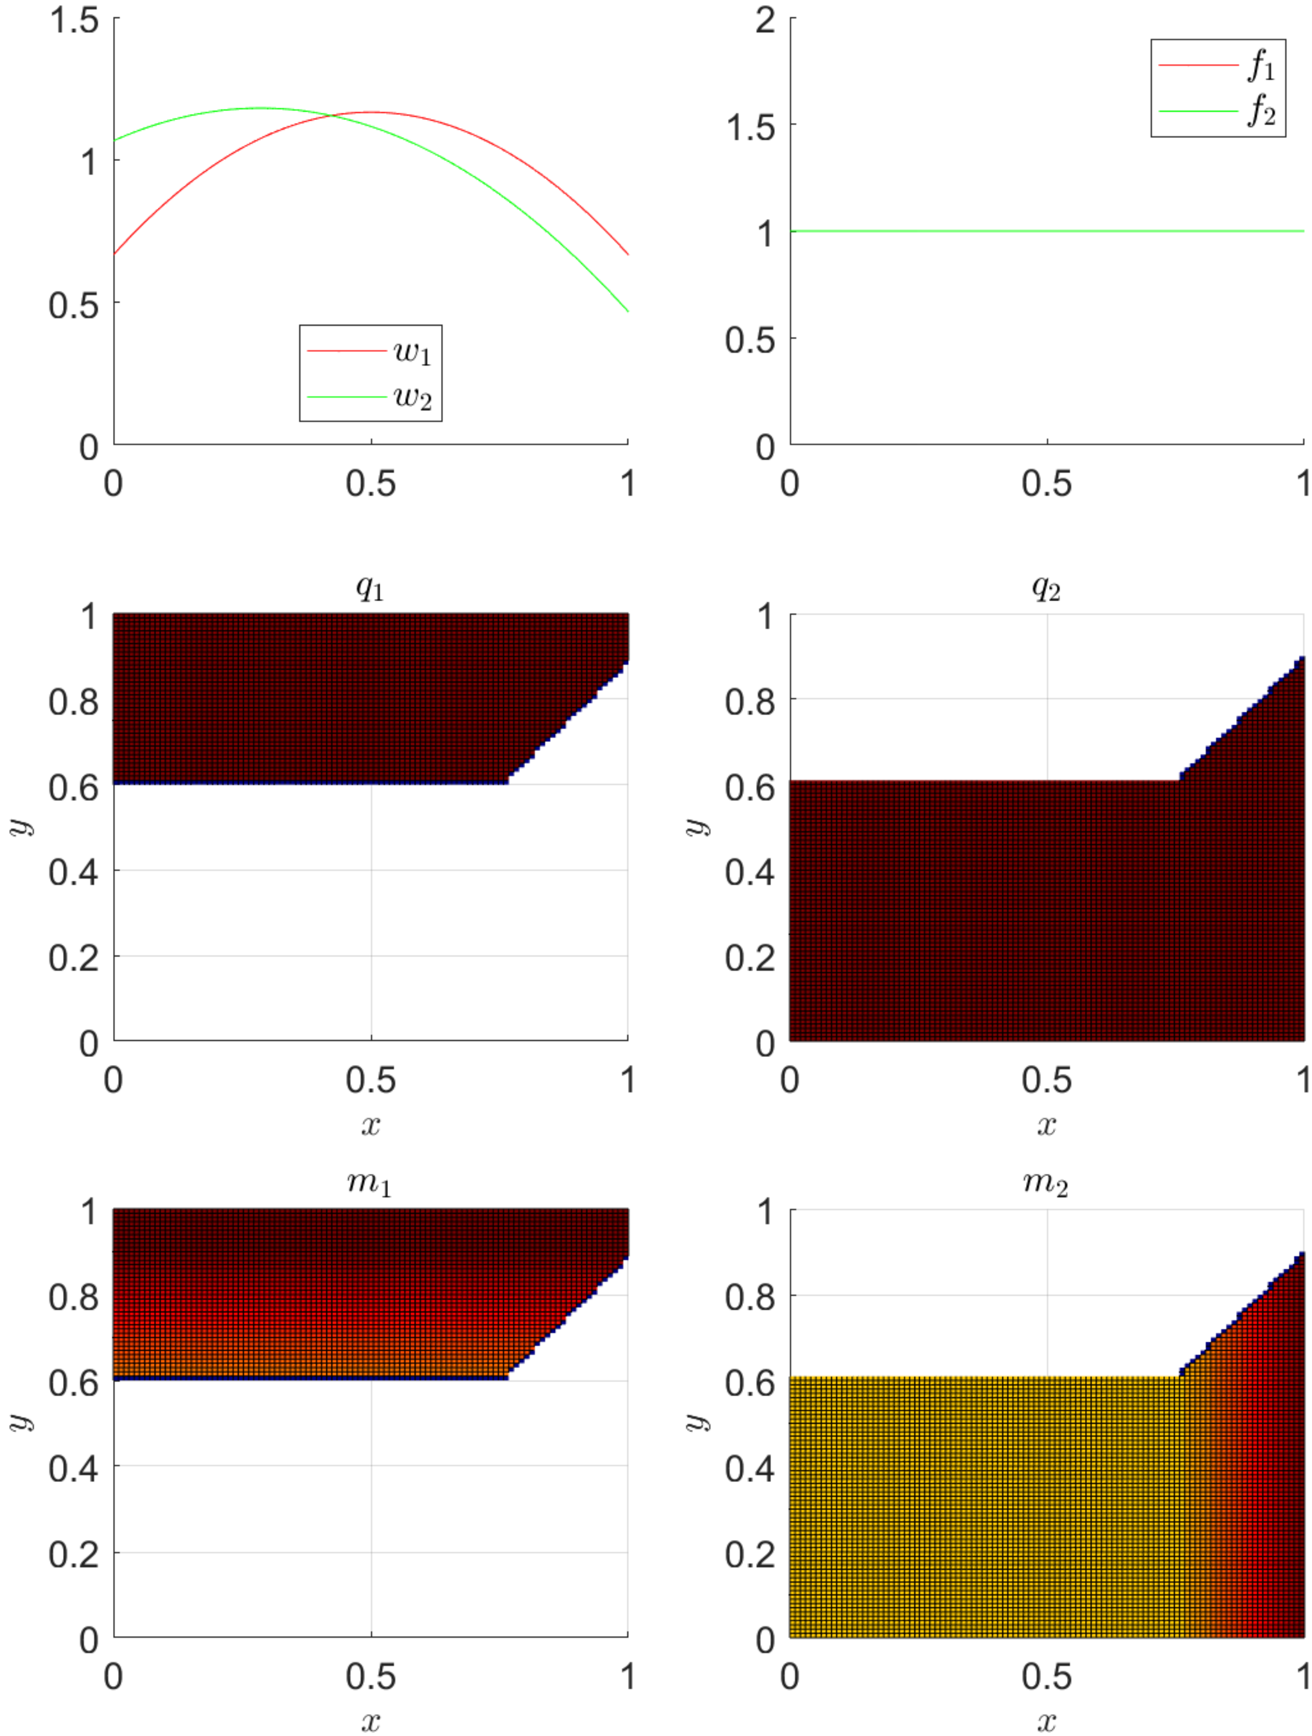

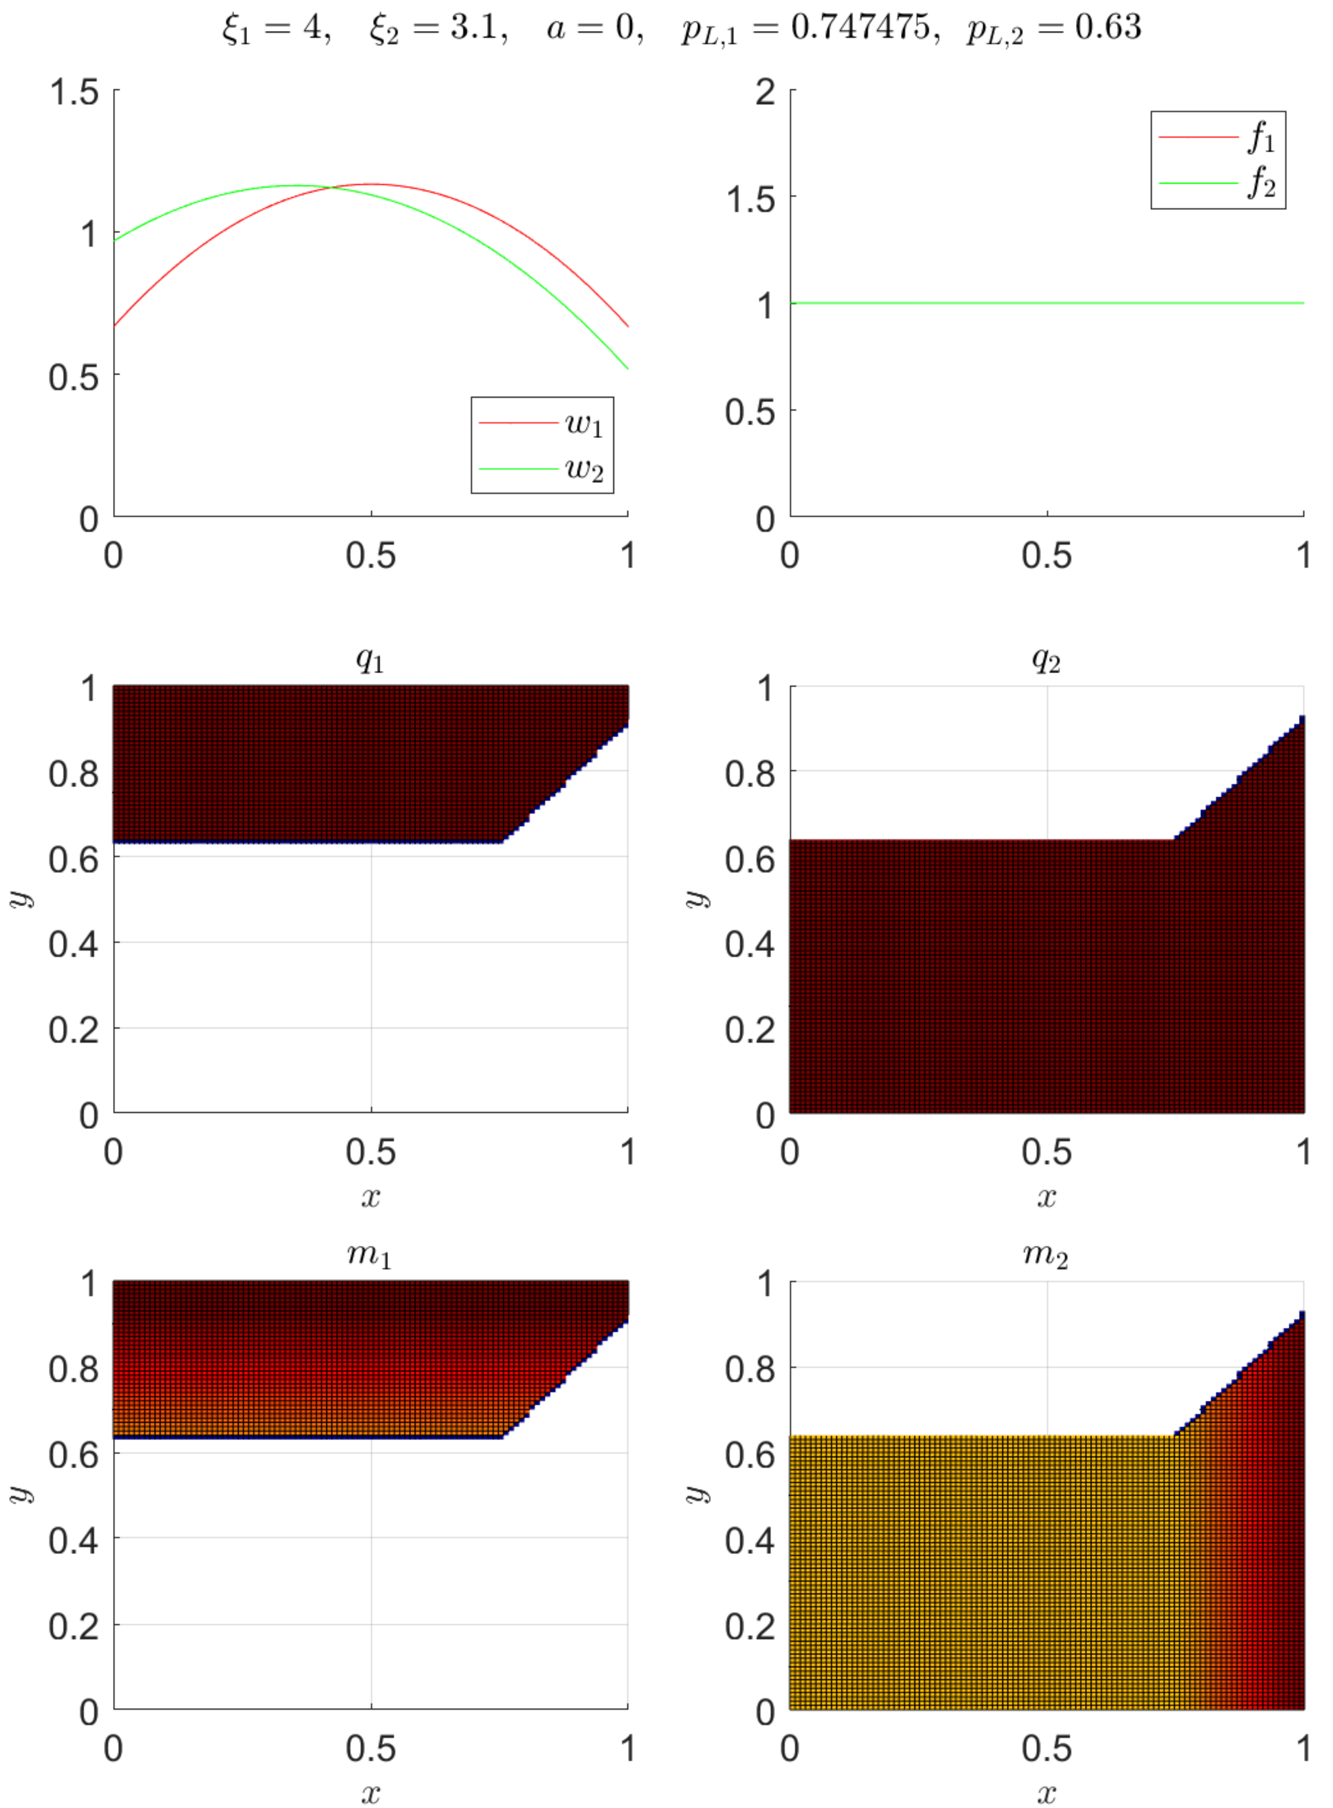

$\xi_1 = 4, \quad \xi_2 = 3.1, \quad a = 0, \quad p_{L,1} = 0.747475, \quad p_{L,2} = 0.63$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

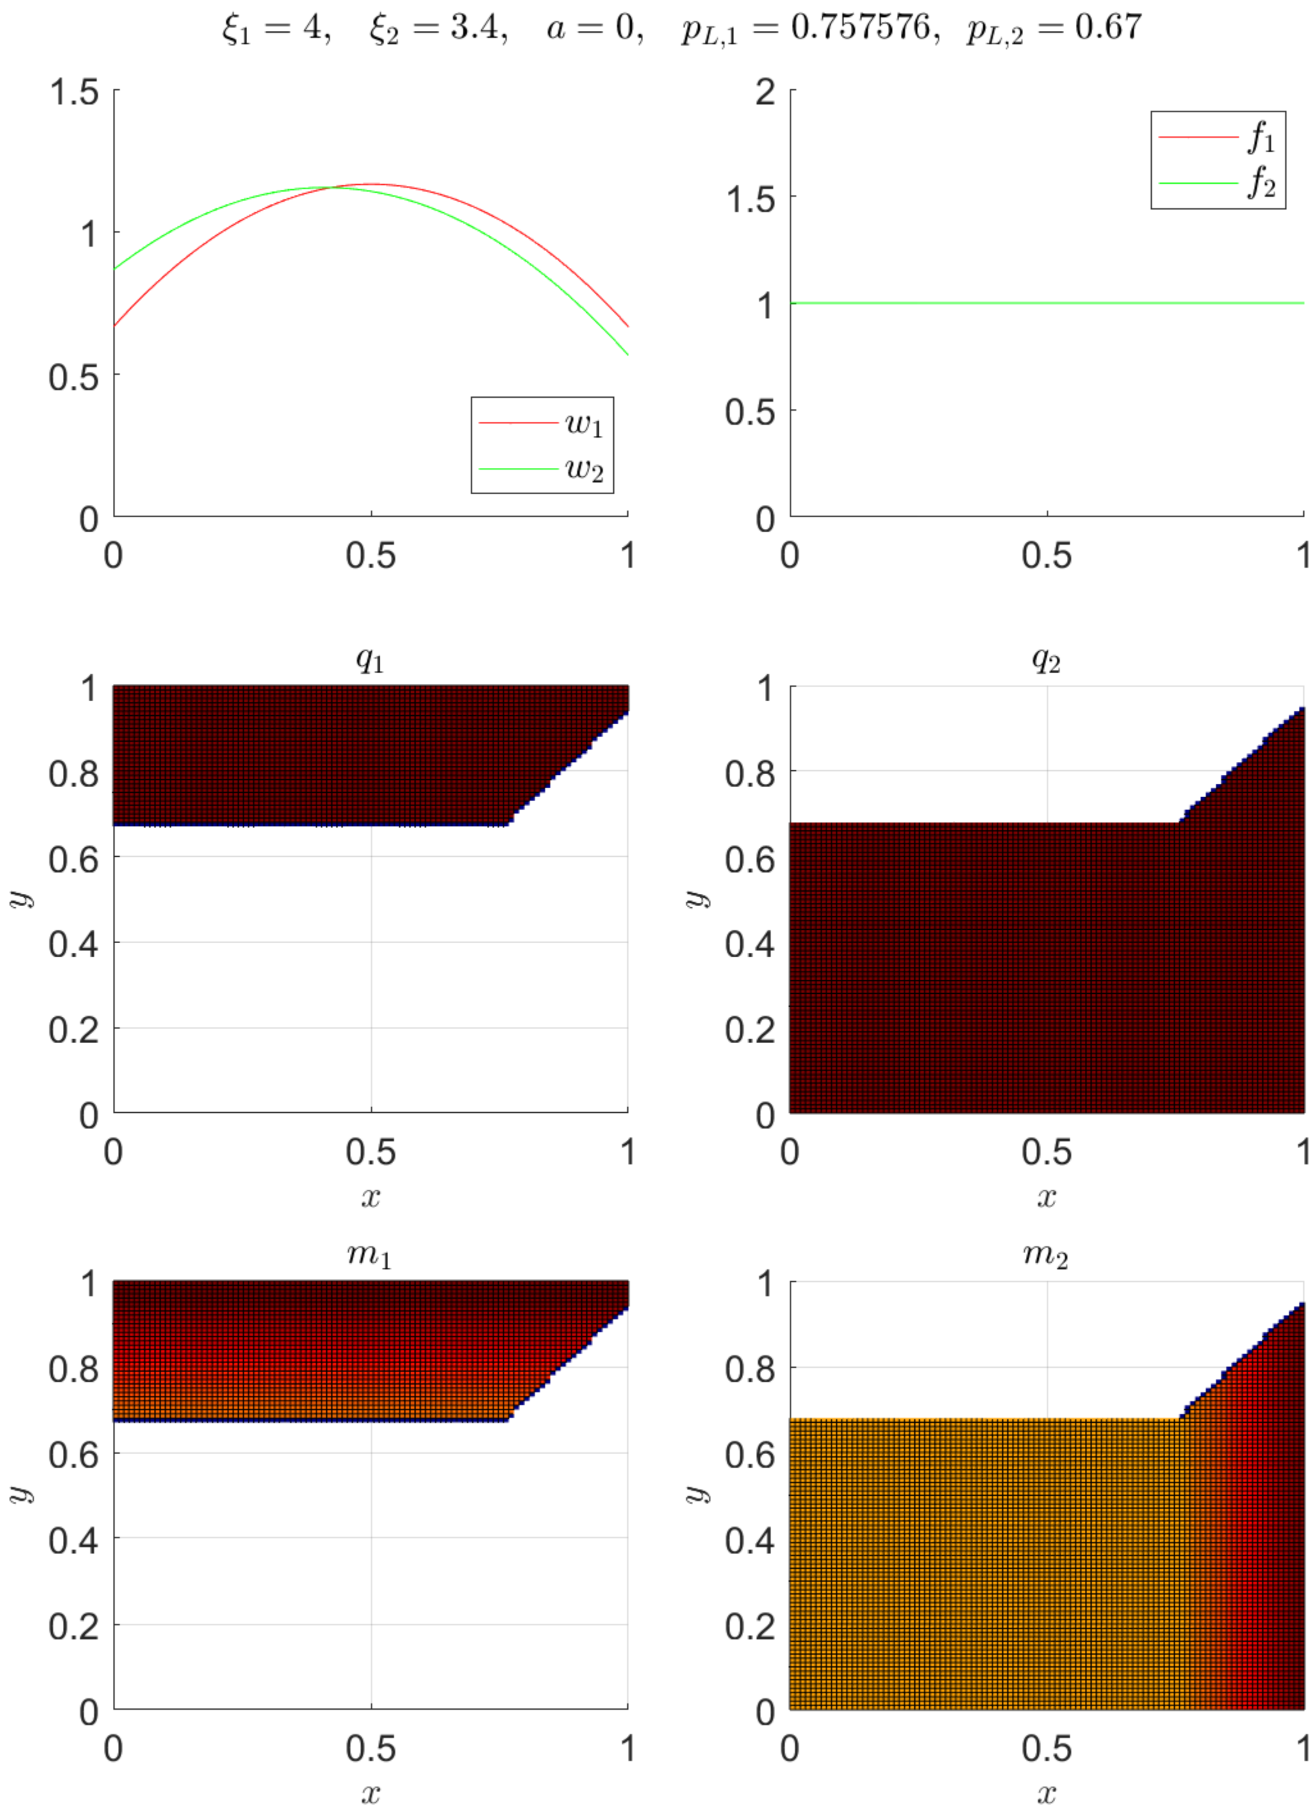
$\xi_1 = 4, \quad \xi_2 = 3.4, \quad a = 0, \quad p_{L,1} = 0.757576, \ \ p_{L,2} = 0.67$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

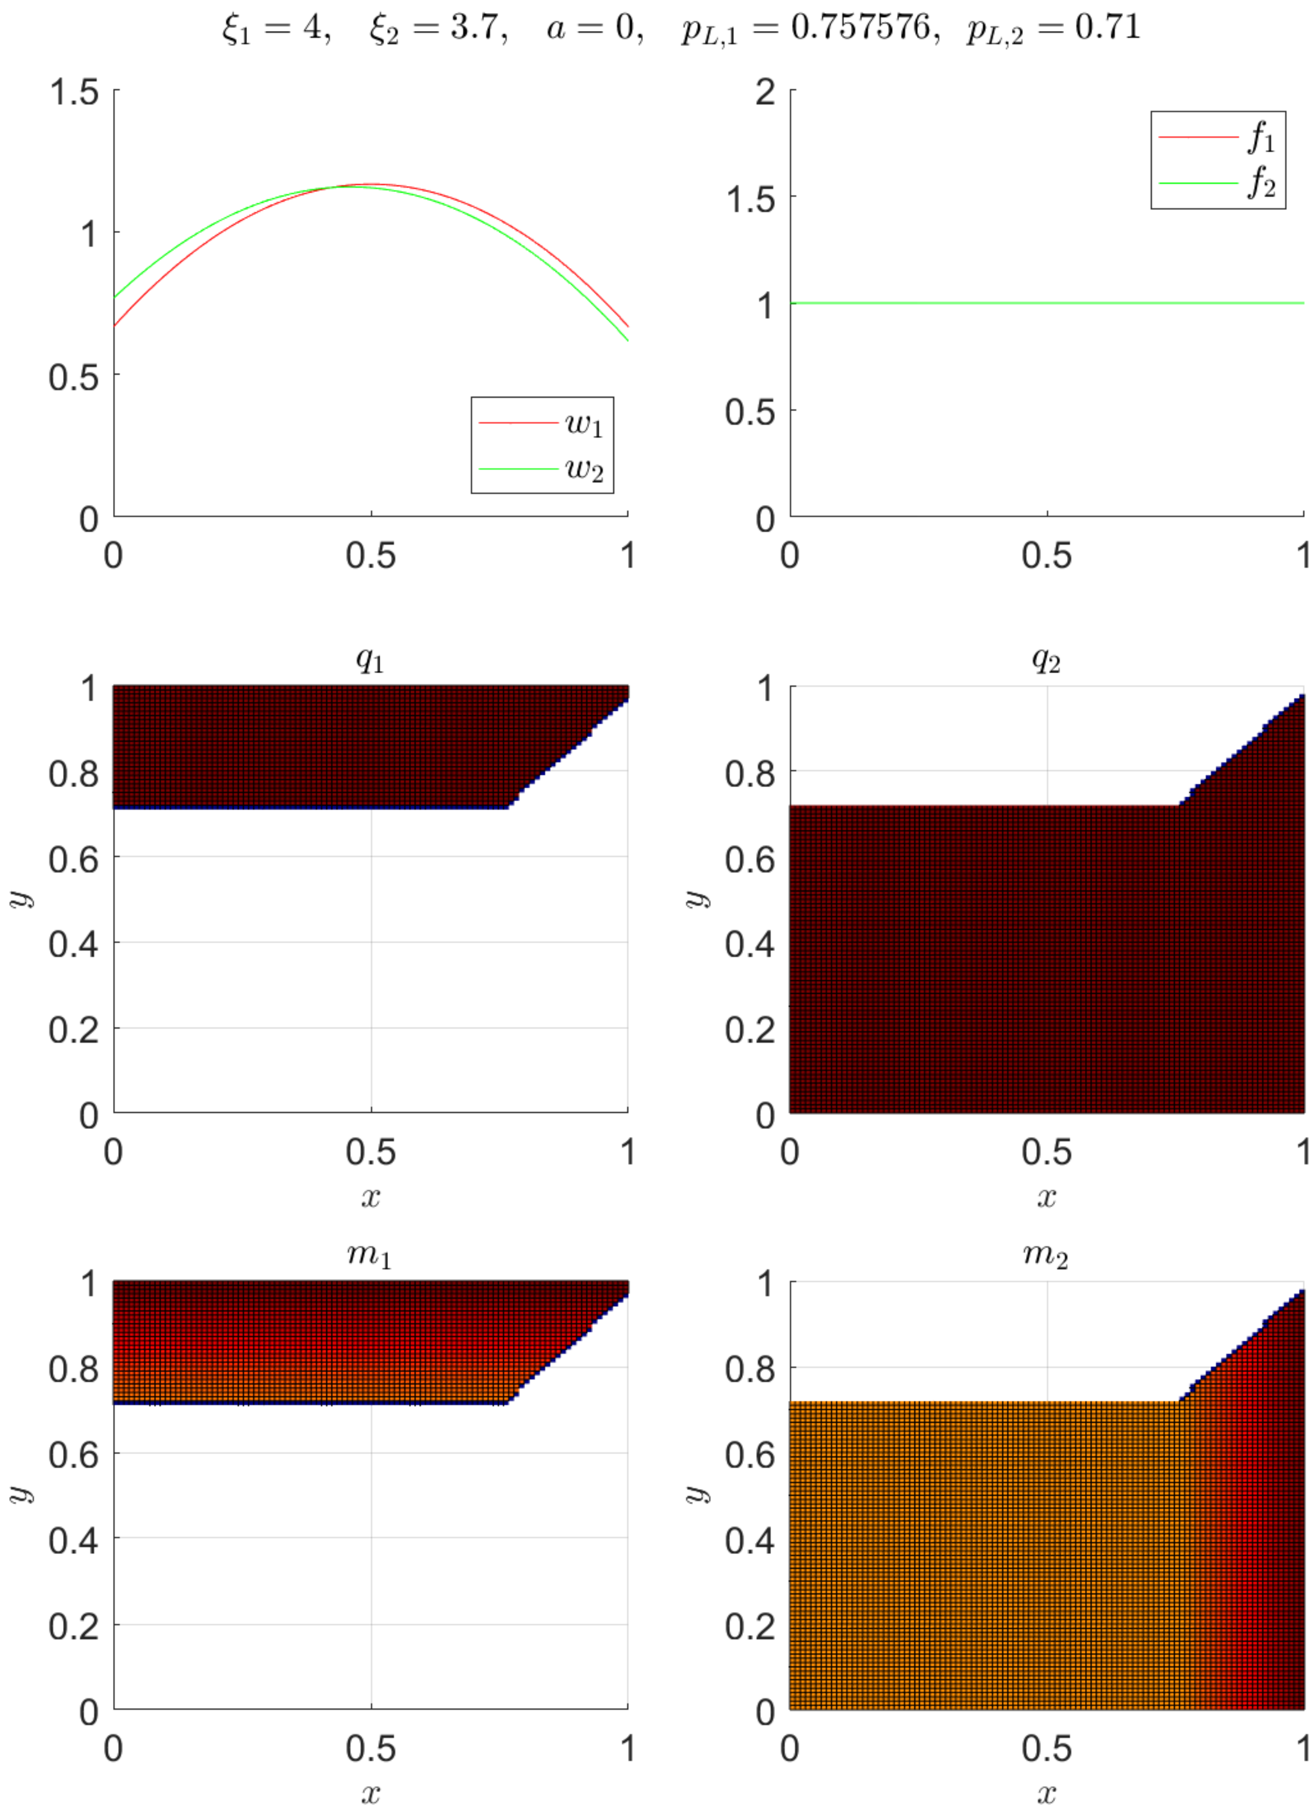

$\xi_1 = 4, \quad \xi_2 = 3.7, \quad a = 0, \quad p_{L,1} = 0.757576, \quad p_{L,2} = 0.71$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

C.2 Asymmetry on f

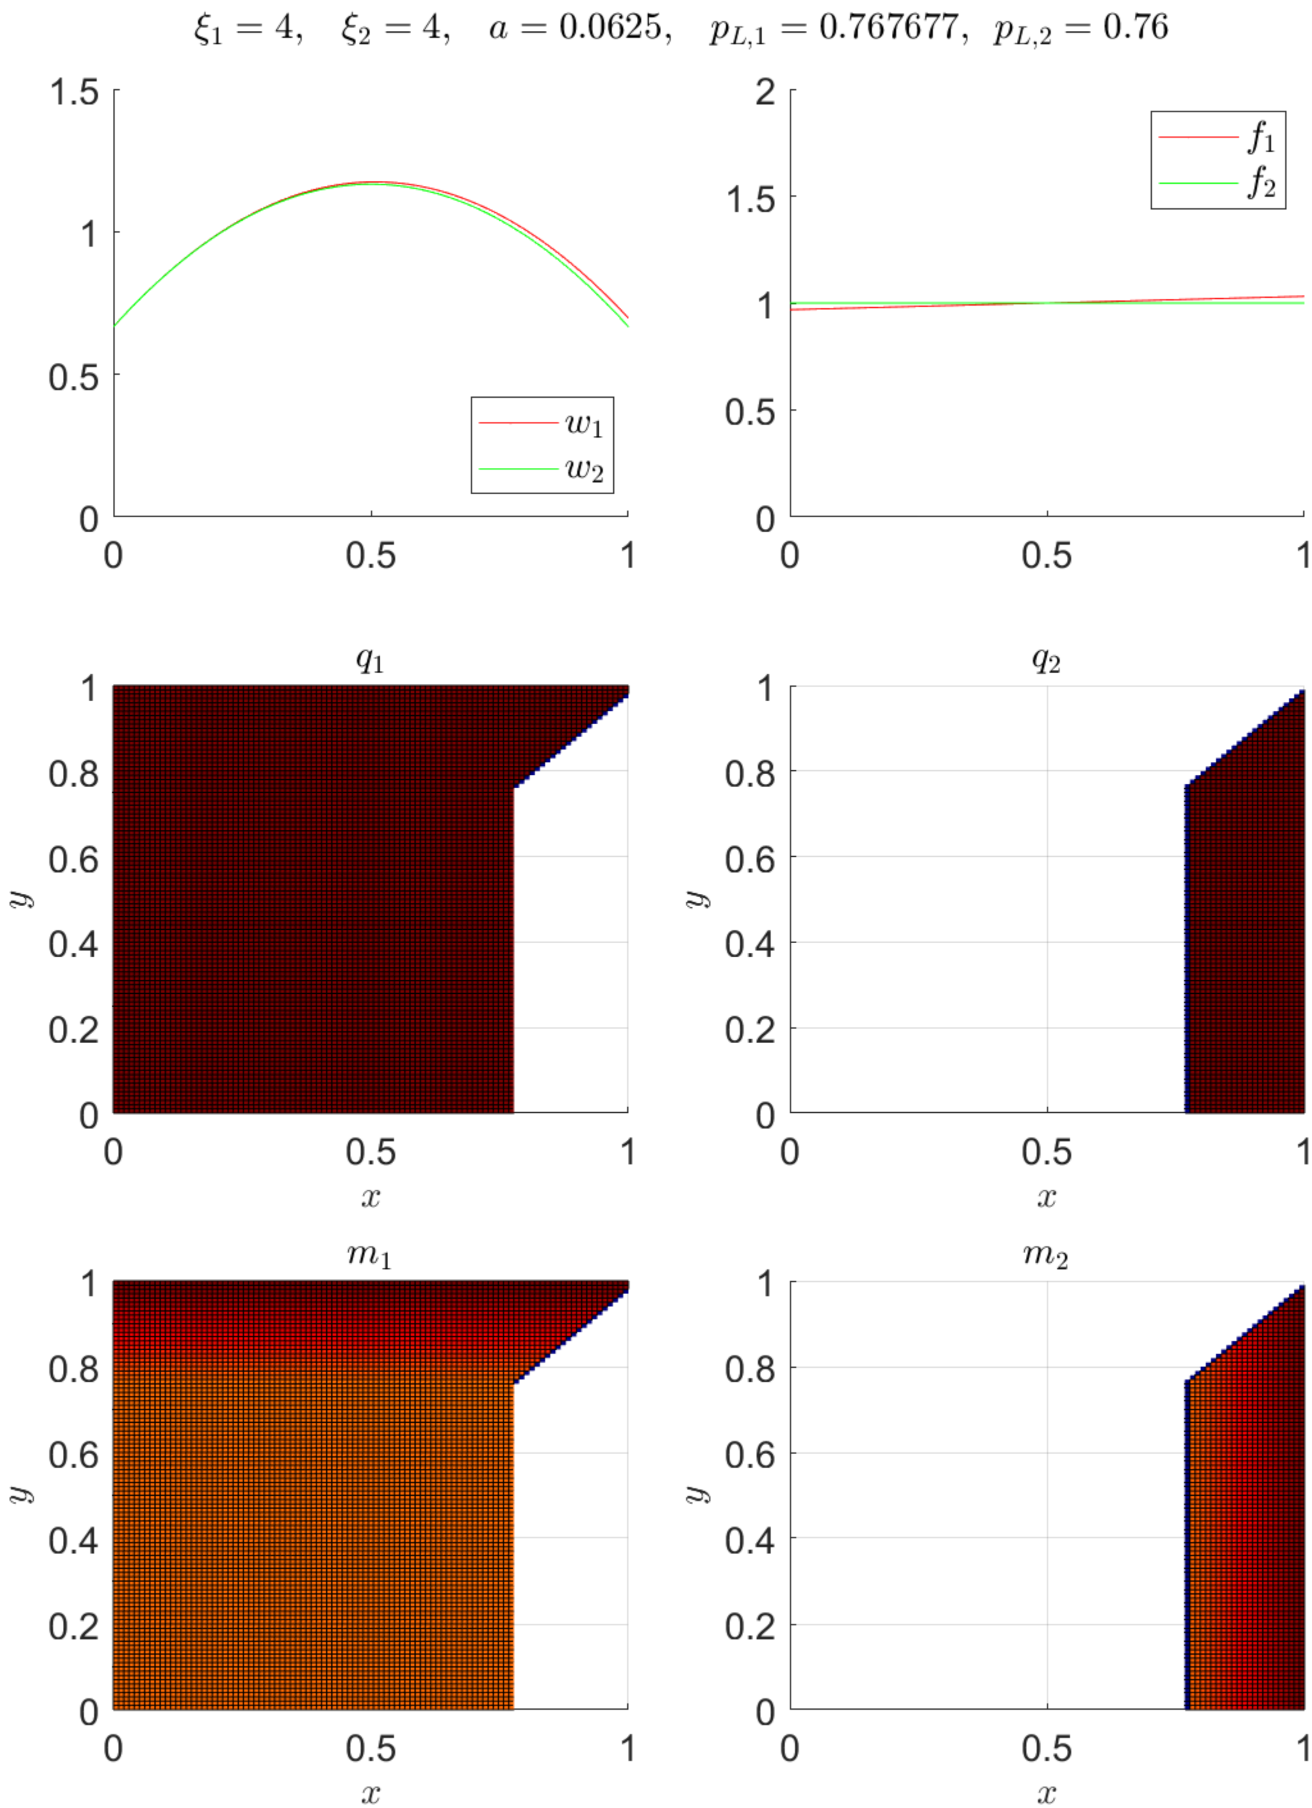
$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.0625, \quad p_{L,1} = 0.767677, \quad p_{L,2} = 0.76$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

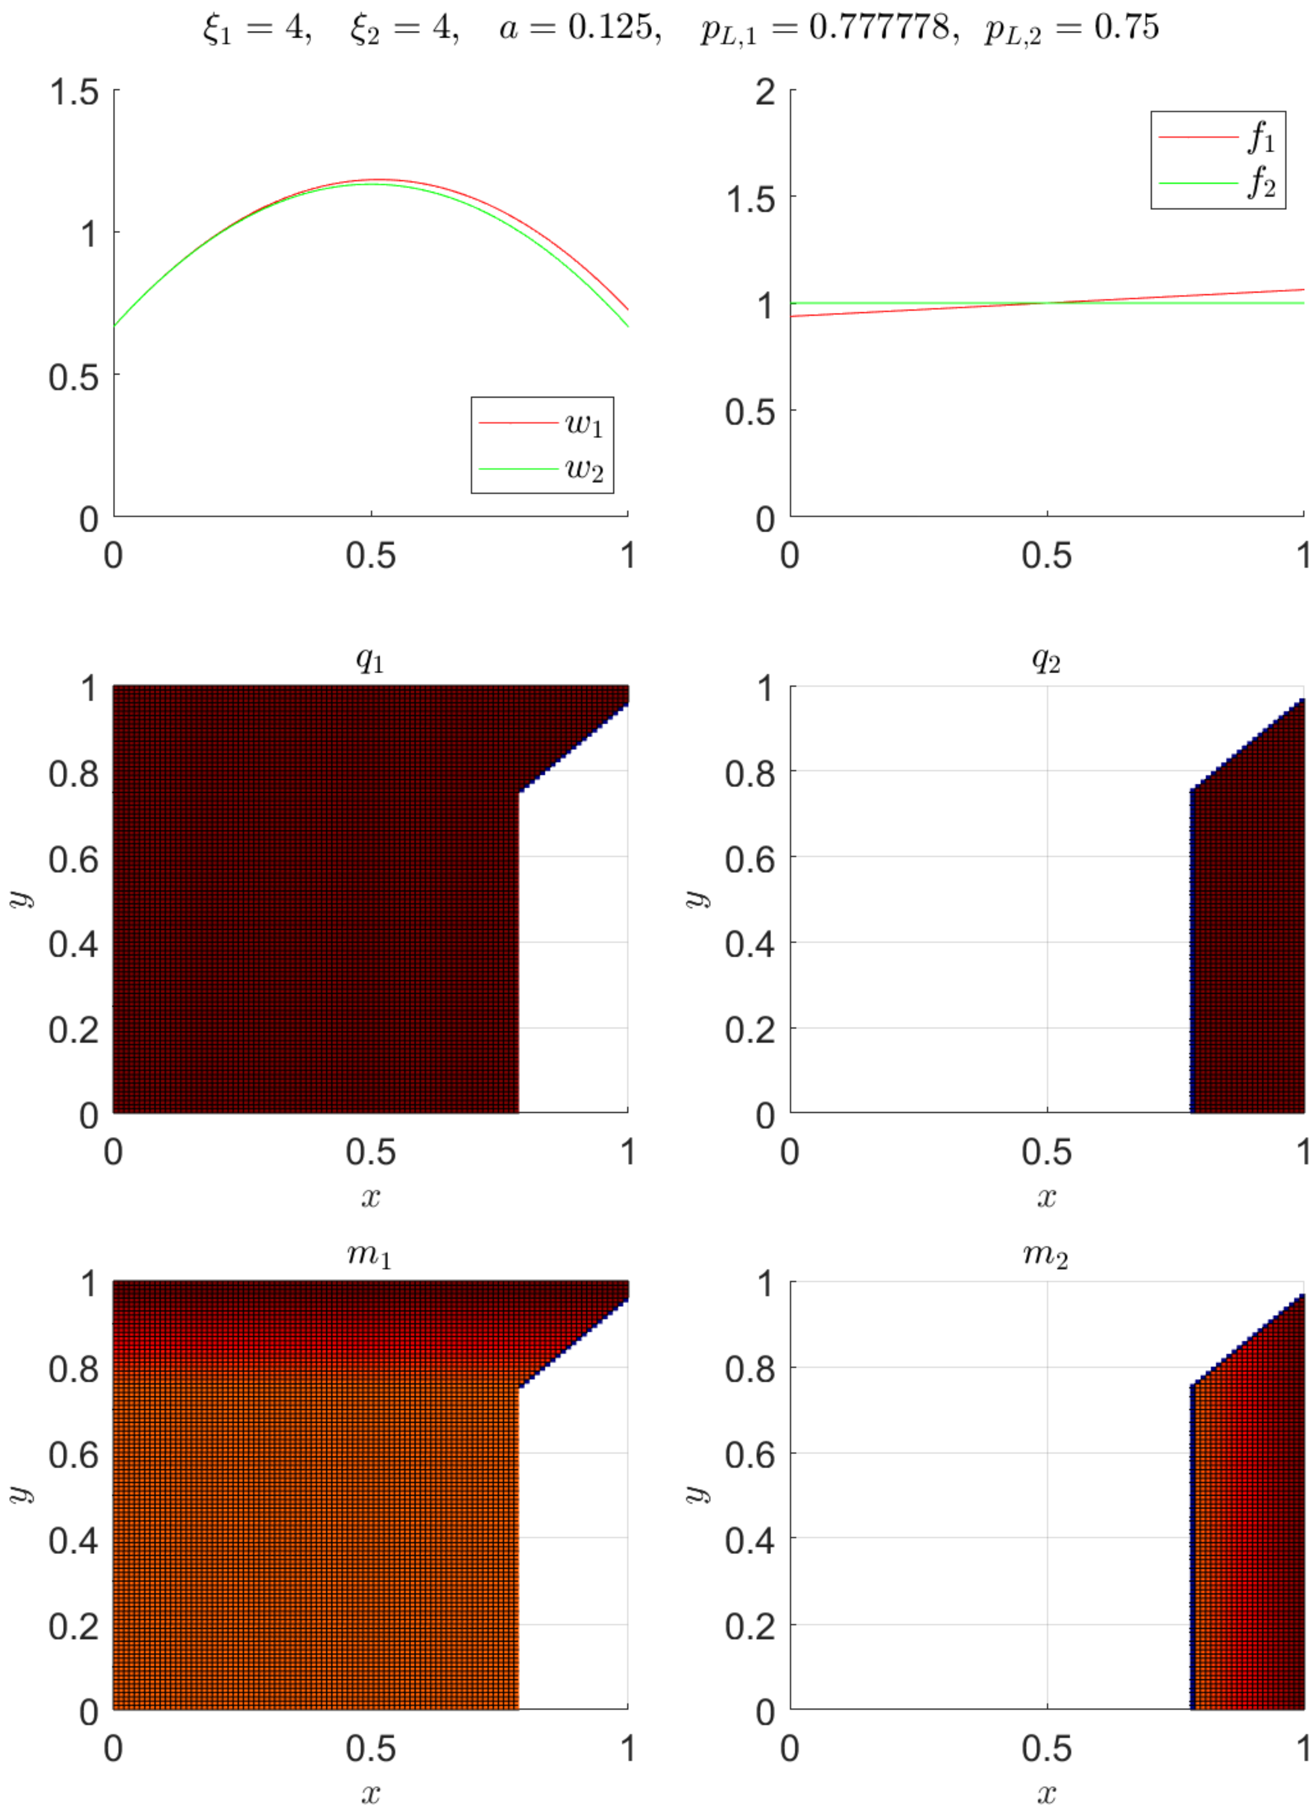

$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.125, \quad p_{L,1} = 0.777778, \; p_{L,2} = 0.75$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

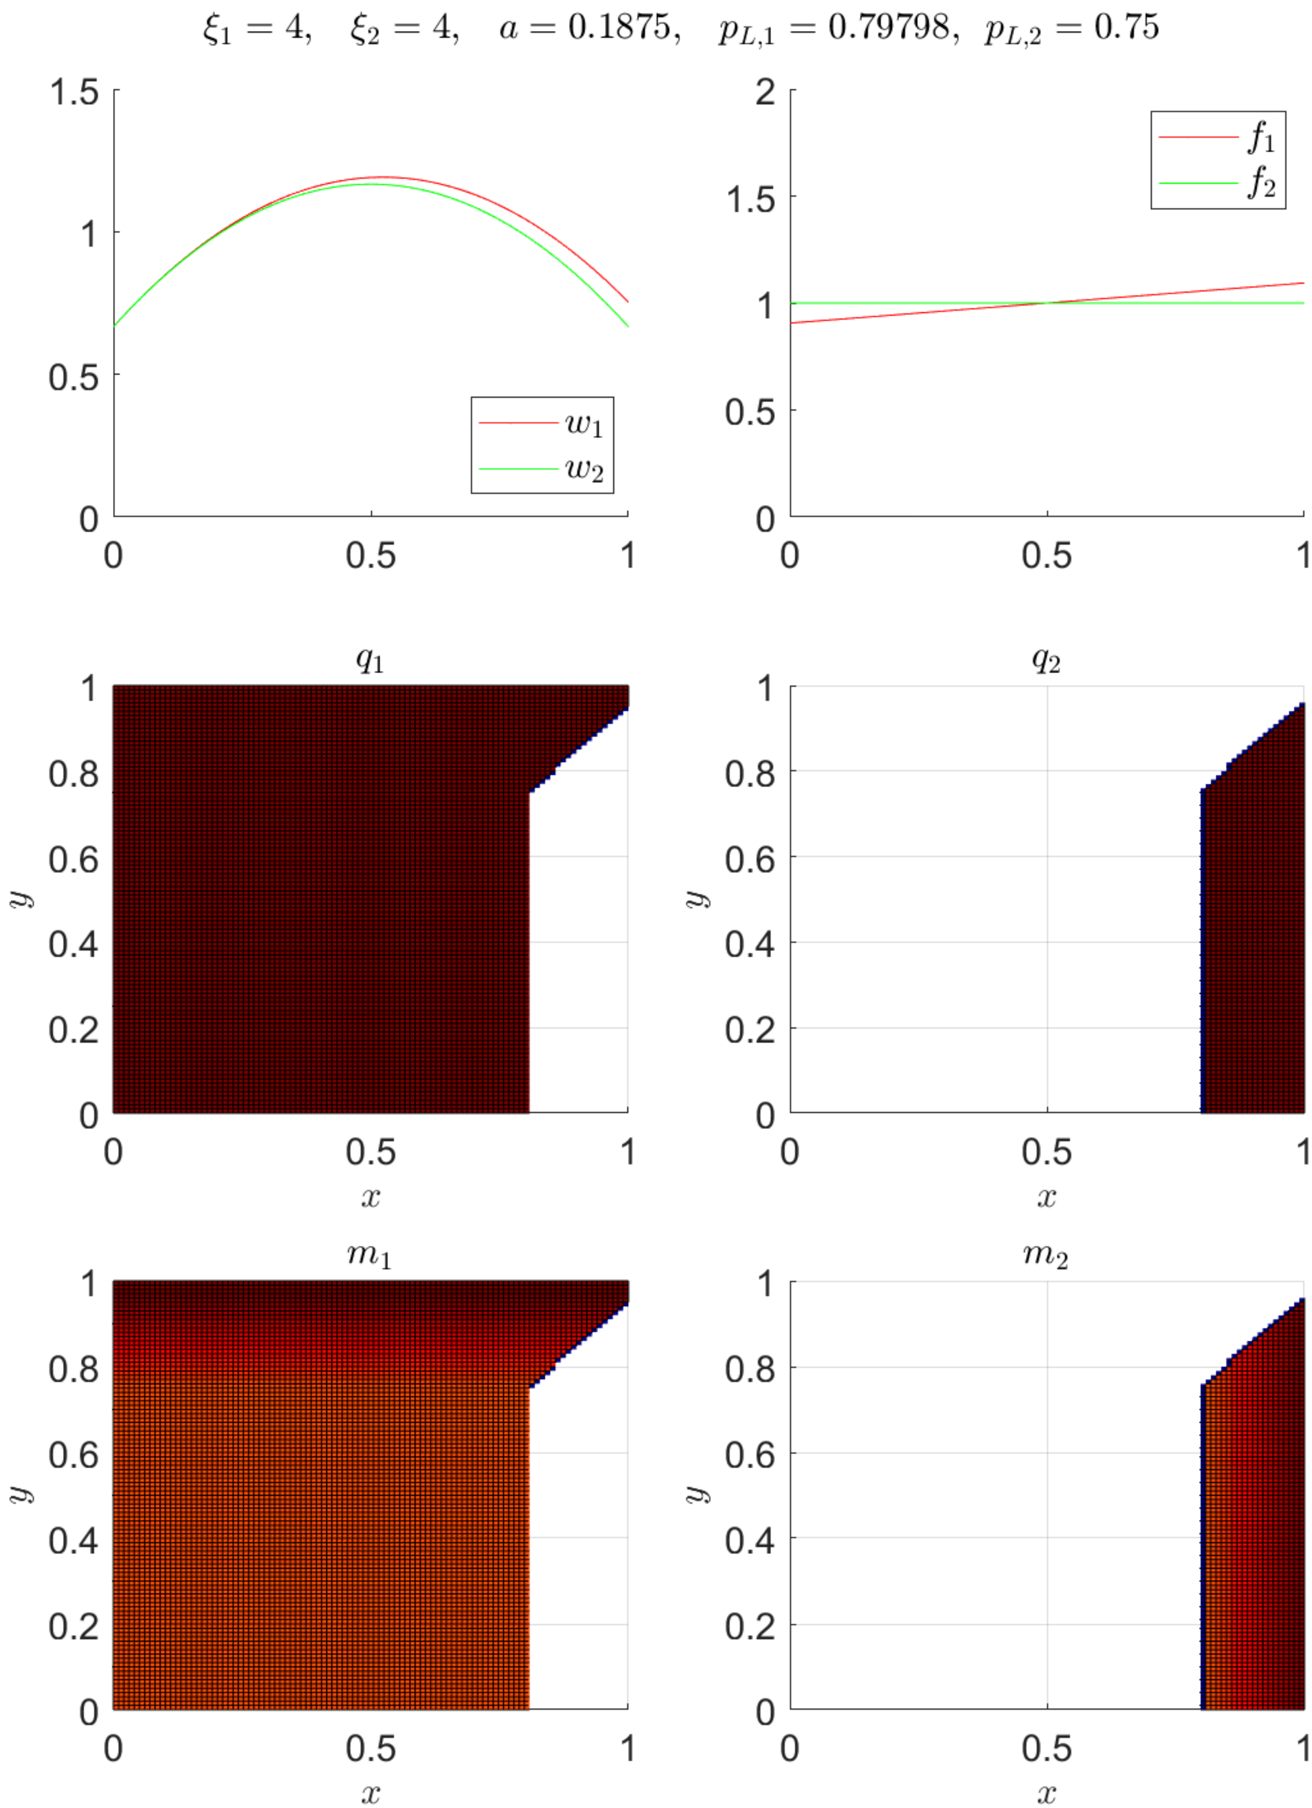

$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.1875, \quad p_{L,1} = 0.79798, \quad p_{L,2} = 0.75$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

$$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.25, \quad p_{L,1} = 0.818182, \quad p_{L,2} = 0.76$$

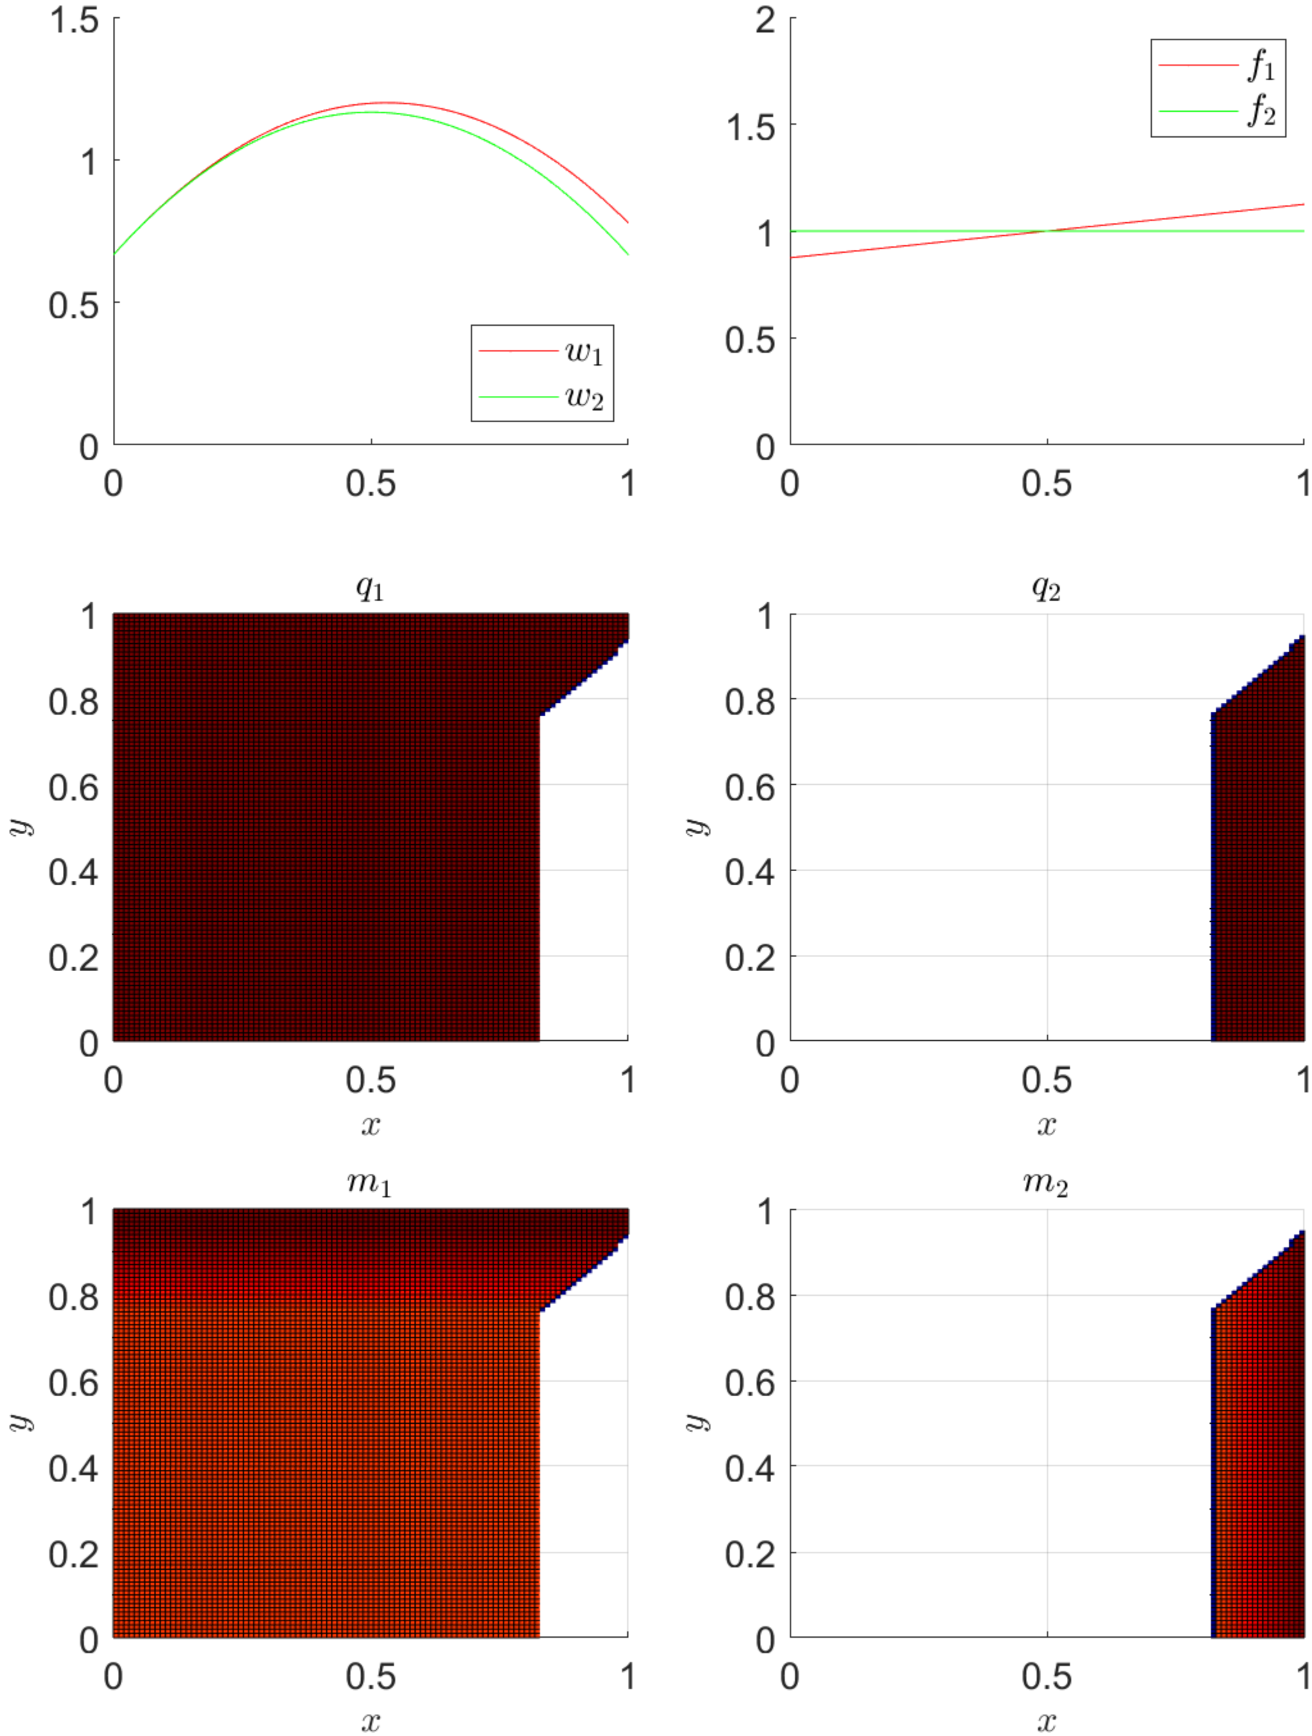

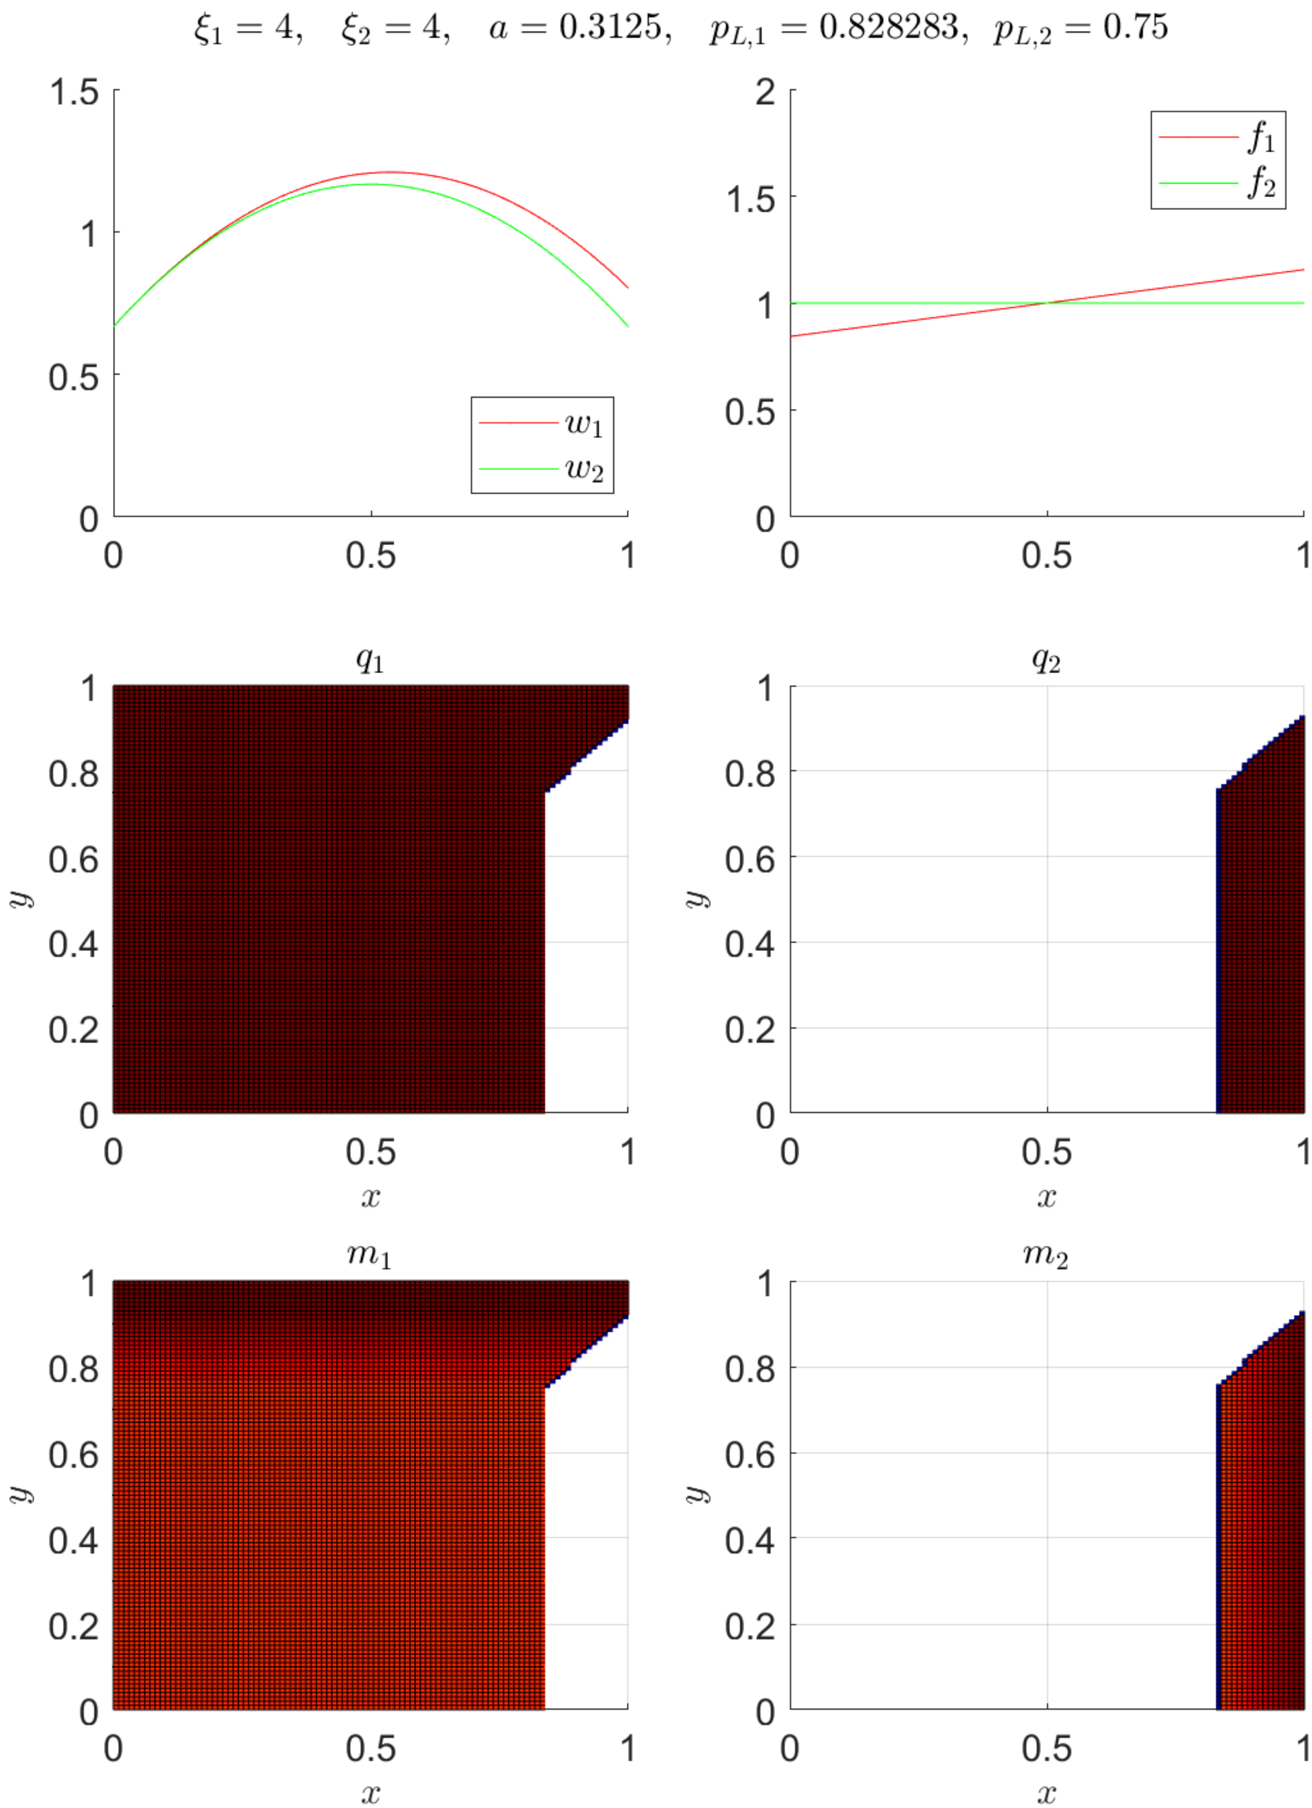

$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.3125, \quad p_{L,1} = 0.828283, \ \ p_{L,2} = 0.75$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

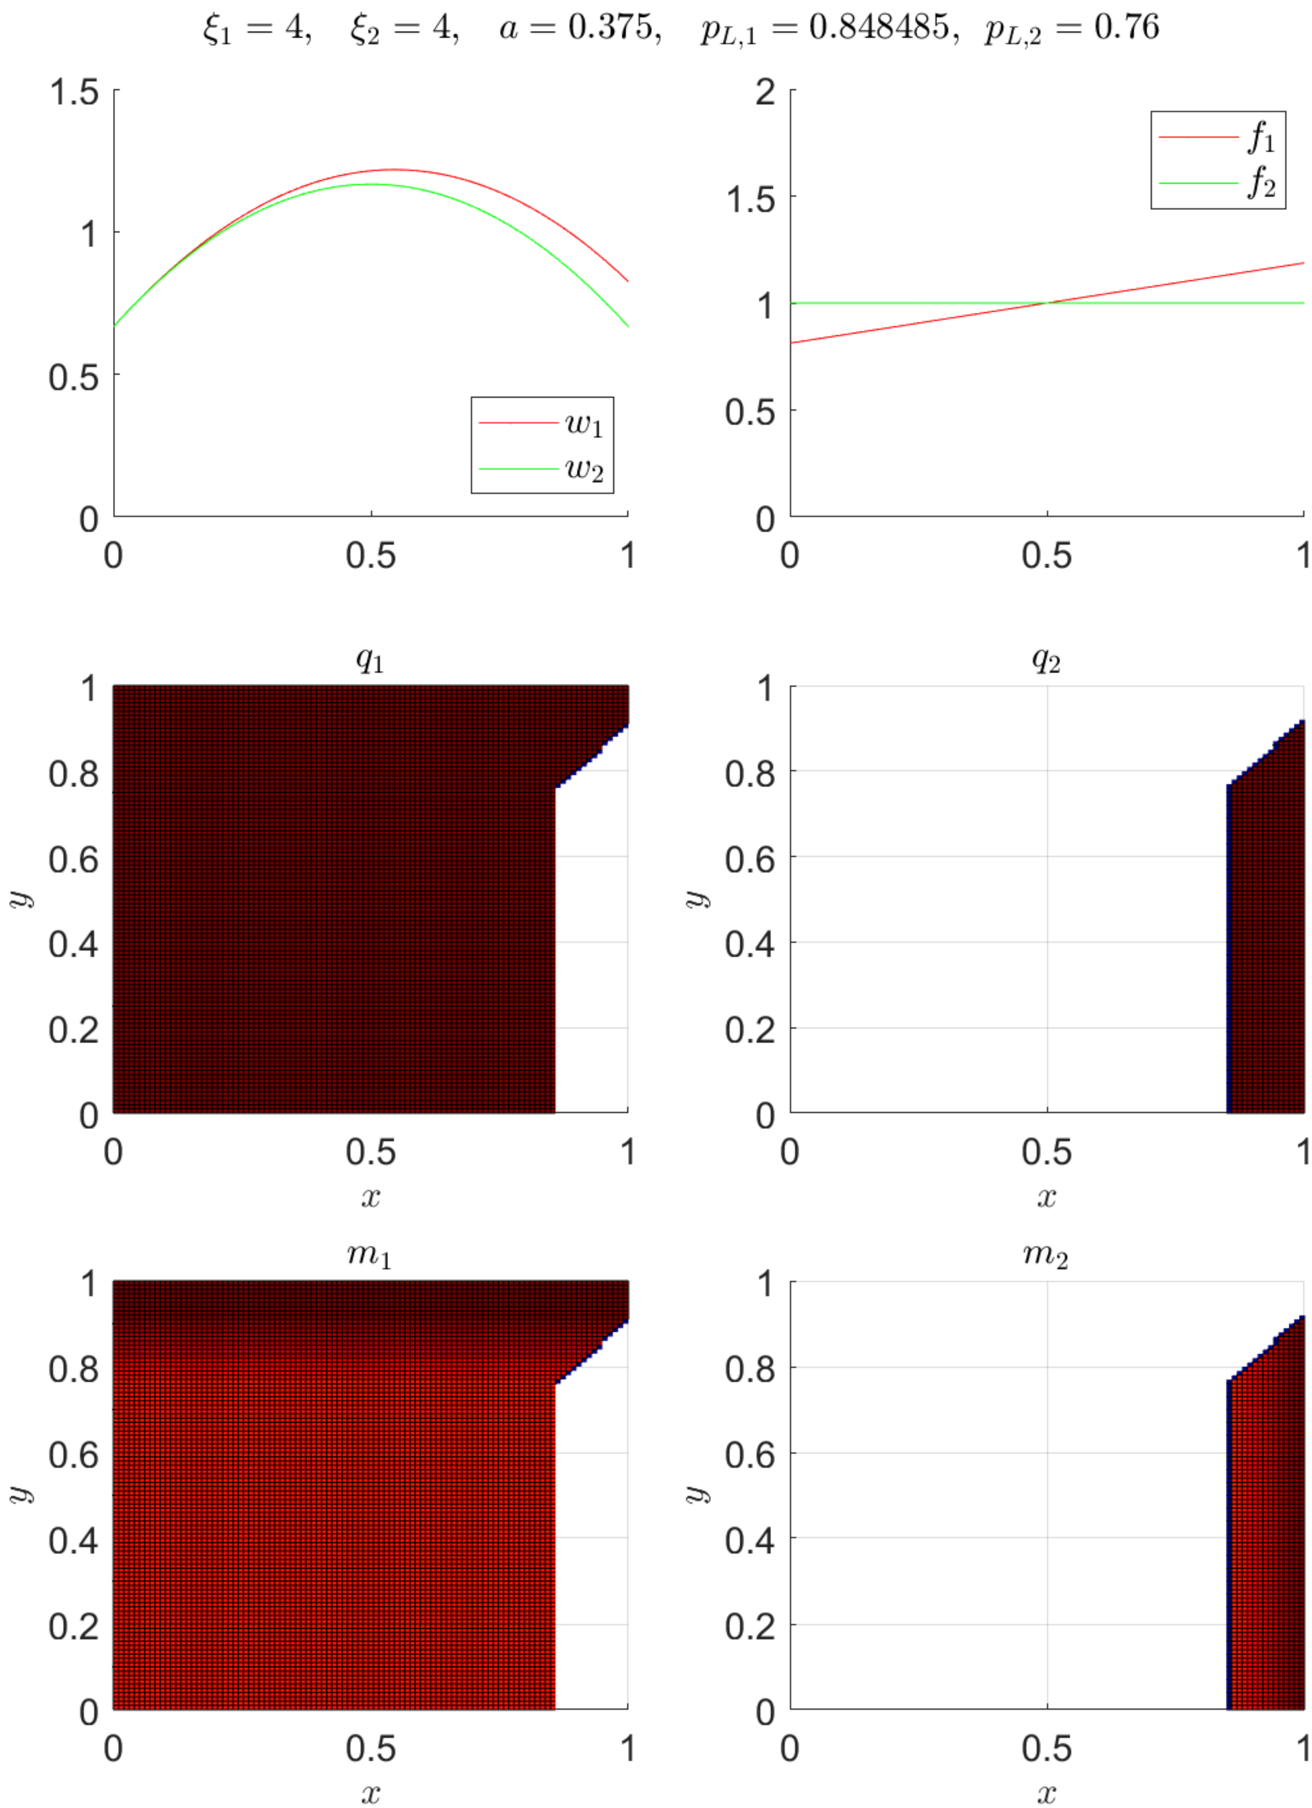

$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.375, \quad p_{L,1} = 0.848485, \quad p_{L,2} = 0.76$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

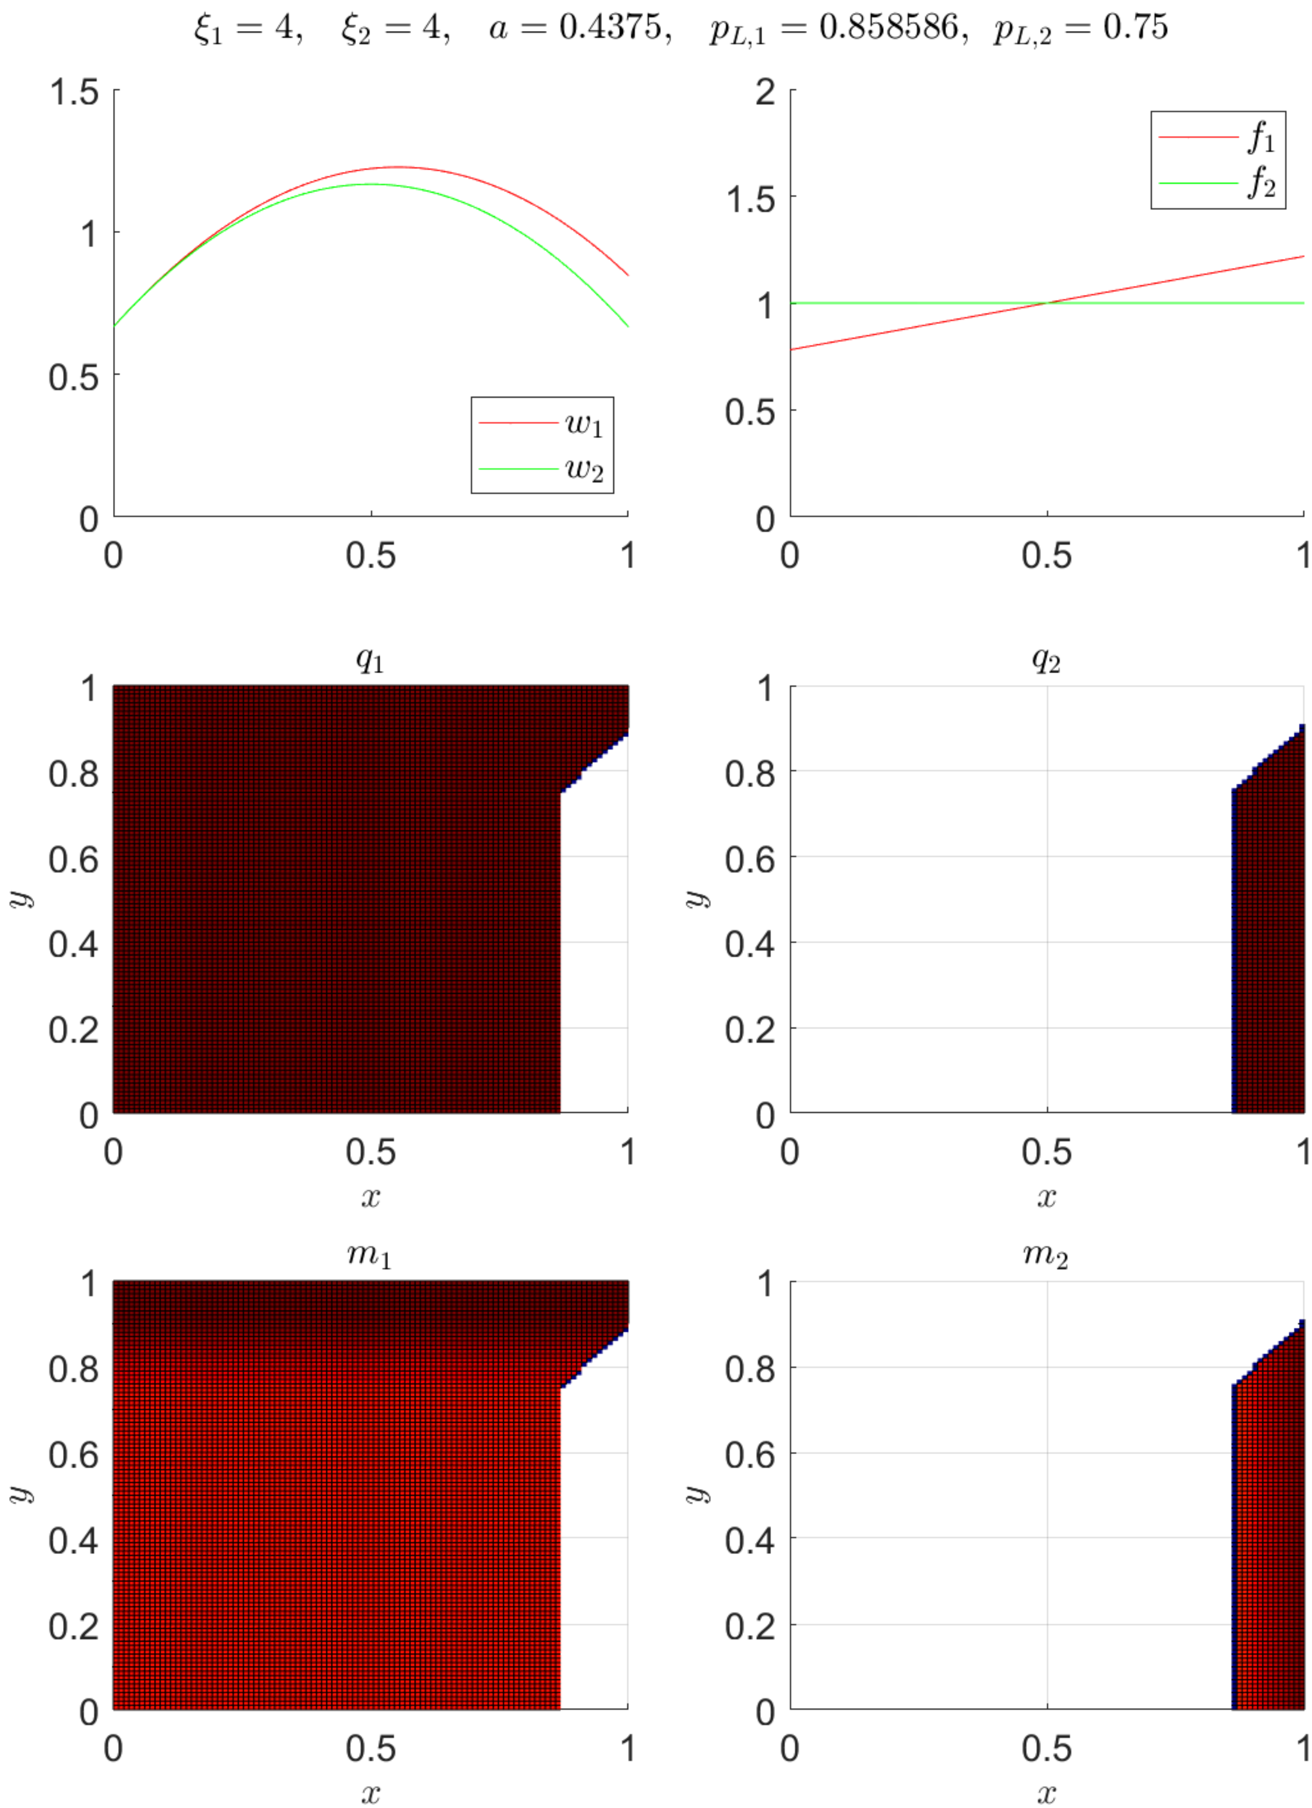
$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.4375, \quad p_{L,1} = 0.858586, \ \ p_{L,2} = 0.75$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

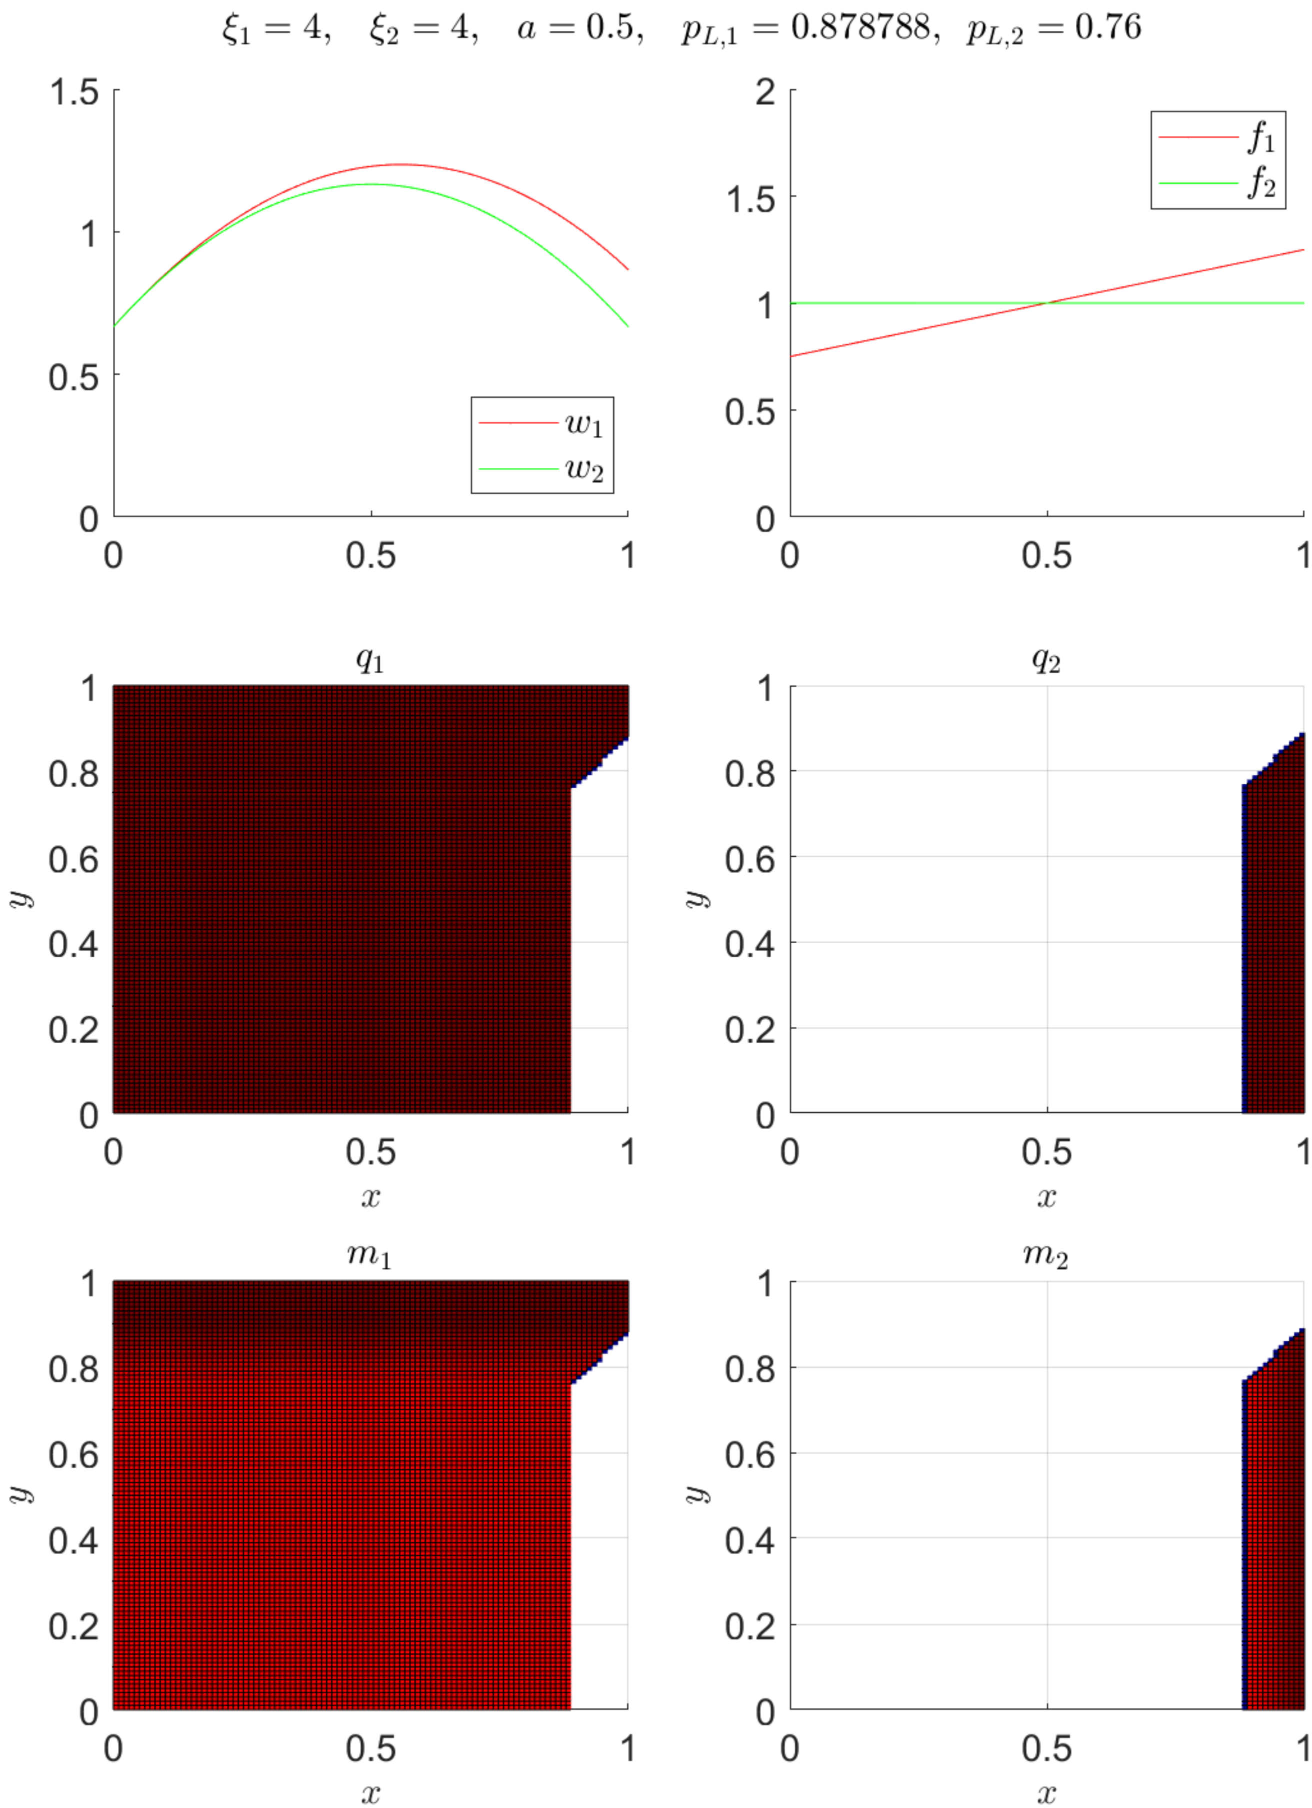
$\xi_1 = 4, \quad \xi_2 = 4, \quad a = 0.5, \quad p_{L,1} = 0.878788, \quad p_{L,2} = 0.76$
w_1
w_2
f_1
f_2
q_1
q_2
m_1
m_2
x
y

D Material for Section 6 (NOT FOR PUBLICATION)

D.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 5. 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 3. ***(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 5.

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 5, 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 7. ***(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.

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

Since

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

then:

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

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 5. 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:

$$\begin{aligned}
\hat{v}(c,\xi) &= \mathit{const} - K\mathbb{E}\left[\hat{D}(c,\xi)+\hat{O}(c,\xi)\right] \\
&= \mathit{const} - (1-\xi)K\mathbb{E}\left[D+O\right] - \xi K\left[\delta(c)+\omega(c)\right] \\
&= \mathit{const}(\xi) - \xi K\left[\delta(c)+\omega(c)\right]. \qquad (59)
\end{aligned}$$

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

From expression (19), the calibrated buyer's payoff reads:

$$\hat{v}(c,\xi) = const(\xi) - \xi K \left[G_D^{-1}\left(\left[1 - \hat{F}(c) \right]^N \right) + G_O^{-1}\left(\left[1 - \hat{F}(c) \right]^N \right) \right]. \tag{60}$$

Our goal is to fully calibrate this function of (c, ξ). The constant $const(\xi)$ reads, from (59):

$$const(\xi) = const - (1-\xi) K \mathbb{E}\left[D + O \right]. \tag{61}$$

We set *const* large enough that the virtual valuation $\hat{w}$ is everywhere positive,[30] and K large enough that, as ξ varies between 0 and 1, the slope of the social welfare function (dashed red line in Figure 6) 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 5, 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 6) yields variations in social surplus (dashed red line in Figure 6) 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 5.

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).

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

Log-concavity of $\hat{F}$

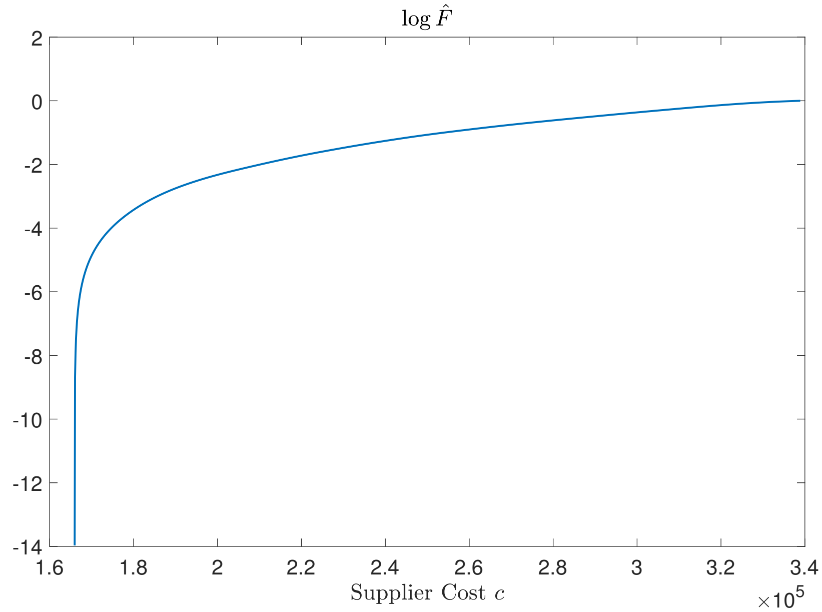

Figure 8: $\log(\hat{F})$ is concave.

E Software applications (NOT FOR PUBLICATION)

This appendix describes two software applications that we have created and made publicly available.[31] 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.

E.1 Software 1

This software is a visually handy procedure realized in Matlab that does not require IBM ILOG CPLEX. 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 that bidder costs are drawn independently from the cost distribution, and requires that $v(c_L) > c_L$. The application's output 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 6). 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.

The user specifies three inputs in an excel spreadsheet called "Input.xlsx", as shown in Figure 9 (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).

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

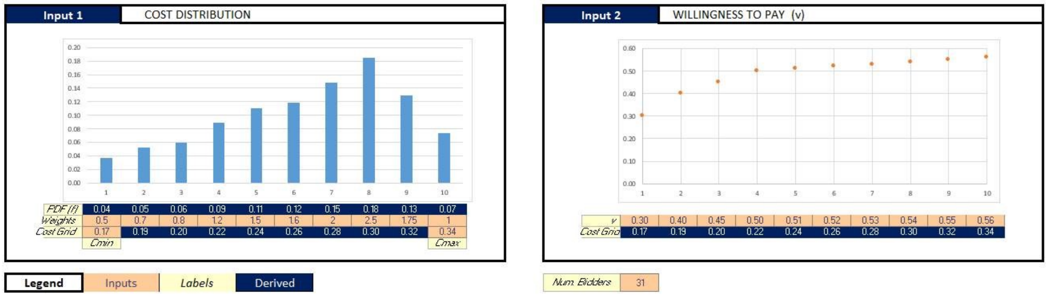

Figure 9: The figure shows the inputs of the visual program that solves for the optimal LoLA among all LoLAs.

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

$$\begin{cases} w_i = v_i - c_i, \quad i = 1 \\ w_i = v_i - c_i - (c_i - c_{i-1}) \cdot \frac{F_i}{f_i}, \quad \forall i > 1 \end{cases} \tag{62}$$

The result for w is showed to the user as in figure 11. The user is asked to check whether w is single-peaked in accordance to assumption 1.

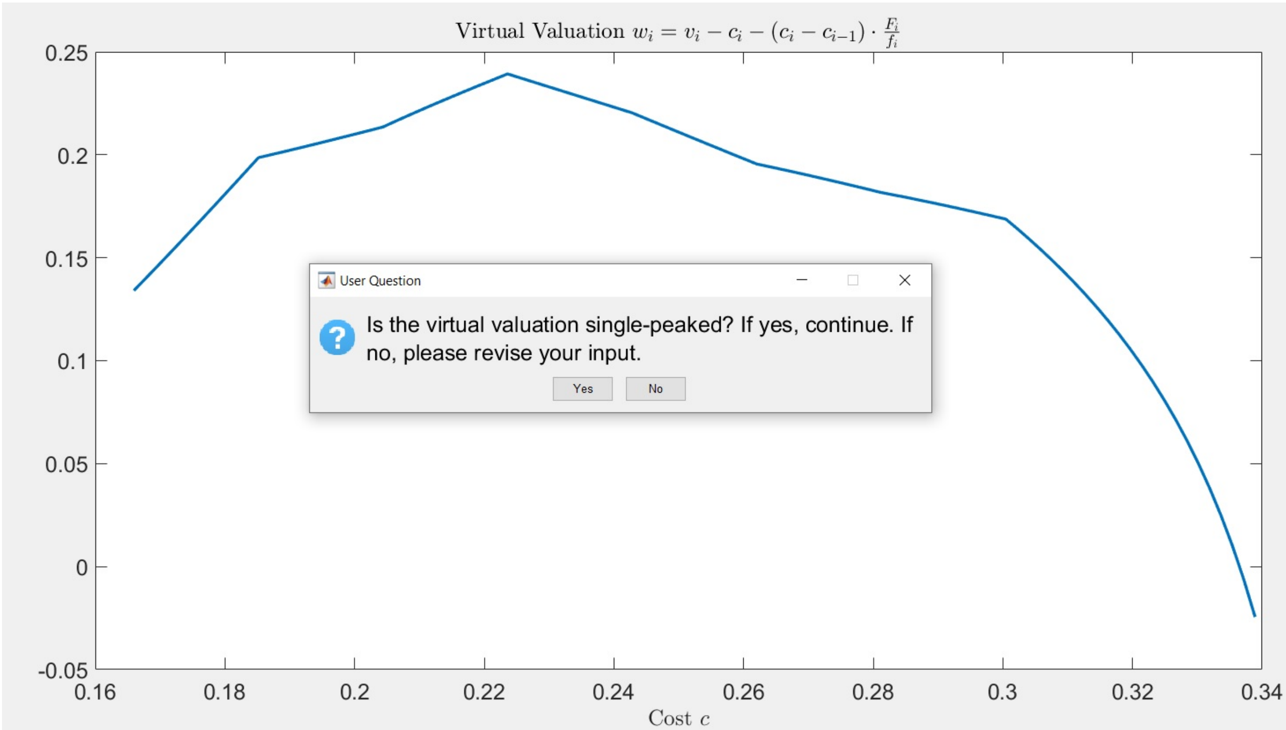

Figure 10: The figure shows w and it involves the user's participation by asking whether or not w is single-peaked.

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 12. Hence, the software asks for the user's confirmation to set the root of w as a reservation price p_H.

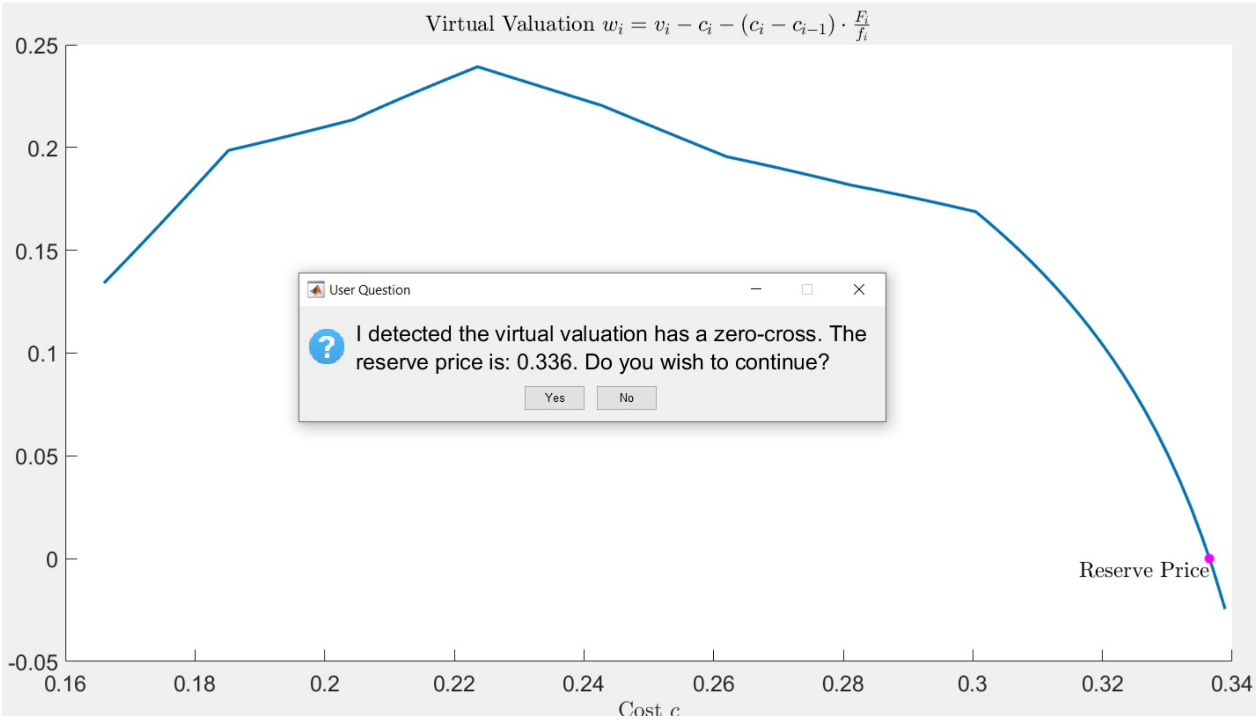

Figure 11: The figure shows the root of w calculated with a solver and using piece-wise linear interpolation on w. It involves the user's participation by acknowledging the root will be used as the reservation price.

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 28 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 12 and 13.

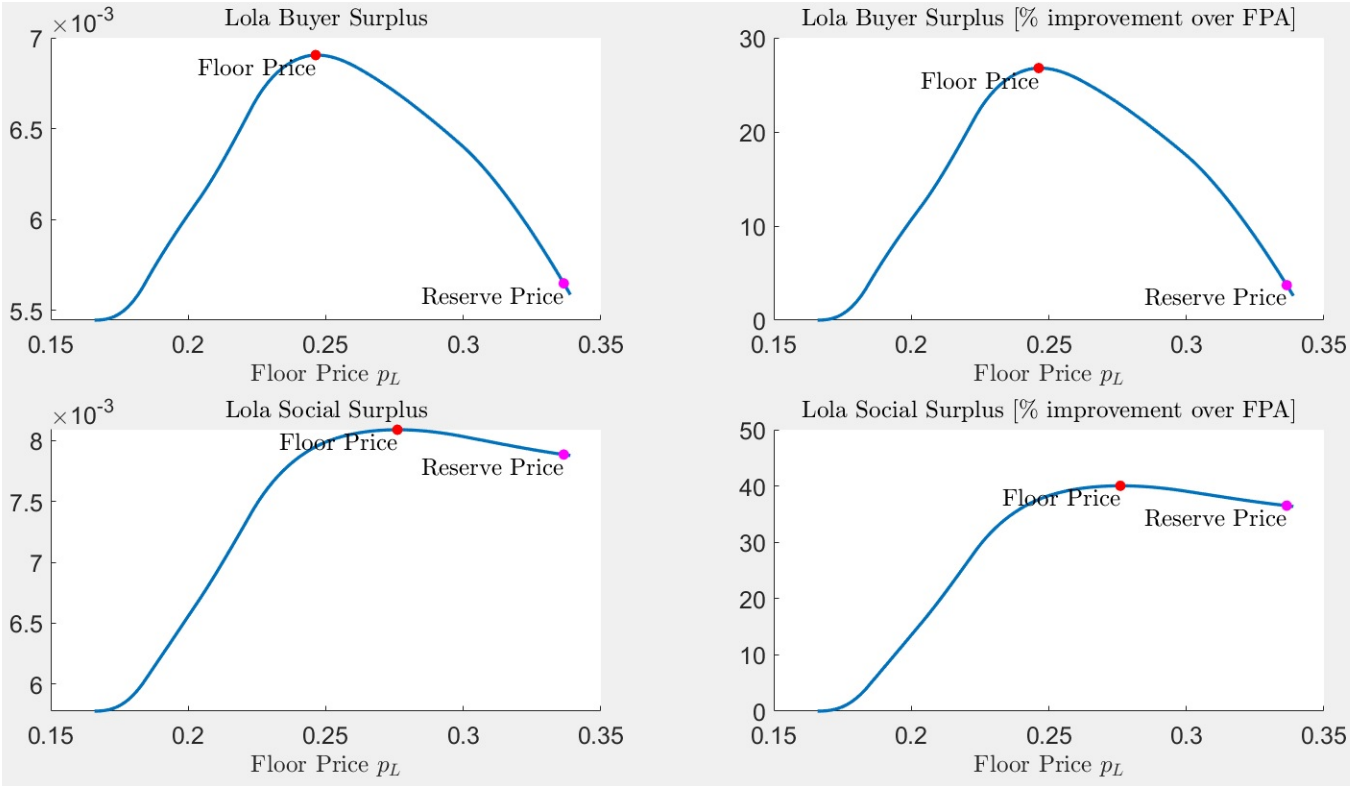

Figure 12: The figure shows the buyer surplus and social surplus in function of the floor price $p_{L,j}$. The points at which these functions are maximized correspond to the respective optimal LoLAs. In addition, the reservation price is also reported.

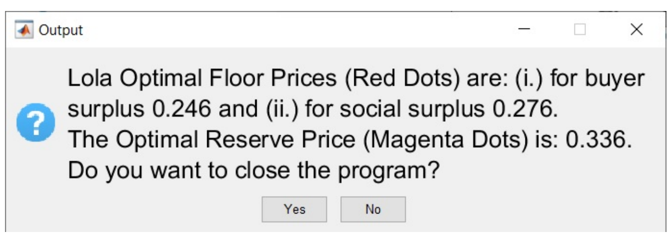

Figure 13: The figure shows the final report with the optimal floor and reservation prices.

E.2 Software 2

This software is realized in Matlab and IBM ILOG CPLEX. The application requires the same inputs as Software 1, and it computes the optimal mechanism *even when that mechanism is not a LoLA*. Therefore, Software 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 the paper are violated, and so Theorem 1 does not apply.

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 (62). 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.[32] The DAT's file specifies all numerical inputs.

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

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