

All-receive procurement auctions*

Simon Loertscher[†] Leslie M. Marx[‡] Patrick Rey[§]

March 25, 2025

Abstract

We develop the procurement analogue to an all-pay auction for an independent private values model with identical distributions. In this *all-receive procurement auction* (ARPA), suppliers simultaneously submit bids. Suppliers with bids below (above) the reserve are paid their bids (are paid and produce nothing). The supplier with the largest bid below the reserve produces the good. With appropriately chosen reserves, which decrease in the number of suppliers, the ARPA is efficient and, given increasing virtual costs, implements the optimal procurement. Appropriately adjusted, ARPAs implement the optimal procurement in general. ARPAs can render supply chains resilient to nonanticipated liquidity shocks.

Keywords: Resilience, liquidity shocks, all-pay auctions, multiple-receive procurement auctions

JEL Classification: D44, D82, L41

*We thank Wouter Dessein, Pierre Dubois, and David Martimort. Financial support from the Australian Research Council Discovery Project Grants DP200103574 and DP250100407, European Research Council Horizon Programme Grant Agreement No. 101055280 DIGECON, and ANR Grant ANR-17-EURE-0010 (Investissements d’Avenir program) is also gratefully acknowledged. Toan Le and Hasini Senadheera provided excellent research assistance.

[†]University of Melbourne

[‡]Duke University

[§]Toulouse School of Economics, University of Toulouse Capitole

1 Introduction

The pandemic and ongoing geopolitical tensions have led policy makers and business leaders to call for more *resilient* supply chains. Typically, reducing reliance on a single, low-cost, but often distant, supplier involves trading off higher costs against a lower risk that the supply chain breaks down. This raises the question of how sustainable efforts to increase domestic production and on-shoring will be. To address this question, we develop procurement auction formats that involve no tradeoff in terms of expected payoffs, yet increase the resilience of the supply chain. Of particular interest is the procurement analogue of the all-pay auction, which we call *all-receive procurement auction* (ARPA).

We study a standard procurement model with independent private values in which a buyer with known value wants to procure one unit from a set of suppliers who draw their costs independently from a common, known distribution. In the ARPA, suppliers submit bids simultaneously. *All* suppliers with bids not exceeding the reserve set by the procurer are paid their bids, and the supplier with the *largest* bid below the reserve is asked to produce the good. Suppliers with bids above the reserve are paid nothing and are never asked to produce. Like first-price and second-price procurement auctions (FPPA and SPPA), with appropriately chosen reserves, the ARPA is efficient and, if the cost distribution exhibits increasing virtual types, an optimal procurement auction. By the payoff equivalence theorem, the ARPA, FPPA and SPPA thus induce the same interim and thus the same ex ante expected payoffs for the suppliers and the same expected payoff for the buyer.

Another feature of the ARPA is that the buyer’s total payment is spread across all suppliers with bids below the reserve. Consequently, relative to other two formats, the ARPA increases suppliers’ liquidity. As we show by considering liquidity shocks, occurring at the end of the game and not anticipated in the auction, this has the advantage of allowing more suppliers to survive a liquidity shock than the FPPA and SPPA. Specifically, there are conditions under which a liquidity shock eliminates all suppliers in the FPPA and SPPA, while in the ARPA it only eliminates the supplier that is selected to produce. Under other conditions, either all or all but the selected producer survive in the ARPA and only the winner in the FPPA and SPPA survives the liquidity shock. With the ARPA, the buyer can thus make the supply chain more resilient—and in particular preserve competition among suppliers—without any sacrifice in expected payoffs.

We moreover show that, with appropriate adjustments of the rules to allow ironing (Myerson, 1981), the payoff equivalence and resilience properties of the ARPA extend to the case in which the virtual cost is not globally monotone; in the process, we also show how to adjust the rules of the FPPA and SPPA to permit ironing. However, the ARPA and its

adjusted variant may degrade the liquidity of the supplier selected to produce. A generalization of the ARPA that mitigates this concern, at the expense of some resilience, would be multiple-receive procurement auctions, in which only a subset of suppliers are paid their bids (with the ARPA and the FPPA being extremal special cases). Interestingly, as we show in an extension, multiple-receive procurement auctions do not necessarily exhibit a monotone equilibrium and can therefore fail to be payoff equivalent to more standard formats. In further extensions, we derive the direct mechanism that is maximally resilient without ever degrading suppliers' liquidity, and address possible concerns about the susceptibility of the ARPA to collusion—showing in particular that the ARPA can actually be less susceptible to collusion than the other formats.

The auction literature is, evidently, too vast to review here; see, for example, Krishna (2010) for an authoritative overview. Revenue equivalence between different auction formats was first observed by Vickrey (1961, 1962) with the driving force behind it, and the deeper result of payoff equivalence, uncovered by Riley and Samuelson (1981) and Myerson (1981). We develop the procurement equivalent to the all-pay auction, the ARPA, and show that, with an appropriate reserve, it is payoff equivalent to standard formats like the FPPA and SPPA. While we do *not* analyze the design of (optimal) mechanisms in the presence of liquidity constraints, which is the subject of papers such as Ganuza (2007) and Arve and Martimort (2024), we analyze the robustness of fixed and otherwise equivalent auction formats to unanticipated liquidity shocks.

In what follows, in Section 2, we introduce the setup and the formal definition of the ARPA. In Section 3, we derive the suppliers' equilibrium behavior in the ARPA. Implementation of the optimal procurement mechanism via the ARPA is studied in Section 4. Section 5 introduces and analyzes liquidity shocks, and Section 6 provides additional discussion and concludes the paper.

2 Setup

There is one buyer, with single-unit demand and value $v > \underline{c}$, and $n \geq 2$ suppliers, each drawing its cost independently from a continuous distribution F with density f and support $\mathcal{C} \equiv [\underline{c}, \bar{c}]$.¹ The buyer and the suppliers are all risk neutral; they have quasilinear utility, and the value of their outside option is 0.

In an *all-receive procurement auction* (ARPA) with reserve R , all suppliers simultaneously

¹A supplier's production cost c can be interpreted either as the opportunity cost of selling an existing asset or as the financial cost of producing an asset or service. For the analysis in Sections 3 and 4, these two interpretations are equivalent. However, for some of the analysis and results in Section 5 on resilience to liquidity shocks, which specification is used matters.

submit bids.² If all bids (strictly) exceed R , the procurement is unsuccessful and there are no payments or production. Otherwise, each supplier i with a bid $B_i \leq R$ is paid B_i , and the supplier with the highest bid produces. If ties occur among the highest bidders below R , then the supplier that has to produce is selected uniformly randomly from among these bidders. This format inherits from the all-pay auction that each agent's transaction price is what it bids. There is however a notable difference between the sale and the procurement contexts. For sale auctions, all-pay and first-price formats have in common that the highest bidder wins. In the procurement context, the ARPA selects the *highest* bidder as supplier, whereas a FPPA or SPPA selects instead the *lowest* bidder. Because incentive compatibility requires an agent's selection probability to be decreasing in its cost, this in turn implies that, in the ARPA, equilibrium bids are *decreasing* in costs.

3 Suppliers' response

We first show that the ARPA elicits a unique symmetric response from the suppliers. For any $c \in \mathcal{C}$, let

$$P_n(c) \equiv [1 - F(c)]^{n-1}$$

denote the probability that $n - 1$ suppliers have costs higher than c and, for any $r \in (\underline{c}, \bar{c}]$ and $c \in [\underline{c}, r]$, define

$$B_n(c; r) \equiv \mathbb{E}_{\mathbf{c}_{-i}}[\min\{\mathbf{c}_{-i}, r\} \mid \min \mathbf{c}_{-i} > c] P_n(c) = c P_n(c) + \int_c^r P_n(x) dx, \quad (1)$$

and

$$R_n(r) \equiv \mathbb{E}_c[\min\{c, r\}] = B_n(\underline{c}; r) = \underline{c} + \int_{\underline{c}}^r P_n(x) dx. \quad (2)$$

Observe that $R_n(r) < r$ because $\underline{c} + \int_{\underline{c}}^r [1 - F(x)]^{n-1} dx < \underline{c} + \int_{\underline{c}}^r dx = r$ and that $R_n(r)$ is decreasing in n because the integral decreases in n .

We have:

Proposition 1. *For any $r \in (\underline{c}, \bar{c}]$, the ARPA with reserve $R = R_n(r)$ induces a unique symmetric Bayesian Nash equilibrium, in which every supplier with cost $c \leq r$ bids $B_n(c; r)$ and every supplier with cost $c > r$ refrains from bidding. This equilibrium replicates the allocation of the FPPA and SPPA with reserve r .*

²Without a reserve, the ARPA would have no equilibrium, as any bidder would find it optimal to submit an infinite bid. But this is not different from an all-pay auction in which, without a reserve (typically taken to be zero), bidders would optimally submit an infinitely negative bid.

Proof. Take $r \in (\underline{c}, \bar{c}]$ as given and consider an ARPA with reserve $R_n(r)$. Standard arguments imply that in any symmetric equilibrium, (i) suppliers produce with positive probability if and only if their cost does not exceed a cost threshold, which we denote by $\hat{r} \in [\underline{c}, \bar{c}]$, and (ii) for a supplier with cost $c \leq \hat{r}$, the probability of being selected is $P_n(c)$ and the bid function $B(c)$ is continuous and differentiable. For $c \leq \hat{r}$, we thus have, by incentive compatibility, $c = \arg \max_{\hat{c}} B(\hat{c}) - cP_n(\hat{c})$. The first-order condition, evaluated at $c = \hat{c}$, yields $0 = B'(c) - cP'_n(c)$, leading to

$$B(c) = k + cP_n(c) + \int_c^{\hat{r}} P_n(x)dx = k + B_n(c; r),$$

for some constant k . Furthermore, because the lowest type produces with probability 1, it cannot be induced to submit any other bid than the maximal one, $R_n(r)$. Thus, $B(\underline{c}) = R_n(r) = B_n(\underline{c}; r)$, implying that $k = 0$. The threshold type \hat{r} must moreover be indifferent between bidding and not, $B_n(\hat{r}; r) - \hat{r}P_n(\hat{r}) = 0$, which implies that $\hat{r} = r$. It follows that a supplier with cost $c \in [\underline{c}, r]$ submits a bid $B(c) = B_n(c; r) \leq R_n(r)$.

To establish sufficiency of the first-order condition, denote by $U(c, \hat{c})$ the utility of a type c that submits the bid $B_n(\hat{c}; r)$. We have:

$$U(c, \hat{c}) = B_n(\hat{c}; r) - cP_n(\hat{c})$$

and

$$\frac{\partial U(c, \hat{c})}{\partial \hat{c}} = (\hat{c} - c)P'_n(\hat{c}),$$

which is positive for $\hat{c} < c$ and negative for $\hat{c} > c$ because $P'_n < 0$. Thus, the first-order condition characterizes the unique maximum.

To conclude the proof, recall that a FPPA or SPPA with reserve r induces suppliers with costs less than or equal to r to submit bids that are increasing in their types and less than or equal to r , ensuring that the lowest-type supplier produces whenever its cost is less than or equal to r , and no suppliers produce when the lowest type is greater than r , generating the same allocation as the ARPA with reserve $R_n(r)$. ■

The equilibrium bidding strategy is *decreasing* in the supplier's cost: $B'_n(c; r) = P'_n(c)c < 0$. This occurs because a higher-cost supplier is less likely to be selected, and is therefore willing to accept a lower (unconditional) payment. (For an illustration, see Supplemental Appendix Figure C.1(a).) Furthermore, a supplier with cost r bids $B_n(r; r) = P_n(r)r$ and thus obtains a zero payoff, as it is asked to produce with probability $P_n(r)$, in which case it

incurs the cost r . Let

$$U_n(c; r) \equiv B_n(c; r) - cP_n(c)$$

denote the resulting expected payoff for a supplier with cost $c \in [\underline{c}, r]$, which thus satisfies $U_n(r; r) = 0$. Because the ARPA bid function is decreasing in c , the highest bid is submitted by the supplier with the lowest cost: $\max_c B_n(c; r) = B_n(\underline{c}; r)$. Hence, the reserve R in the ARPA that corresponds to a reserve r in a FPPA or SPPA is $R = B_n(\underline{c}; r) = R_n(r)$, which satisfies $R_n(r) = r$ for $r = \underline{c}$ and $R_n(r) \in (\underline{c}, r)$ for $r \in (\underline{c}, \bar{c}]$.

Consider now the relation between bidding strategies. In a FPPA with reserve $r \in \mathcal{C}$, a supplier with cost $c \leq r$ bids

$$b_n^{FPPA}(c; r) \equiv c + \frac{\int_c^r P_n(x) dx}{P_n(c)}.$$

Consequently, we have

$$B_n(c; r) = P_n(c)b_n^{FPPA}(c; r). \quad (3)$$

This is the same relationship as between an all-pay and a first-price auction, where, as is well known, the bid in the all-pay auction is equal to the bid in the first-price auction multiplied by the probability of winning. For example, for F uniform on $[0, 1]$ and $r = 1$, we have $b_n^{FPPA}(c; 1) = c + \frac{1-c}{n}$ and $B_n(c; 1) = (1-c)^{n-1}(c + \frac{1-c}{n})$.

In particular, as a supplier with cost \underline{c} is selected with probability 1, it bids the same amount in the ARPA and the FPPA; hence, $R_n(r) = B_n(\underline{c}; r) = b_n^{FPPA}(\underline{c}; r)$. As the latter is decreasing in n , the same applies to $R_n(r)$ (for an illustration, see Figure C.1(b) in the Supplemental Appendix). Furthermore, as bids are decreasing in c in the ARPA and increasing in the FPPA, all bids in the FPPA weakly exceed all bids in the ARPA, as illustrated in Supplemental Appendix Figure C.2.

Finally, by construction, for any $r \in (\underline{c}, \bar{c}]$, $B_n(c; r)$ is continuous and decreasing in c and satisfies $B_n(\underline{c}; r) = R_n(r) > \underline{c}$ and $B_n(r; r) = P_n(r)r < r$. Hence, there exists a unique $\tilde{c}_n(r) \in (\underline{c}, r)$ such that

$$B_n(c; r) \gtrless c \text{ if and only if } c \lesseqgtr \tilde{c}_n(r). \quad (4)$$

This means that suppliers with costs below $\tilde{c}_n(r)$ make a profit even if called upon to produce, whereas suppliers with costs above $\tilde{c}_n(r)$ make a profit only if they do not have to produce (and make a loss otherwise). Because $B_n(c; r)$ is decreasing in n and increasing in r , $\tilde{c}_n(r)$ is also decreasing in n and increasing in r (see Supplemental Appendix Figure C.3).

4 Optimal procurement auctions

Let

$$\gamma(c) \equiv c + \frac{F(c)}{f(c)}.$$

denote the *virtual cost* of a supplier with cost c . We begin with the case in which γ is increasing (the *regular* case hereafter), before addressing the case in which it is not.

4.1 The regular case

When γ is increasing (e.g., when the reverse hazard rate F/f is itself increasing), we can define

$$r^* \equiv \gamma^{-1}(\min\{v, \bar{c}\}).$$

As is well known, in this regular case, the mechanism that maximizes the buyer's expected payoff (the optimal mechanism hereafter) selects a supplier with cost c with probability $P_n(c)$ if $c \leq r^*$, and probability 0 otherwise. Moreover, it can be implemented using an FPPA or SPPA with reserve r^* (see Supplemental Appendix A for details). It follows from Proposition 1 that it can also be implemented using an ARPA with an appropriate reserve:

Corollary 1. *If $\gamma(c)$ is increasing in c , then the optimal mechanism can be implemented using an ARPA with reserve $R_n^* \equiv R_n(r^*)$.*

In the case of ARPA, the optimal reserve is therefore decreasing in the number of suppliers. This comparative static contrasts with the case of a FPPA or SPPA, where the optimal reserve is independent of the number of suppliers, but it is the same as for the efficient and the optimal all-pay selling auction.³

4.2 The nonregular case

Consider now the case in which γ is not strictly monotone in c .⁴ The optimal mechanism then selects agents on the basis of the *ironed* virtual cost function $\bar{\gamma}$, which is continuous, nondecreasing, and differs from γ only on a set of so-called ironing intervals.⁵ Specifically,

³To see the latter, consider for instance the case in which n buyers draw their values independently from F with support $[0, 1]$ and density f , and the auctioneer's opportunity cost of selling the good is $c \in (0, 1)$. The reserve of the efficient all-pay auction is then $F(c)^{n-1}c$, whereas the reserve of the optimal all-pay auction, assuming increasing virtual values, is $F(\tilde{r})^{n-1}\tilde{r}$, where \tilde{r} is such that $\tilde{r} - \frac{1-F(\tilde{r})}{f(\tilde{r})} = c$. Both reserves are decreasing in n .

⁴The virtual cost function is locally increasing at \underline{c} because $\gamma'(c)|_{c=\underline{c}} = 2$. Hence, it is either increasing everywhere (the regular case) or not strictly monotone (the nonregular case).

⁵This is the procurement analogue of the optimal selling mechanism; see Myerson (1981). For a practical implementation of the optimal selling mechanism, see Condorelli (2012).

there exists $K \equiv \{1, \dots, K\}$ and $\{\underline{c}_k, \bar{c}_k\}_{k \in \mathcal{K}}$, satisfying $\underline{c}_1 < \bar{c}_1 < \dots < \underline{c}_K < \bar{c}_K$, such that $\bar{\gamma}(c) = \bar{\gamma}_k \equiv \mathbb{E}_c[\gamma(c) \mid c \in \mathcal{I}_k]$ for $c \in \mathcal{I}_k \equiv [\underline{c}_k, \bar{c}_k]$, and $\bar{\gamma}(c) = \gamma(c)$ for $c \notin \cup_{k \in \mathcal{K}} \mathcal{I}_k$. Note that the ironing intervals, and hence $\bar{\gamma}(\cdot)$, do not depend on the number of bidders. We illustrate the ironing of a nonmonotone virtual cost function in Figure 1.

The optimal reserve price may not be unique if $v = \bar{\gamma}_k$ for some $k \in \mathcal{K}$, as the designer is then indifferent about including or not suppliers with $c \in \mathcal{I}_k$.⁶ Without loss of generality, we break ties in favor of social surplus and include these suppliers; we thus choose

$$r^* = \max\{c \in [\underline{c}, \bar{c}] \mid \bar{\gamma}(c) \leq v\},$$

illustrated in Figure 2(a). Let $\mathcal{K}(v) \equiv \{k \in \mathcal{K} \mid \bar{c}_k \leq r^*\}$ and $\mathcal{I}(v) \equiv \cup_{k \in \mathcal{K}(v)} \mathcal{I}_k$ refer to the included ironing intervals.

For $k \in \mathcal{K}(v)$ and $c_i \in \mathcal{I}_k$, the probability of selecting supplier i is positive but must not vary with c_i . In addition, if multiple agents have the lowest ironed virtual cost, then, without loss of generality, we uniformly randomly select one of these suppliers.⁷ We show in Supplemental Appendix B.1 that the probability of selecting supplier i is then equal to

$$\bar{P}_n(k) \equiv \mathbb{E}_c [P_n(c) \mid c \in \mathcal{I}_k].$$

The selection probability for a supplier with cost c is therefore:

$$P_n^*(c) \equiv \begin{cases} P_n(c) & \text{if } c \in [\underline{c}, r^*] \setminus \mathcal{I}(v), \\ \bar{P}_n(k) & \text{if } c \in \mathcal{I}_k \text{ for } k \in \mathcal{K}(v), \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

as illustrated in Supplemental Appendix Figure C.4.

By the payoff equivalence theorem, the expected payoff of a supplier with cost c , denoted $U_n^*(c)$, is, using $U_n^*(\bar{c}) = 0$,

$$U_n^*(c) = \int_c^{r^*} P_n^*(x) dx.$$

We amend the ARPA to implement the optimal procurement in the nonregular case, referring to the amended mechanism as the ARPA*. By construction, the ARPA* must

⁶This is the procurement analogue to the observation made by Bulow and Roberts (1989) for the problem faced by a monopoly producing with constant marginal costs.

⁷Because they have the same ironed virtual cost, any other cost-independent tie-breaking rule, such as prioritizing agents on the basis of their names or identities, would work as well. Breaking ties uniformly randomly paves the way towards existence of a symmetric equilibrium.

induce a supplier with cost $c < r^*$ to bid $\beta_n^*(c) \equiv P_n^*(c) + U_n^*(c)$. Like the selection probability, the bid function is constant over any ironing interval, and jumps down at both ends of the interval: for $k \in \mathcal{K}(v)$,

$$\beta_n^-(k) \equiv \lim_{c \rightarrow \underline{c}_k^-} \beta_n^*(c) > \bar{\beta}_n(k) \equiv \beta_n^*(\underline{c}_k) = \beta_n^*(\bar{c}_k) > \beta_n^+(k) \equiv \lim_{c \rightarrow \bar{c}_k^+} \beta_n^*(c),$$

as illustrated in Figure 2(b).

The rules of the ARPA* are the same as in the ARPA with reserve $R_n^* \equiv \beta_n^*(\underline{c})$, with the caveat that, for selection purposes, any (off-path) bid, i.e. any b between $\bar{\beta}_n(k)$ and $\beta_n^+(k)$, for $k \in \mathcal{K}(v)$, is treated as $b' = \bar{\beta}_n(k)$. With this adjustment, we have:

Proposition 2. *In the ARPA* with reserve R_n^* , there exists a symmetric Bayesian Nash equilibrium, in which every supplier with cost $c \leq r^*$ bids $\bar{\beta}_n^*(c)$ and every supplier with cost $c > r^*$ refrains from bidding. This equilibrium implements the optimal procurement mechanism.*

Proof. The incentive compatibility of the optimal mechanism ensures that, for any $c \in \mathcal{C}$, bidding $\beta_n^*(c)$ (weakly) dominates any other bid in $\beta_n^*(\mathcal{C})$. It remains to show that doing so also dominates any unexpected bid, that is, any bid in $[-\infty, \beta_n^*(r))$ or in $[\beta_n^+(k), \beta_n^-(k)] \setminus \{\bar{\beta}_n(k)\}$, for $k \in \mathcal{K}(v)$. To see this, it suffices to note that: (i) any bid $b \in [-\infty, \beta_n^*(r))$ (resp., $b \in [\beta_n^+(k), \bar{\beta}_n(k))$) is dominated by $b' = \beta_n^*(r)$ (resp., $b' = \bar{\beta}_n(k)$), which yields the same selection probability, namely, 1 (resp., $\bar{P}_n(k)$, thanks to the adjustment brought to the ARPA* selection rule) but a larger payment; (ii) likewise, any bid $b \in (\bar{\beta}_n(k), \beta_n^-(k))$ is dominated by $b' = \beta_n^-(k)$, which again yields the same selection probability (namely, $P_n(\underline{c}_n(k))$) but a larger payment; and (iii) bidding $b = \beta_n^-(k)$ yields approximately the same payoff as bidding slightly above $\beta_n^-(k)$. ■

5 Resilience to liquidity shocks

The question of why a procurer should use an ARPA or ARPA* rather than a standard format such as the FPPA or SPPA (or the appropriately adjusted formats FPPA* and SPPA* defined below for the nonregular case) is a natural one. To address it, we extend the model and assume that suppliers are liquidity constrained and exposed to liquidity shocks that are not anticipated. We provide conditions under which the ARPA and ARPA* afford resilience to liquidity shocks in the sense that they permit a competitive supply structure to survive where none of the other formats does. In this thought experiment, we keep fixed the direct mechanism that a given auction format implements—efficient or optimal procurement—and

merely ask which format is more resilient in the sense just defined. In particular, this is *not* an exercise in designing optimal mechanisms in the face of liquidity constraints and liquidity shocks.

As mentioned in footnote 1, production costs can alternatively be interpreted as (i) opportunity costs of selling an existing asset or as (ii) the financial cost of producing an asset or service. To fix ideas, for the remainder of this section, we stick to interpretation (ii), that is, the liquidity of a supplier decreases when it is selected to produce. Specifically, we stipulate that suppliers are initially cash constrained and can moreover be hit by an industry-wide liquidity shock $\varepsilon > 0$, which occurs after procurement and production. To ensure that bidding is not affected by the possibility of a liquidity shock, we assume that the probability of the shock occurring is 0, to be interpreted as the limit of a very small shock probability. The eventual cash holding, once the procurement and liquidity shocks have been realized is $p_i - c_i - \tilde{\varepsilon}$ for an agent with liquidity shock $\tilde{\varepsilon} \in \{0, \varepsilon\}$ who is paid p_i and has to produce the good at cost c_i , and $p_i - \tilde{\varepsilon}$ for the same agent if it does not have to produce. Finally, suppliers who end up with negative liquidity go bankrupt.

We will say that an auction format *preserves competition* if the event that two or more suppliers survive the liquidity shock occurs with positive probability.

5.1 Survival in efficient procurement auctions

We begin by considering the case of efficient procurement auctions, in which in equilibrium the supplier with the lowest cost produces with probability 1, provided this cost is below v . The FPPA and SPPA with reserve $r = \min\{v, \bar{c}\}$ are efficient, as is the ARPA with reserve $R_n(r)$, and this holds independently of whether γ is increasing.

Given r , if $\varepsilon \leq R_n(r) - \underline{c}$, there exists a unique $\tilde{c}_\varepsilon(r)$, defined by $B_n(\tilde{c}_\varepsilon(r); r) - \tilde{c}_\varepsilon = \varepsilon$, such that the supplier selected to produce in the ARPA survives after a liquidity shock if and only if its cost is less than or equal to \tilde{c}_ε . If, instead, $\varepsilon > R_n(r) - \underline{c}$, then the producer never survives a liquidity shock in the ARPA (nor in the FPPA or SPPA, as we will see). Similarly, if $\varepsilon \leq R_n(r)$, then a supplier that does *not* produce survives a liquidity shock in the ARPA if and only if its cost is less than or equal to $\tilde{c}_\varepsilon(r) \equiv B_n^{-1}(\varepsilon; r)$, with the inverse being taken with respect to the first argument of $B_n(c; r)$.

In contrast, in the FPPA and SPPA, only the winning supplier has a chance to survive a liquidity shock, and it moreover does so only if its cost is low enough. In particular, for the FPPA, if $\varepsilon \in (R_n(r) - \underline{c}, R_n(r))$, then a liquidity shock kills all suppliers in the FPPA,⁸ but in the ARPA it kills only the supplier that produces and the suppliers with costs above

⁸Recall that $R_n(r) = b_n^{FPPA}(\underline{c}; r)$; hence, if $\varepsilon > R_n(r) - \underline{c}$, a liquidity shock kills the winning supplier even if it has the lowest possible cost, \underline{c} .

$B_n^{-1}(\varepsilon; r)$. For smaller liquidity shocks, the supplier that produces in the FPPA survives only if its cost is less than $\tilde{c}'_\varepsilon(r)$, defined by $b_n^{FPPA}(\tilde{c}'_\varepsilon(r); r) - \tilde{c}'_\varepsilon = \varepsilon$. Because $B_n(c; r) < b_n^{FPPA}(c; r)$ holds for any $c > \underline{c}$ (see (3)), we have for any $\varepsilon \in (0, R_n(r) - \underline{c})$, $\tilde{c}'_\varepsilon(r) > \tilde{c}(r)$. This implies that for liquidity shocks of this form, the supplier that is selected to produce is more likely to survive in the FPPA than in the ARPA. However, because suppliers that do not produce are paid nothing in the FPPA or, for that matter, the SPPA, only the ARPA has the potential to preserve competition.

Summarizing, we have:

Proposition 3. *Assume $r = \min\{v, \bar{c}\}$ and let the liquidity shock ε tend to 0. In the FPPA and SPPA with reserve r , only one supplier—the winner—survives a liquidity shock. In the ARPA with reserve $R_n(r)$, the supplier that produces survives with probability $1 - [1 - F(\tilde{c}_n(r))]^n$, and every other supplier with a cost $c < r$ survives.*

Obviously, if $v < \bar{c}$, then suppliers with $c > v$ are excluded, regardless of the auction format. If, instead, $v \geq \bar{c}$, then a liquidity shock kills all non-selected suppliers in a FPPA or SPPA, and none of them in the ARPA.

Because $\tilde{c}_n(r)$ is decreasing in n , it is not a priori clear whether the probability that the producing supplier survives the liquidity shock when $\varepsilon \rightarrow 0$ increases or decreases in n because, for a fixed c , $1 - [1 - F(c)]^n$ increases with n . As shown in Supplemental Appendix Figure C.5, if costs are uniformly distributed, then the probability that the producing supplier survives such a liquidity shock decreases in n . However, numerical results suggest that it is bounded from below by $1/2$.

5.2 Survival in optimal procurement auctions

We now extend the resilience analysis to optimal procurement auctions. In the regular case, Proposition 3 extends directly to optimal procurement auctions with r replaced by $r^* = \gamma^{-1}(\min\{v, \bar{c}\})$. For the nonregular case, we first need to identify variants of the FPPA and SPPA, referred to as FPPA* and SPPA*, that implement the optimal mechanism characterized in Section 4.2. In doing so, we hold constant the central features of these formats, namely that the winning supplier is paid its bid in a FPPA and that an SPPA endow suppliers with dominant strategies to bid truthfully.

For the FPPA*, we seek to generate a symmetric Bayesian Nash equilibrium in which a supplier with cost $c < r^*$ bids $\beta_n^{FPPA^*}(c)$, equal to $\beta_n^*(c)/P_n^*(c)$; contrary to $\beta_n^*(c)$, this bid function jumps *up* at both ends of any relevant ironing interval, as illustrated in Supplemental Appendix Figure C.6. The rule of the FPPA* is the same as that of a FPPA with reserve

r^* , except that for selection purposes, any (off-path) bid $b \in [\lim_{c \rightarrow \underline{c}_k^-} \beta_n^{\text{FPPA}^*}(c), \beta_n^{\text{FPPA}^*}(\underline{c}_k)]$, for $k \in \mathcal{K}(v)$, is treated as $b' = \beta_n^{\text{FPPA}^*}(\underline{c}_k)$.⁹

Similarly, the rule of the SPPA* is the same as that of the SPPA, unless the second-lowest bid is in \mathcal{I}_k , for some $k \in \mathcal{K}(v)$. In that case, if the lowest cost is also in \mathcal{I}_k , then one of the bidders in \mathcal{I}_k is selected uniformly randomly and paid \bar{c}_k . If, instead, the lowest bid is less than \underline{c}_k , then the lowest bidder is selected and paid

$$p_m(k) \equiv \frac{1}{m+1} \bar{c}_k + \frac{m}{m+1} \underline{c}_k,$$

where m is the number of bids in \mathcal{I}_k . As in the usual second-price auction, truthful bidding is a dominant strategy in the SPPA*.¹⁰ It then follows from the description of the SPPA* that it implements the optimal selection probabilities. Furthermore, incentive compatibility, together with the fact that a supplier with $c > r^*$ obtains zero payoff, ensures that a supplier with cost c obtains $U_n^*(c)$.

With these adjusted formats in hand, we have:

Proposition 4. *Assume $r = r^*$ and let the liquidity shock ε tend to 0. Among the auction formats SPPA, FPPA, and ARPA, only the ARPA preserves competition, and among the auction formats SPPA*, FPPA*, and ARPA*, only the ARPA* preserves competition.*

Proof. In the SPPA and FPPA, as well as in the SPPA* and FPPA*, at most one bidder receives a payment, and so only that bidder has a positive probability of surviving a liquidity shock. In the ARPA, two bidders survive, for example, whenever there are two bidders with costs below $\tilde{c}_n(r)$. Similarly, in the ARPA*, two bidders survive, for example, whenever there are two bidders with costs c such that $c < \beta_n^*(c)$, which occurs with positive probability. ■

As this proposition indicates, when suppliers have low liquidity, if the buyer wants to preserve competition, then the ARPA (or ARPA*) is a forced move because neither first-price nor second-price formats can “save” competition.

Although the propositions in this section apply for $\varepsilon \rightarrow 0$, for larger values of ε , it is still the case that in the FPPA and SPPA (and FPPA* and SPPA*), at most one supplier survives. In contrast, in the ARPA, all suppliers survive a shock of size ε whenever all suppliers’ costs c satisfy $\varepsilon \leq B_n(c; r) \leq R_n(r)$ and the lowest cost is less than \tilde{c}_ε (and analogously for the ARPA*). In that sense, $\varepsilon \rightarrow 0$ is a best-case scenario for the first-price and second-price formats.

⁹The argument establishing that $\beta_n^{\text{FPPA}^*}(c)$ constitutes a symmetric response to the FPPA* is similar to that for the ARPA*.

¹⁰See Supplemental Appendix B.2 for details.

6 Discussion

In this section, we discuss alternative auction formats, address collusion concerns, and provide concluding remarks.

6.1 Multiple-receive procurement auctions and maximal resilience

With the FPPA, the symmetric equilibrium bids are increasing in and exceed costs, implying that no supplier ever regrets ex post being selected to produce; in contrast, with the ARPA the equilibrium bids are decreasing in and eventually fall below costs, implying that the individual rationality constraint can be violated ex post for the selected supplier.¹¹ On the other hand, the ARPA has desirable resilience properties. To balance ex post participation and resilience concerns, one avenue is a *multiple-receive procurement auction*, where in equilibrium the $m \in \{1, \dots, n\}$ bidders with the lowest costs are paid their bids and the lowest-cost bidder is selected to produce—the FPPA corresponding to $m = 1$, and the ARPA to $m = n$.

The above description of the auction rule is unorthodox, in that it relies on agents' types rather than bids. However, without knowing whether the equilibrium bidding strategy is increasing (as for the FPPA) or decreasing (as for the ARPA), this unorthodoxy cannot be avoided. If the bidding strategy is increasing (resp., decreasing), then the m lowest (resp., highest) bids will be paid and the lowest (resp. highest) bidder asked to produce. Interestingly, and maybe unfortunately, we show in Supplemental Appendix B.3 that a monotone equilibrium bidding strategy need not exist for $m \in \{2, \dots, n - 1\}$.

Alternatively, any given nonincreasing allocation (or selection) rule $\tilde{P}_n(c)$, together with associated informational rents, can be implemented in a way that compensates the selected supplier for the cost of production. Specifically, this implementation asks every supplier to report its cost; any bidder with a reported cost c is paid $\int_c^{\bar{c}} \tilde{P}_n(x) dx$ and, if selected to produce, is additionally paid c . To see that this induces truthful reporting in a Bayesian Nash equilibrium, observe that the expected payoff of a bidder with cost c that reports \hat{c} is

$$\int_{\hat{c}}^{\bar{c}} \tilde{P}_n(x) dx + \tilde{P}_n(\hat{c})(\hat{c} - c),$$

whose derivative with respect to \hat{c} is $\tilde{P}'_n(\hat{c})(\hat{c} - c)$. Thus, the first-order and second-order conditions are satisfied at $\hat{c} = c$. This implementation clearly *maximizes resilience*, as it always satisfies the individual rationality constraints ex post and pays every bidder the

¹¹This partly mirrors a property of the all-pay auction, where ex post individual rationality is always violated for all but the winning bidder.

maximum amount possible, subject to this constraint and subject to implementing this selection rule.

6.2 Susceptibility to collusion

The susceptibility of various auction formats to collusion has been a frequent topic in the literature. The general view is that second-price formats are more susceptible to collusion than first-price formats (see, e.g., Kovacic et al., 2006). In the case of the ARPA, a collusive scheme might attempt to have the lowest-cost supplier submit the highest admissible bid, with all others submitting bids slightly below that, thereby securing the highest payment for all bidders. However, a number of factors mitigate the vulnerability of the ARPA to collusion, including the fact that the reserve is there substantially lower, and a strong incentive for the supplier designated to produce to deviate to a lower bid in order to avoid incurring the production cost.

To make things precise, consider the setup of Iossa et al. (2024), in which, in each period over an infinite horizon: (i) a buyer with value v seeks to purchase in two identical markets, in which two suppliers draw costs independently (across suppliers, markets, and time); (ii) all agents discount the future according to a common discount factor $\delta \in [0, 1)$; and (iii) bids submitted in one period are observed prior to the next period. Collusion takes the form of a market allocation in which, for each market $i = 1, 2$, supplier i is the designated supplier and the other supplier bids slightly less aggressively—detected deviations being punished with reversion to competitive bidding in both markets forever.

If the auction format is a FPPA with reserve r , then the designated supplier bids slightly below the reserve if its cost lies below the reserve and at cost otherwise, whereas the non-designated supplier bids the reserve if its cost lies below it and at cost otherwise. If the auction format is a SPPA with reserve r , then the designated supplier bids its cost and the nondesignated supplier bids the reserve if its cost lies below it, and at cost otherwise. As shown in Iossa et al. (2024), consistent with intuition, collusion is profitable in both the FPPA and SPPA, and the SPPA is more susceptible to collusion in the sense of having a lower critical discount factor such that collusion is sustainable.

For the ARPA, at least two collusive schemes are potentially relevant. In the first one, given the reserve $R_2(r)$, the designated bidder bids $R_2(r)$ and the nondesignated bidder bids slightly below $R_2(r)$, regardless of their respective costs; both bidders are then paid $R_2(r)$, and the designated bidder incurs the production cost. In the second scheme, the designated bidder bids $R_2(r)$ when its cost is less than $R_2(r)$ and bids its cost otherwise, whereas the

nondesignated bidder bids slightly below $R_2(r)$ when its cost is less than

$$\hat{c}(r) \equiv \frac{R_2(r)}{1 - F(R_2(r))},$$

and bids its cost otherwise.¹² In the second scheme, the designated bidder with cost c receives $\max\{0, R_2(r) - c\}$, the nondesignated bidder with cost $c < \hat{c}(r)$ has an expected payoff of $R_2(r) - [1 - F(R_2(r))]c$, and the nondesignated bidder with a higher cost receives zero.

There are a number of ways in which the ARPA is less vulnerable to collusion than the FPPA and SPPA (see Supplemental Appendix D for details). First, if the hazard rate is monotone, collusion is always profitable in the FPPA and SPPA (Iossa et al., 2024, Lemma 1); in contrast, it is not necessarily profitable in the ARPA. For example, if $F(c) = c^s$ on $[0, 1]$ with $s \in (0, 1/2)$, and $r = 1$, then ARPA collusion is not profitable under either variant. This is because the payment $R_2(r)$ that the colluding bidders receive approaches zero as s approaches zero, and so collusive payments are outweighed by the inefficiency associated with collusion.

Second, collusion in the ARPA can, in some settings, be harder to sustain than in the other formats. Indeed, to be sustainable, the short-term gain from deviation should be sufficiently lower than the long-term loss from reverting to competition. With the ARPA, the short-term gain from deviation can be quite high, and the long-term loss quite low. For instance, in the first scheme, a designated bidder with cost \bar{c} has a short-term gain of \bar{c} because it avoids production by deviating;¹³ and for uniformly distributed costs on $[0, 1]$, the long-term loss for the second scheme is always below that of the FPPA and SPPA.¹⁴ Thus, the ARPA can be less susceptible to collusion than the other formats—for specific settings, however, it can be more susceptible (see Supplemental Appendix D.5).

Finally, while market allocations in the FPPA and SPPA ensure that trade occurs whenever it would occur under competition, i.e., if and only if at least one supplier has a cost below the reserve r , this is not the case in the ARPA for $r < \bar{c}$. In variant 1, trade always occurs, which means that there is trade even when both suppliers have costs in $(r, \bar{c}]$ and so there would be no trade under competition. In variant 2, there is trade when both costs are in $(r, \bar{c}]$ and the nondesignated supplier's cost is less than $\hat{c}(r)$, and there is no trade when the designated supplier's cost is in $(R_2(r), r)$ and the nondesignated supplier's cost is

¹²In a third variant, the designated bidder bids $R_2(r)$ when its cost is less than $R_2(r)$ and bids its cost otherwise, and the nondesignated bidder bids slightly below $R_2(r)$ regardless of its cost. Our simulations suggest that this variant is dominated in terms of profitability and sustainability by at least one of the other variants.

¹³In the second scheme, the most profitable deviation is by a designated supplier with cost $R(c)$ that deviates to a bid below that of the nondesignated supplier, for a gain of $F(\hat{c}(r))R(r)$.

¹⁴See supplemental appendix D.

in $(\hat{c}(r), \bar{c})$, even though in that case there is a supplier with cost below r .

Thus, while the ARPA, like other auction formats, is potentially vulnerable to collusion, it does not appear to be significantly more vulnerable than other commonly used auction formats.

6.3 Concluding remarks

Motivated by the recent upsurge of interest in supply chain resilience, this paper proposes a new family of auction formats, called all-receive procurement auctions, in which all suppliers with bids below the reserve are paid their bids and the highest bidder is selected to produce. These formats are payoff equivalent to standard formats such as the first-price and second-price procurement auction and permit implementing both efficient and optimal procurement auctions. Because in equilibrium more suppliers are paid than in the standard formats, the liquidity holdings of the suppliers that are not selected to produce are larger than in the standard formats, making them more resilient to certain forms of nonanticipated liquidity shocks. Because they are payoff equivalent, all-receive procurement auctions may afford resilience without any cost to the procurer.

Rather than paying multiple suppliers while having a single supplier produce, a procurer might consider spreading production across multiple suppliers. This would allow the buyer to obtain parts of the input even if a subset of the suppliers fail to deliver, for example, because of a breakdown in their production facility or problems associated with shipment. Notwithstanding the intuitive appeal, multisourcing of this form is inefficient in independent private values models and not optimal in the regular case because, with probability one, a single supplier has the lowest (virtual) cost. Consequently, under these conditions, multisourcing comes at a cost to the procurer. However, without regularity, multisourcing becomes optimal without any cost to the procurer because rather than breaking ties randomly among suppliers with the lowest ironed virtual cost to select a single supplier to produce, ties can equivalently be broken to have each of the suppliers involved in the tie produce the same fraction of the total. While away from efficiency, our focus has been on profit-maximizing auctions, in general (i.e., without sufficiently strong regularity) ironing and tie-breaking remain part of the procurement mechanisms that maximize a convex combination of social surplus and the buyer's profit, provided that the weight on the latter is positive. Thus, multisourcing for the purpose of increasing supply chain resilience can be optimal quite generally.

References

- ARVE, M. AND D. MARTIMORT (2024): “Auctioning Long-Term Projects under Financial Constraints,” *Review of Economic Studies*.
- BULOW, J. AND J. ROBERTS (1989): “The Simple Economics of Optimal Auctions,” *Journal of Political Economics*, 97, 1060–90.
- CONDORELLI, D. (2012): “What Money Can’t Buy: Efficient Mechanism Design with Costly Signals,” *Games and Economic Behavior*, 75, 613–624.
- GANUZA, J.-J. (2007): “Competition and Cost Overruns in Procurement,” *Journal of Industrial Economics*, 633–660.
- IOSSA, E., S. LOERTSCHER, L. M. MARX, AND P. REY (2024): “Coordination in the Fight against Collusion,” *American Economic Journal: Microeconomics*, 16, 224–261.
- KOVACIC, W. E., R. C. MARSHALL, L. M. MARX, AND M. E. RAIFF (2006): “Bidding Rings and the Design of Anti-Collusive Measures for Auctions and Procurements,” in *Handbook of Procurement*, ed. by N. Dimitri, G. Piga, and G. Spagnolo, Cambridge and New York: Cambridge University Press, 381–411.
- KRISHNA, V. (2010): *Auction Theory*, Elsevier Science, Academic Press.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- RILEY, J. G. AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.
- VICKREY, W. (1961): “Counterspeculation, Auction, and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- (1962): “Auctions and Bidding Games,” in *Recent Advances in Game Theory*, Princeton, NJ: Princeton University Press, vol. 29, Princeton Conference Series, 15–27.