London Business School

LBS Research Online

D R Beil, Q Chen, I Duenyas and B D See
When to deploy test auctions in sourcing
Article

This version is available in the LBS Research Online repository: `http://lbsresearch.london.edu/1021/`

When to Deploy Test Auctions in Sourcing

Damian R. Beil, Qi (George) Chen, Izak Duenyas

Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109

dbeil, georgeqc, duenyas@umich.edu

Brendan D. See

A.T. Kearney, New York, New York 10036

see.brendan@gmail.com

We investigate when a buyer seeking to procure multiple units of an input may find it advantageous to run a "test auction" in which she has incumbent suppliers bid on a portion of the desired units. The test auction reveals incumbent supplier cost information that helps the buyer determine how many entrants (if any) to recruit at a cost prior to awarding the remaining units. The optimal number of entrant suppliers to recruit follows a threshold policy that is monotonic in the test auction's clearing price unless the underlying supplier cost distribution is not regular. When setting her reserve price in the test auction, the buyer uses supplier recruitment as her "outside option": If the reserve price is not met in the test auction, the buyer recruits new suppliers and runs a second auction. We compare the attractiveness of the test auction procedure relative to the more conventional procedure in which the buyer auctions off her entire demand in one auction. Since the buyer can choose ex ante which procedure to use, we propose using whichever has lower ex ante total (purchase plus recruitment) cost. Finally, using an optimal mechanism analysis, we find a lower bound on the buyer's cost, and use that cost as a benchmark to show that our proposed sourcing strategy performs well given its ease of implementation.

Key words: new supplier recruitment; procurement; sequential auctions; test auctions; reserve price; mechanism design; optimal mechanism.

1. Introduction

Many firms award supply contracts via competitive bidding, whereby suppliers fiercely compete against each other in a "reverse auction". In this paper we address the following question: How many bidders should a buyer recruit for a competitive bidding event? Surprisingly, most research on auctions takes the existence of bidders as their point of departure. More bidders increases competition and may result in better prices, but it is not necessarily simple to identify new potential suppliers (bidders).

Recently the authors worked with a Fortune 500 manufacturer (the buyer) who conducted an extensive search for potential suppliers for a new part. The part is used in highly engineered applications and is produced by high tolerance milling of bar stock on multi-spindle lathes. The buyer gathered supplier names by scouring various sources, including

2

Beil, Chen, Duenyas, and See: *When to Deploy Test Auctions in Sourcing*
Article submitted to *Manufacturing & Service Operations Management*; manuscript no. MSOM-14-313

library and industry databases, web searches, and other business units within the firm. They also contacted multi-spindle lathe manufacturers to get customer lists. After searching the globe for suppliers, and discarding those who were too small or lacked ultrasonic washing capabilities (the parts themselves needed to be extremely clean to be usable in the buyer's assembly processes), the buyer was left with a handful of new prospective suppliers.

The approaches a buyer can use to recruit prospective suppliers are not limited to identifying existing suppliers, but could involve identifying novel sources: At the buyer firm we interacted with, for parts the buyer firm typically bought from bar stock finishers, suppliers capable of manufacturing parts out of cold-headed blanks instead of bar stock were also considered for part segments where the material and geometry allowed a near net shape blank to be cold headed and then finished through machining. Buyers can also draw on third party service providers — a large cottage industry of firms exists to help buyers identify potential suppliers. Trading on their knowledge and familiarity with suppliers and their capabilities, these firms provide interested buyers with a prospective supplier list, often in return for a fee. Even after locating a supplier, the buyer would still need to explain its specific requirements and communicate extensively with the supplier to ensure that there is no misunderstanding about the specifications as well as the expectations on both sides. This can be very time consuming when dealing with new suppliers. Also, oftentimes the buyer representatives spend some time to be reasonably confident (ahead of a formal qualification process) that the supplier does have the technology that they claim.

Since recruiting new suppliers entails costs for the buyer (in tracking down prospective suppliers, identifying novel sources, or paying a third party), the decision of how many bidders to recruit for the auction is not trivial: The number of suppliers to include in an auction should balance the benefits the buyer anticipates from intensified bidding competition with the costs of recruiting more suppliers.

In fact, the buyer firm we worked with had a few existing suppliers ("incumbents") already in the supply base who they knew were capable of making the new part. However, the fact that they may have made similar parts in the past does not mean that the buyer will know how much the supplier will charge for a new part — changes in specifications, supplier utilization, and economic factors can significantly affect costs. Moreover, the buyer decided to recruit new suppliers ("entrants") in order to heat up competition for producing

the part and drive down pricing. The buyer eventually sourced the parts by running a reverse auction in which incumbent and entrant suppliers competed for the business.

After observing the above, and contemplating the time and resources that a buyer invests in recruiting suppliers, we were left wondering if there could be a way to manage these supplier recruitment costs. In fact, given how simple it is to run a reverse auction, a buyer firm could potentially auction off a portion of her demand in order to gain a better understanding of the pricing in the incumbent market (whether it is very competitive or not), and based on this determine how many new suppliers to recruit. Moreover, when the buyer uses a reserve price, supplier recruitment can serve as the buyer's "outside option": If the reserve price is not met in the first auction, the buyer recruits new suppliers and runs a second auction. Although this is a very natural question that would arise in a variety of industries and settings (after all, exerting time and effort to recruit new suppliers is something that many buyers do prior to competitive bidding events), we could find little prior research addressing it. Filling this gap is the goal of this paper.

Other papers have addressed the process of dealing with new suppliers once they have already been added to buyer's supply base. In the operations management literature, these papers typically deal with supplier qualification. The supplier qualification process refers to the act of verifying a potential suppliers' production abilities, history with previous customers, financial stability, and other attributes. The main distinction between new supplier recruitment and supplier qualification stems from a potential supplier's ability to place a competitive bid: not-yet-qualified suppliers could hypothetically place a bid before undergoing the qualification process; by contrast, in our paper's case suppliers who have not been recruited yet are unable to bid because they are not aware of the auction and the buyer has not discovered them. The supplier qualification literature (e.g., Wan and Beil (2009) and Wan et al. (2012)) considers whether it is optimal to delay qualification screening until after the auction. In our paper, the buyer cannot delay recruiting suppliers until after the auction, so the problems faced are fundamentally different. We focus on the idea of test auctions to inform the buyer's decision on how many suppliers to recruit.

Specifically, in our paper we model the following trade-off. To ensure competition and realize the lowest possible cost, a buyer ideally wants as many suppliers to compete for the contract as possible. However, because recruiting additional suppliers is expensive, the buyer faces a difficult decision. She can either absorb new supplier recruitment costs and

recruit more suppliers to increase price competition, or instead she can choose not to add any additional suppliers to her supply base and hope that the existing suppliers can produce the input at a low cost. Usually the procurement process must be completed in a timely manner because once a buyer has decided to put a particular contract out for bid (e.g., because it won new business from an upstream customer), there is a limited time window within which the supplier(s) for the contract must be selected. To accommodate the time-consuming nature of the bid preparation process (e.g., understanding the product specifications, getting trained on the buyer auction software platform, etc.), for each contract, a buyer can usually perform one recruitment round whereby she recruits new suppliers who then prepare their bids in parallel.

Recognizing these complicating factors, we examine the following approach when answering the question of how many bidders to recruit for an auction. The buyer may hold an initial auction, which we call a "test auction," for a portion of the units she needs in order to get pricing information from the existing supply base. Then, the buyer can use her updated knowledge about the existing suppliers to help her decide how many additional suppliers to recruit for a final auction in which the remainder of the units will be sourced. We note that it is fairly common for firms to pick a supplier, give them some business and consider them for additional business if their price performance meets firm targets. Several firms we have worked with already do this. Thus, our proposed test auction procedure would not be out of place in such environments. In our setting, the unsettled question resolved by the auction is cost/price performance.

In analyzing the buyer's problem, we answer the following research questions:

1. When is it a good strategy for the buyer to run a test auction versus recruit additional suppliers and run one auction for all units?

2. How many suppliers should the buyer recruit under each scenario, and in the test auction scenario how does this decision depend on the outcome of the test auction, and how can the buyer best implement a reserve price? Furthermore, how well do these strategies perform compared to a theoretically optimal (but perhaps hard to implement) mechanism (along the lines of Myerson (1981)) that serves as a lower bound of the optimal cost?

3. How do cost distributions, the size of the minimal quantity that has to be auctioned, the number of incumbent suppliers, and other business factors affect these decisions?

There has been a fair amount of literature on sourcing policies that address both auctions and other types of contracting mechanisms; Elmaghraby (2000) provides a good review. The main contribution of our paper is the novel "test auction" procedure, where the buyer may choose to use sequential auctions separated by a new supplier recruitment round. The test auction serves as a partial cost-discovery mechanism which the buyer uses to inform her decision regarding how many entrant suppliers should be added to the supply pool to compete in a subsequent auction. Peleg et al. (2002) study a setting where the buyer chooses how many suppliers to recruit for an auction, but they do not study the possibility that the buyer runs a test auction to update her information about existing suppliers, which is the focus of our paper. de Boer et al. (2000) find the "Economic Tender Quantity" when there are costs associated with sending RFQs, evaluating suppliers' tenders, and communicating the results of the competition. In our model, we study the same type of cost, but incorporate incumbent suppliers and allow multiple auctions.

Many buyers split their contract among different suppliers. Papers on multi-sourcing generally focus on mitigating supply disruption risks, e.g., Tomlin (2006), Federgruen and Yang (2009), Chaturvedi and Martínez-de-Albéniz (2011), Yang et al. (2012). Our buyer also can multi-source, but for an entirely different reason. In our paper, multi-sourcing can be a consequence of the buyer discovering information about the incumbents' costs.

To our knowledge, ours is the first paper to consider the use of test auctions to help manage new supplier recruitment and sourcing costs. Our paper should assist firms to better manage their total (contracting+new supplier recruitment) procurement costs.

2. Model

We model a risk-neutral, cost-minimizing buyer who seeks to procure Q units of a certain input. For example, Q could be in the tens or hundreds of thousands, corresponding to several years of supply of a certain component. We initially assume that the contracts for these units will be awarded using open-bid descending-price procurement auctions, which proceed as follows: The auction price falls continuously until all but one bidder drops out; the last remaining bidder wins the auction and is paid the auction ending price. (Auctions with a continuously falling price are also known as "reverse clock auctions"; see Ausubel and Cramton (2006) for discussions about clock auctions in practice.) Of course, the buyer

could use other mechanisms to award the contract, but we focus on open-bid descending-price auctions since they are ubiquitous in practice and are easy to explain and implement. (In §4.1, we address the other common auction format, the first-price sealed-bid auction.)

We assume that whenever the buyer runs an auction she has to auction off at least y units. Auctions for very small quantities may not generate much interest from suppliers and that is why we model a minimum order quantity. (For example, an auction to buy one standard bolt is not going to entice any suppliers to bid.) Different industries may have different minimum order quantities, and our model allows for any $0 \leq 2y \leq Q$ (note that, in order to use the test auction approach that involves *two* auctions, the buyer's total quantity needs to be no smaller than twice the minimum order quantity y, i.e., $Q \geq 2y$).

We assume that there are two groups of suppliers: incumbent suppliers and as-of-yet-unknown potential suppliers. The *incumbent suppliers* can be viewed as suppliers whom the buyer has recruited in the past and thus further recruitment is unnecessary. This may be because the supplier has produced a previous generation of the product or has supplied a similar product. Companies we have worked with will typically have at most a few incumbents for any given part (e.g., there were only two incumbents for many parts at the buyer firm mentioned in the Introduction). We let $n \geq 2$ denote the number of incumbent suppliers. There are also suppliers who are as-of-yet unknown; we call these *potential* or *entrant suppliers*. To recruit these suppliers, the buyer incurs a cost. For example, the cost may stem from scouring the globe for existing suppliers, developing novel sources, or paying a third party a finder's fee. The expected cost to recruit m entrant suppliers is $k(m)$. We assume that $k(\cdot)$ is an increasing convex (possibly weakly convex) function. This captures the fact that each successive entrant is typically more costly for the buyer to identify because the buyer will exhaust the most cost-effective avenues of recruitment first. Further, we assume that $k(1) > 0$ — that is, the buyer cannot recruit an entrant for free (alternatively, cases where $k(1) = 0$ can be handled by our model by considering the entrants who cost zero for the buyer to recruit as incumbent suppliers).

As is common in the literature (e.g., Chen 2007) we assume each supplier i is risk-neutral and has a linear production cost function $x_i \cdot q$, where q is the number of units to be produced. The x_i's are independent and identically distributed random variables with cumulative distribution F, probability distribution f, and support $[a, b]$. Let $X_{(i:j)}$ denote the i^{th}-lowest value among j independent draws from distribution F. Further, we assume

that the cost distribution F is *regular* (i.e., $x + \frac{F(x)}{f(x)}$ is increasing in x). Note that the regularity condition is satisfied by many distributions such as the uniform, normal, and Pareto distributions and is a common assumption in the auction literature. It is important to note that the new supplier recruitment process is aimed at identification of suppliers, not cost discovery. This is why the buyer needs to run an auction among the recruited suppliers for cost discovery. Each supplier's variable cost x_i is their private information, but the distribution F is common knowledge. We assume that the buyer must transact with a supplier; this captures cases where the buyer is purchasing a component that is needed in order to assemble her products, but the buyer does not have in-house production capabilities for the component. We discuss relaxing this assumption in §5.

In general once a buyer has decided to put a particular contract out for bid, there is a limited time window within which new contracts must be struck. To accommodate the fact that time is required for entrant suppliers to be identified, understand the buyer's specifications, get trained on the buyer's auction software platform, etc. within the sourcing time window, we are considering a setting where the firm can only recruit suppliers once (although the firm can recruit multiple, m, entrant suppliers in parallel). This eliminates the possibility of the buyer spending many months identifying and waiting while an entrant prepares to bid in an initial auction, then spending many more months doing the same for a second entrant followed by a second auction, etc. Such a "sequential screening" setting has been studied in McAfee and McMillan (1988), but the common practical setting that we are studying allows for recruitment only once during the procurement cycle.

2.1. Main Trade-Off and the Timelines of the Test and No-Test Strategies

The buyer must decide if she wants to (i) recruit additional entrant suppliers up-front or (ii) auction off a portion of the units among the incumbent suppliers to inform her decision regarding how many entrants to recruit, and then procure the remaining units through a subsequent auction. We will call (i) the **no-test strategy** and (ii) the **test strategy**. On the one hand, the no-test strategy increases competition for all Q units. However, this may not be optimal for the buyer. If the buyer auctions off a subset of the units before recruiting any entrants, she may discover important cost information. If the incumbent suppliers' cost realizations are small, then spending additional money to recruit entrant suppliers will not be as attractive as it was ex ante, and thus recruitment costs could be saved. Likewise, if the buyer discovers that the incumbent suppliers have relatively high

costs, she may choose to recruit *more* entrant suppliers than she would have prior to discovering this information. Thus, deploying an initial auction for a subset of the units informs the buyer's decision regarding how many additional suppliers to recruit. However, with the test strategy, the buyer may end up paying a high price for the units she buys in the first auction, while recruiting additional suppliers before running this auction would have resulted in potentially lower prices for these units. Below, we explain the mechanics and the timelines of the two strategies in more detail.

Test strategy. Under the test strategy, the buyer first holds an open-descending auction, what we call a *test auction* with a reserve price r_1, to auction off a portion of the Q units with only the n incumbent suppliers. As will be discussed in §2.2, we assume that the buyer promises to exclude from any subsequent auction those incumbents who do not meet r_1 in the test auction. There are three possible scenarios that can happen in the test auction: (a) no incumbent meets the reserve price; (b) exactly one incumbent meets the reserve price, or; (c) at least two incumbents meet the reserve price. In (a), the test auction concludes without a transaction, so the buyer has to auction off *all* Q units in a second auction where she recruits entrant suppliers and invites *only* those entrants to compete. In the second auction, the auction price is initially set at b (the buyer cannot use a lower reserve price in the second auction because she does not have time to perform a second recruitment round which would be necessary if zero entrants meet the reserve price in the second auction) and continuously descends until a single entrant remains in the auction, and that entrant wins all Q units at the per unit auction ending price. In (b), the single incumbent that met the test auction reserve price wins *all units of the test auction* at the reserve price r_1 per unit. At this point, the buyer recruits entrant suppliers to compete alongside this incumbent for the remaining units. When the second auction begins, its descending price clock starts at the ending price of the test auction, r_1, and continuously descends until a single supplier remains in the auction, and that supplier wins the remaining units at the ending price of the second auction. In (c), more than one incumbent meets the test auction reserve price, so the auction price clock continuously descends until a single incumbent remains in the auction, and that incumbent wins *all units of the test auction* at the test auction ending price. Similar to (b), the remaining units are auctioned off among the incumbents who met the test auction reserve price and any newly recruited entrants in a second auction, but with the descending price clock starting at the ending price of the test auction instead of

r_1. Note that in (b) and (c), if the clearing price of the test auction is sufficiently low, the buyer may choose to not recruit any entrants and instead simply purchase all Q units from the winning incumbent supplier at the test auction's clearing price. Note that when the test auction clears, i.e., in (b) and (c), its clearing price (ending price) acts as a de facto reserve price for the second auction as this price serves as the opening bid in the second auction. The buyer cannot use a lower reserve price in the second auction because the buyer does not have time to perform a second recruitment round which would be necessary if zero suppliers meet the reserve price. Figure 1 provides a graphical illustration of the test strategy timeline. (In the figure, "bidder $(i:j)$" refers to the bidder with the i^{th} lowest cost among j bidders.) The costs denoted in the figure are a consequence of the fact that, as will be established in the next section by Proposition 1, in equilibrium suppliers drop out at their true cost. Moreover, Proposition 1 also establishes that, given our model, the optimal quantity for the test auction is y units, ensuring that the buyer spends the minimal amount in order to learn the incumbents' costs.

Test without reserve price strategy. In the test strategy discussed above, the buyer can find the best test auction reserve price via optimization (see next section for more details). If the buyer is not that sophisticated and simply sets $r_1 = b$, then she is essentially running the test auction without a reserve price, and we call the resulting strategy *test without reserve price strategy.* We note that, in this strategy, the test auction always clears since $X_{(1:n)} \leq b$ (so the lower part in Figure 1 is never reached), and the clearing price of the test auction equals $x = \min\{b, X_{(2:n)}\} = X_{(2:n)}$.
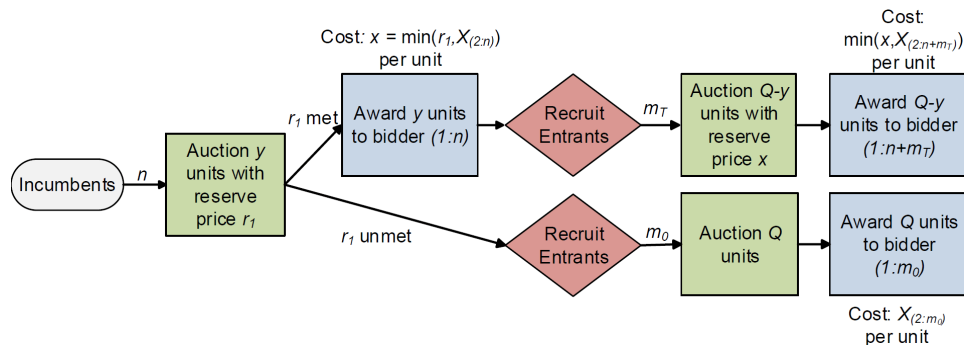


Figure 1 Timeline of the test strategy.

No-test strategy. Under the no-test strategy, the buyer first recruits entrants who compete alongside the incumbents in an open-descending auction. (Note that the buyer would

10

Beil, Chen, Duenyas, and See: *When to Deploy Test Auctions in Sourcing*
Article submitted to *Manufacturing & Service Operations Management*; manuscript no. MSOM-14-313

just hold one single auction for all Q units because she will not be able to initiate a second recruitment round and the competing suppliers all have linear cost structures.) The auction price is initially set at b, and continuously descends until a single supplier remains in the auction, and that supplier wins all Q units at the ending auction price per unit. Figure 2 depicts the timeline for the no-test strategy. That $X_{(2:n+m_N)}$ is the auction clearing price when m_N entrants are recruited will be established in §3.1 by Proposition 1.

Figure 2   Timeline of the no-test strategy.

2.2.   Discussion of Assumptions

Having explained the detailed mechanics and timeline of both strategies, we briefly discuss the assumptions. First, we explain why, in the test strategy, the buyer discards incumbents who do not meet the test auction reserve price. Allowing the incumbents to join the second auction even if they do not meet the test auction reserve price would put more bidders into competition in the second auction; however, it has the disadvantage that it reduces the information that the buyer gains in the test auction because when an incumbent does not meet the reserve price in the test auction, the buyer does not know if the incumbent's cost was higher than the reserve price or if instead the incumbent was being strategic and decided to forgo the chance to clinch the units auctioned off in the test auction in the hopes of getting a higher price in the second auction. In our test auction formulation where incumbents are excluded if they do not meet the reserve price, the buyer does not suffer from this informational ambiguity, and the test auction can be seen as a way to cleanly gather information about the incumbents before deciding how many entrants to recruit. However, the above discussion reveals that various test auction formulations are possible, and the analyses could become quite complex. This underscores the usefulness of finding a lower bound of the optimal cost (which will be investigated in §3.3) as a benchmark against which to evaluate the performance of our proposed test auction formulation where the incumbents who do not meet the test auction reserve price are discarded.

Second, although a reserve price can typically help the buyer avoid paying a high cost, the buyer must consider the possibility that the reserve price she sets in the test auction might not be met and thus she might not transact with the suppliers in the test auction. In practice, many buyers we have dealt with need to procure items that they are *not* capable of producing themselves. They rely on suppliers for production. Although the auction literature commonly uses the notion of an "outside option cost" (needed when setting a reserve price), such a cost typically is assumed to be exogenous; for instance, the cost of forgoing the contract. However, would the buyer forgo the contract if the reserve price is not met? In reality, if the reserve price is not met and the buyer has sufficient time to recruit new entrants, a buyer could re-run the auction with a new set of suppliers, as is captured in our model. To our knowledge, this is the first paper to inform this decision by endogenizing the cost of recruiting new suppliers and re-running an auction.

Lastly, we want to highlight that using the test auction's clearing price as a reserve price for the second auction could be executed via the following contract between the buyer and the winner of the test auction: the contract states that the buyer will purchase the units that are up for bid in the test auction at its clearing price, and has the right to later purchase additional units (i.e., the *remainder* of the Q units) at that same price. Such an arrangement assumes that the winner of the test auction has enough capacity available to supply additional units up for bid in the second auction and his cost does not change drastically between the two auctions. This reflects, for example, long-term fixed-price contracts lasting multiple years commonly seen in many industries such as automotive and energy infrastructure where the capacity is planned out in advance, minimum and maximum supply quantities are set, and the supplier agrees to pricing that is typically fixed over multiple years. For settings where one portion of the cost is stable, but the other portion of the cost may be very volatile such as the raw material cost of a commodity input, multi-year fixed-price contracts could be too risky for suppliers. Instead, the buyer and supplier can use a modified approach — an "index contract" — in which the buyer reimburses the supplier for the volatile cost of the raw material used, and the buyer could still use a test auction in supplier selection.

3.   Analysis

In this section, we first formulate the buyer's expected costs under the test and no-test strategies and then characterize situations when one is preferred to the other.

3.1.   Characterizing the Test and No-Test Strategies' Expected Costs

Proposition 1 below will show that in equilibrium suppliers remain in an auction until either winning or reaching their true cost, whichever happens first, and when using the test auction it is optimal for the buyer to auction off the minimum order quantity, y. We will use these facts when deriving the strategies' expected costs.

Test strategy. The buyer's optimal test auction reserve price depends on her expected "outside option" cost – the cost which she expects to incur if zero incumbents meet the reserve. In this case (lower path of Figure 1), as discussed in §2.2, we assume that the buyer removes all the incumbent suppliers from her supplier pool to enforce the inherent threat of the reserve price, and recruits new entrants and then procures all Q units from them in the second auction. Hence, the buyer's "outside option" is purchasing the Q units from the entrants, which in expectation costs her

$$C_0 \triangleq \min_{m_0 \geq 1} \left\{ \mathbb{E}[X_{(2:m_0)}]Q + k(m_0) \right\}, \tag{1}$$

where $X_{(2:1)} \triangleq b$, the upper bound of the suppliers' unit cost. Note that when $m_0 = 1$, the buyer recruits a single entrant supplier and sets a reserve price equal to b in order to award the Q-unit contract to the sole entrant at the upper bound of his unit cost distribution. (The buyer facing a single entrant in the second auction cannot do better than this, since she must transact but does not know the entrant supplier's true cost; see Myerson (1981).)

We now evaluate the buyer's expected cost (the sum of auction payment and supplier recruitment costs) under the test strategy when there is at least one incumbent whose cost is below the reserve price, i.e, $X_{(1:n)} \leq r_1$ (upper path of Figure 1). When the test auction clears at $x = \min\{r_1, X_{(2:n)}\}$, the buyer's payment for the first y units is x per unit. After holding the test auction, the buyer must decide how many additional suppliers to recruit. The cost information revealed in the test auction allows the buyer to make a more informed decision than she would if she did not hold a test auction. The expected cost of the second auction (including recruitment costs) when the test auction clearing price is x is

$$C(x) \triangleq \min_{m_T} \left(\mathbb{E}[\min\{x, X_{(2:n+m_T)}\} | X_{(1:n)} \leq x = \min\{r_1, X_{(2:n)}\}](Q - y) + k(m_T) \right), \tag{2}$$

and the optimal number of entrants to recruit $m_T^*(x)$ equals

$$m_T^*(x) \triangleq \arg\min_{m_T} \left(\mathbb{E}[\min\{x, X_{(2:n+m_T)}\} | X_{(1:n)} \leq x = \min\{r_1, X_{(2:n)}\}](Q - y) + k(m_T) \right). \tag{3}$$

Note that the expectation is written as the minimum of x and the second-lowest among $n + m_T$ draws, conditioned on the fact that the outcome of the test auction among the n incumbents was clearing price x (note that when the test auction clears we must have $X_{(1:n)} \le r_1$, and hence $X_{(1:n)} \le \min\{r_1, X_{(2:n)}\} = x$). Note that this is equivalent to the expected cost from an auction that opens at price x, with one bidder whose cost is at most x (the winning incumbent from the test auction) competing for the contract alongside m_T entrants. This is because Proposition 1 below shows that in equilibrium bidders in the test auction drop out of the auction at their true cost.

We now characterize the optimal reserve price for the test auction, r_1^*; specifically,

$$r_1^* \triangleq \underset{r_1 \in [a,b]}{\arg\min} \left(\overline{F}(r_1)^n C_0 + n F(r_1) \overline{F}(r_1)^{n-1} (r_1 \cdot y + C(r_1)) + \int_a^{r_1} n(n-1) f(s) F(s) \overline{F}(s)^{n-2} (s \cdot y + C(s)) ds \right)$$

(4)

where $\overline{F}(x) \triangleq (1 - F(x))$ is the complementary cumulative distribution function. In (4), the first term corresponds to the case where zero incumbents meet the reserve price, resulting in the buyer eliminating the incumbents from competition for the units and resorting to her "outside option" of recruiting new suppliers. The second term represents the case where a single incumbent meets reserve price r_1^*, resulting in a clearing price equal to the reserve price. The final term covers cases where at least two incumbent suppliers meet r_1^*.

No-test strategy. If the buyer bypasses the test auction and selects the no-test strategy (Figure 2), she first chooses how many entrant suppliers to recruit and a reserve price, and then holds a single open descending auction for all Q units. Note that the buyer must set the reserve price at the suppliers' per unit cost upper bound b (i.e., the auction opens at an auction price b) because she would not be able to recruit additional suppliers if the reserve price was unmet. Since all suppliers drop out at their true cost (as Proposition 1 below shows), the buyer's expected cost (the sum of the auction payment and recruitment costs) when she recruits m_N additional entrant suppliers is $\mathbb{E}[X_{(2:n+m_N)}]Q + k(m_N)$. Let m_N^* denote the minimizer of this cost; the buyer's total *ex ante* expected cost under the no-test strategy is thus $\mathbb{E}[X_{(2:n+m_N^*)}]Q + k(m_N^*)$.

The structure of the test and no-test strategies is characterized by the following proposition (proofs for our results are in the Online Supplement).

PROPOSITION 1. *In equilibrium, under the test and no-test strategies, suppliers in an auction will remain in it until either the price reaches their true cost or they win the*

auction, whichever happens first. Under the no-test strategy, the buyer's total cost function is discrete convex in the number of entrant suppliers she recruits, m_N. Under the test strategy, it is optimal for the buyer to auction off just y units in the test auction. If the test auction reserve price r_1 is met and the clearing price is $x = \min\{r_1, X_{(2:n)}\}$, the expected cost of recruitment and the second auction payment is discrete convex in the number of entrant suppliers the buyer recruits, m_T, and the optimal number of entrants to recruit, m_T^*, is nondecreasing in the realization of the clearing price of the test auction x. If the test auction reserve price is not met, in the second auction, the buyer will recruit the same number of entrants as if she were using the no-test strategy to auction off all Q units with zero incumbents to start with.

The proposition proves that the buyer's expected cost under both strategies is discrete convex in the number of entrants to recruit, which simplifies the buyer's recruitment choice (because if the buyer prefers recruiting m entrant suppliers to recruiting $m + 1$ entrant suppliers, she also prefers recruiting m to any $m' > m + 1$). Further, the result states that under the test strategy the buyer will recruit a relatively small number of entrants if she realizes that the price she has to pay the winning incumbent is low, and she will be willing to recruit a larger number of entrants if she realizes the incumbents have high costs. This occurs because the expected benefit of recruiting an entrant supplier (namely, the reduction in the expected clearing price of the second auction for $Q - y$ units) increases as the clearing price of the test auction increases.

Although the result that the number of entrants to recruit is nondecreasing in the clearing price of the test auction may seem intuitive at first, it may not hold if our model's regularity assumption is violated. For example, consider the case where the buyer needs to source 1,000,000 units and the suppliers' cost distribution pdf $f(x)$ equals: 0.1 if $x \in [8, 9]$; 0.9 if $x \in [11, 12]$; 0 elsewhere. This cost distribution results in a uniformly-distributed cost on $[8, 9]$ with probability 0.1 and a uniformly-distributed cost on $[11, 12]$ with probability 0.9. This represents the case where most suppliers are likely to have a relatively high cost, but a few may have a lower cost (e.g., because they have lower labor costs or have better-suited production technology). Suppose the buyer has already sourced 300,000 units through an auction with the incumbent suppliers and she is deciding how many entrants to recruit before holding the second auction for the remaining units, and let the recruitment cost be $k(m) = 154,000 \cdot m$ — that is, the buyer's marginal cost of recruiting an entrant is constant.

If the clearing price of the test auction is $x = 11$, the buyer will find it optimal to recruit one entrant prior to the second auction. However, if the clearing price of the test auction is $x = 11.1$, the buyer will find it optimal to recruit zero entrants — the number of entrants to recruit *decreases* when the clearing price of the test auction increases!

When the clearing price is $x = 11$, the buyer is certain that the winning incumbent supplier has a cost in $[8, 9]$ with probability 1, and is willing to recruit an entrant because with probability 0.1 the entrant will also have a low cost, resulting in significant savings in the second auction for the buyer. However, when the clearing price is $x = 11.1$, the buyer is uncertain of whether the winning incumbent is in the $[8, 9]$ or $[11, 11.1]$ cost region (in fact, the winning incumbent has a cost in the higher region with probability 0.47). Thus, the buyer does not find it optimal to recruit entrant suppliers because the expected unit cost savings are less than the associated recruitment cost. This example reveals that although it is intuitive that the number of entrants to recruit should increase in the clearing price of the test auction, this is not a foregone conclusion. What Proposition 1 shows, however, is that m_T^* increases in x under our regularity assumption, namely $x + \frac{F(x)}{f(x)}$ increases in x.

We finally note that it is never optimal for the buyer to use the no-test strategy and recruit $m_N^* = 0$ entrant suppliers. Such a policy is weakly dominated by using the test strategy: Instead of choosing to not recruit additional suppliers prior to procuring the entire batch of Q units, the buyer could set $r_1 = b$ and use the test strategy and do no worse in expectation. Thus, buyers who typically use only existing incumbent suppliers for new components without testing their prices first or recruiting entrants may have an opportunity to lower costs by running a test auction.

3.2. Comparison of the Strategies

To understand the trade-off between test and no-test, consider the following numerical example. Assume the suppliers' unit costs are uniformly distributed between \$5 and \$15, the buyer has 5 incumbent suppliers, needs to source 50,000 units, the minimum order quantity is 2,500 units, and the cost to recruit each entrant supplier is \$17,500 (e.g., this could be the cost of time spent gathering supplier lists, contacting suppliers and sifting through them for fit, working with engineering to explain the RFQ to them, training them on the auction platform, etc.). Under this scenario, if the buyer chooses the no-test strategy, she will find it optimal to recruit $m_N^* = 2$ entrant suppliers, and then hold one auction for

all 50,000 units with the 7 suppliers. As a result, she will pay \$35,000 in recruitment costs and will expect to pay \$7.5 per unit, for a total expected cost of \$410,000.

Under the test auction approach, if the buyer is not sophisticated enough to use the optimal reserve price, but uses the test without reserve price strategy instead, she will first hold an auction for the minimum order quantity with the 5 incumbent suppliers and then decide how many entrant suppliers to recruit based on the clearing price of the test auction. Per Proposition 1, the optimal number to recruit increases in the outcome of the test auction. In this example, she will recruit anywhere from 0 to 5 entrant suppliers, where the optimal number to recruit is illustrated by Figure 3. In this case, even when the buyer is not sophisticated enough to optimize the test auction reserve price, the test without reserve price strategy results in an expected total cost (expected payment in both auctions and expected recruitment costs) of \$397,013 — a *net savings*, defined as the percent savings after subtracting the minimum possible cost $a \cdot Q$ (the buyer will incur $a \cdot Q$ regardless of her strategy, and thus we do not consider it when comparing her options), of 8.1%.

Figure 3   Optimal number of entrants to recruit given the clearing price of the test auction.

If the buyer uses the more sophisticated strategy of optimizing her test auction reserve price, she will solve (4) and use the optimal reserve price $r_1^* = \$7.5$ in the test auction. If the reserve price is met, then the clearing price of the test auction is no higher than \$7.5; so by Figure 3, in this case the buyer would recruit zero additional suppliers and source all Q units from the lowest-cost incumbent. On the other hand, if the reserve price is not met, the buyer would find it optimal to recruit 7 entrants and hold an auction for all Q units among the entrants. In this case, the test strategy results in an expected total cost of \$386,410 — a net savings of 14.7% over the no-test strategy.

Finally, since buyers in practice may be tempted to just auction a contract without running a test auction or recruiting additional suppliers, we consider the case where the buyer holds a single auction, only inviting the incumbents. In this case, implementing the test without reserve price strategy (resp. test strategy) results in a net savings of 11.8% (resp. 18.2%) over this suboptimal policy. Once again, in an environment where 2-3%

savings can amount to millions of dollars a year, the advantages of a more sophisticated strategy as proposed in this paper can be significant.

The following comparative statics provide insights that are helpful for buyers wondering if a test auction approach is appropriate for their situation.

PROPOSITION 2. *(i) Holding all other parameters constant, the buyer prefers the test without reserve price strategy over the no-test strategy if: (a) The minimum order quantity, y, is sufficiently small; (b) The number of incumbents, n, is sufficiently large; (c) The new supplier recruitment cost, $k(m)$, grows proportionately large enough (i.e., the recruitment cost is taken to be $\gamma \cdot k(m)$ and $\gamma \in \mathbb{R}^+$ is sufficiently large); (d) The cost distribution variance factor α is sufficiently small, where $\tilde{f}(x) = \frac{1}{\alpha} f(\frac{x+\Delta}{\alpha})$ with $\Delta \in \mathbb{R}$ and $\alpha \in \mathbb{R}^+$.*

(ii) Moreover, if the buyer prefers the test without reserve price strategy to the no-test strategy, then the buyer prefers the test strategy to the no-test strategy.

To understand part (a), note that a smaller y makes the test auction more attractive because the risk with a test auction only among the incumbents is that an incumbent wins an order at a high price. However, as y decreases, the amount risked becomes smaller. Similar intuition explains part (b): When the number of incumbents is large, there is more competition during the test auction stage and the buyer is less likely to pay a high price for the first y units, increasing the attractiveness of the test auction. Part (c) follows because when it is expensive to recruit suppliers, the buyer needs to be more judicious about how many entrants to recruit. Using a test auction helps her decide the exact number of entrants to recruit and avoids the risk of recruiting too many, which would be expensive when the function $k(m)$ grows proportionately larger. Part (d) appears counterintuitive at first sight. After all, one may think that price discovery is more beneficial when prices are more variable. The flip side of the coin is that when prices are less variable, the firm has a lower chance of paying too much for the first y units in a test auction and the test auction is less risky. Finally, part (ii) of the proposition follows from the fact that allowing the buyer to optimize the test auction reserve price never makes her worse off.

Finally, note that our proposed strategy is the minimum of the test and no-test strategies since in practice, based on the problem's parameters, the buyer can compute a priori the expected costs under both test and no-test strategies and choose the cheaper strategy to implement. Although our proposed strategy is easy to implement, one may wonder whether

more sophisticated auction mechanisms can achieve significant savings over our proposed strategy. We investigate this in the next two subsections.

3.3. A Lower Bound of the Optimal Cost

In the previous subsection, we investigated when the buyer can benefit from implementing the test strategy versus the more conventional no-test strategy. Although we saw that employing test auctions can provide significant benefits, it would be interesting to see how well our proposed strategy performs compared to the optimal cost. Note that, as we mentioned in §2.2, although we have assumed in the test strategy that the buyer discards incumbents who do not meet the test auction reserve price, if the buyer were designing a mechanism from scratch, she could potentially invite incumbents who do not meet the reserve price to join the second auction. Therefore, in this subsection, we employ an optimal mechanism analysis and seek the optimal mechanism among a *broader* class of feasible mechanisms where the only constraint we place is that the buyer must award at least y units to a supplier if she awards any units to him, and can only run one round of recruiting (this is because, as stated in the Introduction and the Model sections, there is only time for one round of recruitment, so the buyer *cannot repetitively* run test auctions and recruit more entrants after each test auction depending on the cost realizations, but can only do this once). Note that this class of feasible mechanisms subsumes the test strategy and the other mechanism discussed above where all incumbents are always allowed into the second auction. However, it also includes mechanisms that are potentially very hard to implement in practice (e.g., the buyer uses the test auction solely as a cost discovery scheme and then awards all Q units to the winner of the second auction *after* recruiting additional entrants). Hence, the optimal mechanism serves as a lower bound of the lowest possible cost the buyer can ever achieve in practice. (For ease of reference, we call it the "optimal mechanism" even though it should be understood as a lower bound of the optimal cost.)

We consider a buyer with full commitment power. As explained in the Online Supplement, applying a multi-stage version of the revelation principle (McAfee and McMillan 1988) allows us to confine our search for an optimal mechanism to truthful mechanisms with two stages of communication: in stage 1 the incumbents report their true cost to the buyer; in stage 2 the buyer decides how many entrants to recruit, and these entrants report their true cost to the buyer; the buyer then makes contract allocation and payment decisions based on the costs of the incumbents, and recruited entrants (if any).

To formalize, let $x^I \triangleq (x_1, \ldots, x_n)$ denote the vector of all incumbent costs and let x^I_{-i} denote the cost vector excluding incumbent i for any $i = 1, \ldots, n$. Moreover, let $m_T(x^I)$ denote the buyer's decision on the number of entrants recruited in stage 2, given the vector of costs reported by incumbents in stage 1. When l entrants are recruited (i.e., when $m_T(x^I) = l$) we define $x^{E,l} \triangleq (x_{n+1}, \ldots, x_{n+l})$ as the vector of the recruited entrants' costs and let $x^{E,l}_{-i}$ denote the cost vector excluding entrant i for any $i = n+1, \ldots, n+l$. Based on the stage 1 and stage 2 reports, when l entrants are recruited, the payment and allocation given to suppliers $i = 1, \ldots, n+l$ are denoted $P_i(x^I, x^{E,l})$ and $Q_i(x^I, x^{E,l})$. For vector z, let $f(z) = \prod_{i=1}^{|z|} f(z_i)$ be the joint distribution of $(z_1, \ldots, z_{|z|})$. Let $u_i(s_i)$, $p_i(s_i)$, $q_i(s_i)$ denote supplier i's expected profit, payment, and allocation when his strategy is to report s_i when his true cost is x_i, and all other suppliers report truthfully. Specifically, $u_i(s_i) \triangleq p_i(s_i) - x_i q_i(s_i)$, and, by the law of total expectation, for any incumbent $i = 1, \ldots, n$,

$$p_i(s_i) \triangleq \sum_{l=0}^{\infty} \int_{\{x^I_{-i} | m_T(s_i, x^I_{-i}) = l\}} \int_{x^{E,l}} P_i((s_i, x^I_{-i}), x^{E,l}) f(x^{E,l}) dx^{E,l} f(x^I_{-i}) dx^I_{-i},$$

$$q_i(s_i) \triangleq \sum_{l=0}^{\infty} \int_{\{x^I_{-i} | m_T(s_i, x^I_{-i}) = l\}} \int_{x^{E,l}} Q_i((s_i, x^I_{-i}), x^{E,l}) f(x^{E,l}) dx^{E,l} f(x^I_{-i}) dx^I_{-i},$$

where $\{x^I_{-i} | m_T(s_i, x^I_{-i}) = l\}$ is the set of cost realizations x^I_{-i} such that, given (s_i, x^I_{-i}), the buyer will recruit l entrants under the recruitment rule m_T. The expressions for the expected payment and allocation for any entrant supplier are similar, and we omit them for brevity. The buyer's mechanism design problem is

$$\min_{Q_i(\cdot), P_i(\cdot), m_T(\cdot)} \sum_{l=0}^{\infty} \int_{\{x^I | m_T(x^I) = l\}} \left[k(l) + \sum_{i=1}^{n+l} \int_{x^{E,l}} P_i(x^I, x^{E,l}) f(x^{E,l}) \mathrm{d}x^{E,l} \right] f(x^I) \mathrm{d}x^I \quad (5)$$

$$\text{s.t.} \quad u_i(x_i) \geq 0 \qquad\qquad\qquad\qquad\qquad \forall x_i, \forall i \quad (6)$$

$$u_i(x_i) \geq u_i(s_i) \qquad\qquad\qquad\qquad \forall s_i \neq x_i, \forall x_i, \forall i \quad (7)$$

$$\sum_{i=1}^{n+m_T(x^I)} Q_i(x^I, x^{E, m_T(x^I)}) = Q \qquad\qquad \forall x^I, \ \forall x^{E, m_T(x^I)} \quad (8)$$

$$Q_i(x^I, x^{E, m_T(x^I)}) \in \{0, y, y+1, \ldots, Q\} \qquad \forall x^I, \ \forall x^{E, m_T(x^I)}, \ \forall i. \quad (9)$$

The objective function (5) states that the buyer minimizes her expected total cost (recruitment costs plus payments to suppliers) by choosing the quantity allocation rule, payment rule, and entrant recruitment rule, $m_T(x^I)$. Note that, by law of total expectation,

we write the expected total cost in (5) by conditioning on how many entrants are recruited under the recruitment rule. Constraint (6) ensures individual rationality, while (7) ensures incentive compatibility, i.e., truthful reporting is a Bayesian Nash equilibrium. Constraint (8) states that the desired Q units must be procured, while (9) ensures that as in §2, an incumbent or entrant who is successfully awarded business gets at least the minimum order quantity, y. The quantity allocation, payment, and recruitment rules that solve this constrained cost minimization problem will serve as a lower bound on the expected cost the buyer could achieve regardless of sourcing mechanism — whether it be the test strategy, no-test strategy, or any other sourcing mechanism (e.g., negotiation, first-price sealed-bid auction, Dutch auction, etc.). Let $\psi(x_i) = x_i + \frac{F(x_i)}{f(x_i)}$ denote the virtual cost function, where $\psi(\cdot)$ is increasing because of the regularity assumption.

PROPOSITION 3. *The following constitutes an optimal mechanism: Based on the incumbent suppliers' reported cost vector x^I, the number of entrants recruited by the buyer is*

$$m_T^*(x^I) = \underset{m_T(x^I) \in \mathbb{N}}{\arg \min} \left\{ k(m_T(x^I)) + \mathbb{E}[\psi(X_{(1:n+m_T(x^I))}) | x_{(1:n)}^I] \cdot Q \right\} \tag{10}$$

where m_T^* is nondecreasing in $x_{(1:n)}^I$ and constant in $x_{(2:n)}^I, x_{(3:n)}^I, \ldots, x_{(n:n)}^I$. Define

$$\tilde{x}^I \triangleq \begin{cases} [t_i, x_{(2:n)}^I, x_{(3:n)}^I, \ldots, x_{(n:n)}^I] & \text{if} \quad x_{(1:n+m_T^*(x^I))} = x_{(1:n)}^I, \\ x^I & \text{otherwise.} \end{cases}$$

Let i be the lowest-cost supplier, i.e., $x_i = x_{(1:n+m_T^*(x^I))}$. The buyer pays i

$$P^*(x^I, x^{E, m_T^*(x^I)}) = x_{(1:n+m_T^*(x^I))} \cdot Q + \int_{x_{(1:n+m_T^*(x^I))}}^{x_{(2:n+m_T^*(x^I))}} (1 - F(t_i))^{(m_T^*(\tilde{x}^I) - m_T^*(x^I))} \cdot Q \, dt_i \tag{11}$$

to supply the Q units while the other suppliers are not awarded any units and are paid zero.

Under this mechanism, the buyer views the incumbents' vector of reported costs, x^I, and recruits $m_T^*(x^I)$ entrants based on these reports. After the $m_T^*(x^I)$ entrants report their costs, the buyer awards the Q-unit contract to the supplier with the lowest report and pays him $P^*(x^I, x^{E, m_T^*(x^I)})$. The first term of (11) is the winning supplier's true cost, while the second term represents a markup that rewards the supplier for truthfully revealing his cost. If the winning supplier is an entrant supplier, we note that $\tilde{x}^I = x^I$ and the markup is equal to the difference between the second-lowest report and the winning supplier's report. Thus,

in the case where an entrant is awarded the contract, $P^*(x^I, x^{E,m_T^*(x^I)}) = x_{(2:n+m_T^*(x^I))} \cdot Q$ and the winning entrant supplier is paid the second-lowest reported cost.

However, if the winning supplier is an incumbent, then the markup represents the amount by which the incumbent supplier could have inflated his reported cost and still won the contract. The important point is that a winning incumbent supplier's reported cost also influences the buyer's recruitment decision. Hence, the markup of an incumbent supplier who eventually wins the contract must account for the increased competition (i.e., $m_T^*(\tilde{x}^I) - m_T^*(x^I)$ additional entrants) associated with an inflated reported cost.

Although the optimal mechanism described in Proposition 3 results in the minimum expected cost for the buyer, the buyer may be unable to use such a mechanism in practice. The described mechanism requires the incumbent suppliers to reveal their true cost *before* the buyer recruits suppliers to compete for the units up for bid. Revealing such information prior to the recruitment round will not result in a contract (even for the minimum order quantity) for the incumbents — it will only help the buyer make a more accurate decision on how many entrants to recruit. A buyer may find it difficult to convince a supplier to reveal his cost information just so the buyer can decide how many suppliers to recruit to compete against him. In contrast, in the test strategy, the guaranteed award of y units to the winning incumbent supplier (given he meets the reserve price) provides an incentive for the incumbents to reveal such information — it is the minimum quantity required to induce suppliers to bid. Moreover, even in the case when $y = 0$, the buyer may still be unable to implement this optimal mechanism since the payment rule for the incumbents is very complicated and would be hard to explain to practitioners. Finally, we note that in the test strategy bidders have to beat their competitors' pricing but do not have to reveal their true cost to win; under the mechanism described in Proposition 3, the lowest-cost supplier still has to reveal his cost.

Regardless of whether the buyer would be able to implement the optimal mechanism in practice, the optimal mechanism can be used as a theoretical benchmark to quantify the relative performance of the test and no-test strategies; the "optimality gap" between the optimal mechanism and the test and no-test strategies which we investigate in §3.4 can help the buyer gauge whether the cost and effort of attempting to implement more complex mechanisms than the test and no-test strategies are worth the savings.

3.4. Numerical Studies

In this section, we use an extensive numerical study to investigate two important questions: (i) how effective is our proposed strategy compared to the lower bound of the optimal cost derived in the previous section (what we will refer to as the optimal mechanism), and; (ii) how and why does the answer to (i) depend on the underlying problem characteristics.

Effectiveness of our proposed strategy. To address (i), we calculated the test and no-test strategies' performance relative to the optimal mechanism in a wide range of situations by generating 840 problem scenarios with different model parameters (i.e., production cost distribution, recruitment cost functions, number of incumbents and minimum order quantity). Table 1 below summarizes the factorial design of our numerical study.

Table 1 Factorial design of the numerical study. Distribution types are short coded as follows: U1, U2, U3 are uniform distributions with supports [\$7.5, \$12.5], [\$5, \$15], [\$2.5, \$17.5] and coefficient of variations of 0.14, 0.29, 0.43 respectively; N1, N2, N3 are normal distributions with mean \$10 and coefficient of variation of 0.3, 0.4, 0.5 respectively which are then truncated to have support [\$2.5, \$17.5]; E is exponential distribution with mean \$10 which is then truncated to have support [\$2.5, \$17.5].

Parameters	Values
Distribution type	$U1, U2, U3, N1, N2, N3, E$
Number of incumbents	2, 5, 8, 11
Minimum quantity	500, 2,500, 4,500, 10,000, 20,000
Recruitment cost function	$\$7,500m, \$12,500m, \$17,500m, \$22,500m, \$27,500m, \$7,500m^6$

We report key summary statistics of the net savings of the optimal mechanism over our proposed strategy for all 840 scenarios in Table 2. Note that the optimal mechanism only provides an average net savings of 1.6% over our proposed strategy (i.e., the minimum of the test and no-test strategies); even at the 99^{th}-percentile, the net savings is no more than 3%. This suggests that, although our proposed strategy is simple and easy to implement, it provides near-optimal results and there is very limited room for improvement by using a more sophisticated mechanism.

Moreover, although the no-test strategy may outperform the test strategy in some cases (as the 1^{st}-percentile of the net savings of the proposed strategy over the no-test strategy equals *exactly* 0%), we want to stress the significance of incorporating the test strategy by noting that using the no-test strategy alone is not sufficient to achieve near-optimal results. In fact, the optimal mechanism provides a substantial average net savings of 8.2% over the no-test strategy (this is to be compared with the 1.6% under our proposed strategy), and

Table 2   **Summary statistics for the net savings of proposed strategy over the no-test strategy, net savings of the optimal mechanism over the no-test strategy, and net savings of the optimal mechanism over the proposed strategy**

	Proposed over No-Test	OPT over No-Test	OPT over Proposed
Mean	6.8%	8.2%	1.6%
Standard deviation	5.0%	5.5%	0.7%
99^{th}-percentile	17.3%	19.7%	3.0%
75^{th}-percentile	10.9%	12.6%	2.0%
50^{th}-percentile	6.2%	7.4%	1.6%
25^{th}-percentile	2.0%	3.4%	1.0%
1^{st}-percentile	0.0%	0.2%	0.2%

our proposed strategy provides an average net savings of 6.8% over the no-test strategy. By incorporating the test strategy in her sourcing toolkit, the buyer can significantly reduce her expected procurement cost. Another way to illustrate the usefulness of the test strategy option is to compute the fraction of the potential absolute savings the optimal mechanism offers over the no-test strategy that can be captured by our proposed strategy. To compute this metric, we aggregate the savings of our proposed strategy over the no-test strategy across all scenarios, and divide this by the aggregated potential savings using the optimal mechanism. (This aggregated performance metric is more useful than simply taking the average of the aforementioned fraction for all scenarios because aggregation allows us to capture the heterogeneity of the absolute potential savings across scenarios, i.e., our proposed strategy capturing 60% out of an absolute potential savings of 10 cents is vastly different from capturing 60% out of an absolute potential savings of 100 dollars.) Consistent with our previous observations, our proposed strategy performs very well, capturing 83.5% of the aggregated potential savings from the optimal mechanism which as we mentioned would be hard to implement in practice.

The take-away from the numerical results above is important to reiterate: The buyer can use an easily-implementable strategy with an auction mechanism very familiar to practitioners that will provide average net savings of 6.8% compared to the no-test strategy, and furthermore this easily-implementable strategy captures 83.5% of the aggregated absolute savings that the optimal mechanism can offer over the no-test strategy. Thus, we believe that our proposed strategy is a powerful sourcing mechanism for buyers.

Sensitivity analysis. Our previous numerical study does not reveal how the performance metric varies across all scenarios. To investigate that, we conduct sensitivity analyses below by varying the parameters one at a time for a wide range of values.

First, we considered linear recruitment cost functions whose (constant) marginal cost of recruiting an entrant was varied from \$2,000 to \$120,000: Figure 4(a) shows that the net savings of the optimal mechanism over our proposed strategy is increasing and then decreasing in the recruitment cost, and this pattern is consistent across different distributions. For extremely low recruitment costs, the buyer will use the no-test strategy and recruit a large number of entrants prior to holding an auction; likewise, under the optimal mechanism the buyer will (most likely) recruit many entrants after viewing the incumbents' costs. As the recruitment cost increases, the value of viewing the incumbents' costs increases until a certain point, after which the value decreases because the buyer will be less inclined to recruit entrants due to the recruitment cost.

Second, the number of incumbent suppliers was varied from 2 to 20. Figure 4(b) shows that the net savings of the optimal mechanism increases and then decreases in the number of incumbents, and this pattern is consistent across different distributions. Using intuition similar to the recruitment cost case, for very small and large n, the buyer's recruitment decision hinges less on the incumbents' cost information; the buyer will be more inclined to recruit many (in the case of very small n) or few (in the case of large n) entrants regardless of the cost information, so the value of the optimal mechanism with respect to our proposed strategy decreases.

We next wanted to investigate how the minimum order quantity y affected the suboptimality of our proposed strategy. Conceivably, if the optimal reserve price of the test strategy is low enough, then whenever the test auction clears, the buyer will simply award all Q units to the winning incumbent and skip recruiting entrants. In this case, the cost of the test strategy and the net savings of the optimal mechanism over our proposed mechanism will be insensitive to y. Indeed, in the special case of linear recruitment cost function, this happens as illustrated in Figure 5(a). However, this insensitivity does not hold in general: In contrast, Figure 5(b) shows that when the recruitment cost function is strictly convex (i.e., $k(m) = 7500m^6$), the net savings of the optimal mechanism first strictly increases as y increases but then remains constant when y is large. As the minimum order quantity increases, the cost of the test strategy increases because (i) fewer units are left to enjoy a low price in the second auction, and more units are exposed to the risk of a high clearing price in the test auction, and (ii) the chance of resorting to the "outside option" when the test auction fails to generate a winner increases due to a lower optimal reserve price.
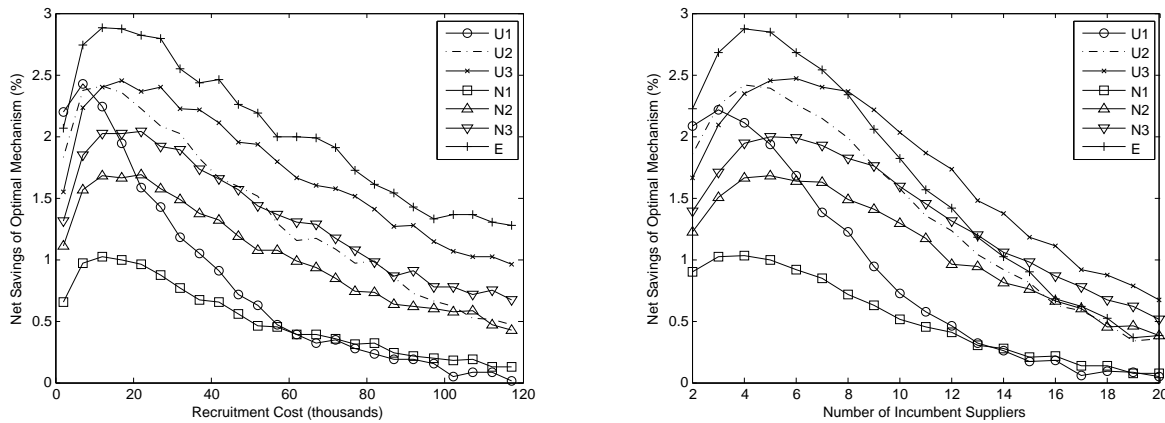
(a) Net savings as a function of the constant marginal recruitment cost.

(b) Net savings as a function of the number of incumbent suppliers.

Figure 4   The buyer's net savings from employing the optimal mechanism instead of the minimum of the no-test and test strategies. Parameters are $Q = 50{,}000$, $y = 4{,}500$, $k(m) = \$17{,}500m$, and $n = 5$, unless otherwise noted.

(a) Net savings as a function of the minimum order quantity for a linear recruitment cost function $k(m) = \$17{,}500m$.

(b) Net savings as a function of the minimum order quantity for a strictly convex recruitment cost function $k(m) = \$7{,}500m^6$.

Figure 5   The buyer's net savings from employing the optimal mechanism instead of the minimum of the no-test and test strategies. Parameters are $Q = 50{,}000$ and $n = 5$.

Therefore, the net savings of the optimal mechanism, which gathers cost information for free (and is thus insensitive to y), is increasing in y when our proposed strategy chooses the test strategy (i.e., when y is small). When y is large, our proposed strategy chooses the no-test strategy whose cost is independent of y, so the net-savings is constant in y.

(a) Net savings as a function of the coefficient of variation of uniform distributions whose means equal $10.

(b) Net savings as a function of the coefficient of variation of normal distributions (before truncation) whose means equal $10 and are truncated to have the same support [$2.5, $17.5].

Figure 6   **The buyer's net savings from employing the optimal mechanism instead of the minimum of the no-test and test strategies. Parameters are** $Q = 50,000$, $y = 4,500$, $k(m) = \$17,500m$, **and** $n = 5$.

Finally we wanted to explore how the coefficient of variation (CV) of the underlying cost distribution for suppliers affected the suboptimality of our proposed strategy. Exponential distributions always have a coefficient of variation of 1. But for the uniform and normal distribution cases, we can vary the CV by adjusting the width of the support for the uniform case, and the standard deviation for the normal case. This created a series of distributions with varying CVs. For the normal case, we then truncated the distribution to have support [$2.5,$17.5] to ensure a finite support. Figures 6(a) and 6(b) show that, for both the uniform and the normal distributions, the optimal mechanism's net savings grow at a decreasing rate as the CV of the suppliers' cost distribution increases; as the variability of the incumbents' costs increase, the value of being able to view the incumbents' cost information without having to allocate any quantity to incumbents (which is what the optimal mechanism does) increases due to the increased uncertainty in cost.

4.   Extensions
4.1.   Other Auction Mechanisms
We studied the open-bid descending-price auction format due to its prevalence in industry, the intuitive nature of its mechanics, and the attractive simplicity of the suppliers' optimal bidding strategy (as proved in Proposition 1). However, there are a variety of different auction formats that could be used in lieu of the open-bid descending-price auction.

Another auction format that is studied in the literature and also used in practice is the first-price sealed-bid auction. Unfortunately this auction format is generally intractable analytically for our test-strategy setting due to bidder asymmetry resultant from bidders' belief updating: Suppose an incumbent supplier wins the test auction at clearing price x. The buyer then decides how many entrant suppliers to recruit (say m^*) to compete against this winning incumbent supplier in the second auction for the remaining $Q - y$ units. The second auction will then be an auction with ex ante asymmetric bidders. Namely, each supplier is aware that at least one other supplier has a different cost distribution: Given the number of recruited entrants and the opening price (i.e., the reserve price) of the second auction, the entrant suppliers can update their belief regarding the incumbent's cost.

In general, an equilibrium bidding strategy cannot be found in such asymmetric first-price sealed-bid auctions (see, e.g., Hartline et al. (2014), Maskin and Riley (2000)). A variety of papers examine the two-bidder case and find equilibrium bidding strategies for the two asymmetric bidders (e.g., Mares and Swinkels (2014) and the references therein). Of course, in our case there may be more than two bidders. Lebrun (1999) finds positive results for the n bidder case where suppliers' cost distributions have supports equal to the same interval. This is not the case in our setting, as the winning incumbent supplier's cost distribution will have support on $[a, x]$ while the entrant suppliers' cost distribution will have support on $[a, b]$. Finally, the case closest to our model — where all bidders except one have the same valuation probability measure — was numerically examined by Marshall et al. (1994). However, to our knowledge no analytical results exist for this case. The same applies for the reverse Dutch auction (i.e., the other auction format the authors have seen used in practice whereby a price clock rises until a bidder ends the auction by accepting the contract), as this format is strategically equivalent to a first-price sealed-bid auction.

Regardless of these intractability issues, for these other auction formats used in practice the following result generalizes Proposition 2, which described the buyer's preference between the test and no-test strategies. To account for the role of belief updating during the test strategy, we will assume the buyer has discretion over whether to announce the number of entrants she recruits for the second auction and whether to announce an opening bid for the second auction, and can commit to such strategies, as such announcements reveal information that the entrants could use when updating their beliefs.

PROPOSITION 4. *If the buyer uses sealed-bid first-price or reverse Dutch auctions instead of open-bid descending-price auctions, then Proposition 2 holds.*

Finally, we note that the relatively small average optimality gap (1.6% net savings in our numerical study) between the second-price open-bid auction formats and the optimal mechanism studied in §3 should reassure practitioners that the use of other contracting mechanisms would provide at most minimal savings over our proposed strategy.

4.2. Test Auction Strategy with Non-Independent Costs

We have emphasized the buyer's opportunities for cost savings by using a test auction strategy that allows the buyer to gain more information regarding incumbent suppliers' costs before making a recruitment decision. The preceding analysis assumed a private values framework, where suppliers' costs are independent and identically distributed according to cumulative distribution F with probability distribution f. Further, the cost distribution F is known to all suppliers and the buyer. Thus, in our model the incumbent cost information gathered in the test auction does not influence the buyer's belief regarding the entrant suppliers' costs. However, there may be certain situations where this is not the case; the buyer may be able to update her belief on the entrant suppliers' cost distribution based on the incumbent suppliers' bid-down-to levels in the test auction. For example, all suppliers may share similar costs for things like commodity input prices or shipping/freight costs. Thus, suppose the buyer has an initial belief that the suppliers' cost distribution is F, and to update her belief after the test auction whose reserve price is set at r_1, the buyer refines her belief to the distribution $F(\cdot|r_1; b_1, b_2, \ldots, b_{n-1})$ with density $f(\cdot|r_1; b_1, b_2, \ldots, b_{n-1})$, where b_i represents the i^{th} drop-out bid (i.e., the data the buyer observes) in the test auction. (If an incumbent does not meet the reserve price, then his drop-out bid is "N/A" for not applicable.)

In our original model we showed that incumbents had a dominant strategy of remaining in an auction until either winning or reaching their true cost. This bidding strategy was the bidders' best response to the buyer's recruitment strategy, where the number of recruited entrants increases in the test auction's clearing price. However, in the setting where the buyer updates her belief on the underlying supplier cost distribution, things become more complex. In such a case, even if the incumbents bid down to their true cost, the buyer's recruiting strategy need not have the same property as before. In fact, the number of

entrants the buyer recruits could conceivably be decreasing in the test auction clearing price. For example, imagine a case where the clearing price of the auction is high, but the buyer's updated belief confirms that an entrant supplier will be unlikely to have lower costs than the clearing price of the test auction. Even though the clearing price is high, the buyer will recruit fewer entrant suppliers than she might under an alternative scenario where the clearing price is low but the buyer anticipates that entrant suppliers will provide a cost savings. In the context of this example, one can see how incumbent suppliers may want to alter their bidding strategy to influence the buyer's recruitment decision. Therefore, to ensure that incumbents still find it optimal to bid aggressively in the test auction, the buyer adds an additional element to her auction policy: She promises to only allow the test auction winner to participate in the second auction. (By contrast, despite being mild, this assumption was not required in the independent costs setting studied in §3.1–3.2 so we did not impose it there.) Under this more stringent policy, the incumbent bidders can be shown to have a dominant strategy of remaining in the auction until either winning or reaching their true cost. We also note that, unlike before, the buyer's recruitment strategy need not depend on just the test auction's clearing price; indeed, one can easily see that the buyer's recruitment policy depends on her updated belief about the underlying cost distribution, $F(\cdot | r_1; b_1, b_2, \ldots, b_{n-1})$, which would in general depend on the drop-out bids of the incumbents that the buyer observes in the test auction and the test auction reserve price r_1. The following result formalizes these insights.

PROPOSITION 5. *Suppose that under the test strategy the buyer uses r_1 as the test auction reserve price and only allows the test auction winner (if any) to compete with entrant suppliers in the second auction. Then, when the buyer can update her belief about the suppliers' cost distribution:*

(i) Suppliers in an auction will remain in it until either the price reaches their true cost or they win the auction, whichever happens first.

(ii) Under the test strategy, if the test auction clears (resp. does not clear), then the optimal number of suppliers to recruit, m_T^* (resp. m_0^*), is a function of the test auction reserve price r_1 and the drop-out bids $b_1, \ldots, b_{n-1}$ of the $n-1$ incumbent suppliers who do not win the test auction, where b_i equals $X_{(n-i+1:n)}$ if $X_{(n-i+1:n)} \leq r_1$, or "N/A" otherwise.

Unlike Proposition 1, which stated that for our original model the number of entrants the buyer recruits is solely dependent on the test auction's clearing price, we note that both

the test auction reserve price r_1 and the drop-out bids of the $n-1$ incumbent suppliers who do not win the test auction influence the number of entrants to recruit under the case of information updating. This provides evidence of an additional benefit of using a test auction when the buyer updates her belief about the cost distribution: Beyond refining her knowledge of the *relative* cost differences of the incumbents, the buyer learns about the underlying cost distribution of the suppliers. But as before, the primary benefit of the test auction is reducing the buyer's exposure to relatively high-cost incumbent suppliers.

Despite these interesting differences, we conclude by noting that the comparative statics insights from Proposition 2 continue to hold under the case of information updating.

PROPOSITION 6. *When the buyer updates her beliefs about the suppliers' cost distribution given the drop-out bids from the test auction according to $F(\cdot|r_1; b_1, b_2, \ldots, b_{n-1})$ where r_1 is the reserve price used in the test auction, and only allows the incumbent supplier (if any) who is awarded units in the test auction to compete with entrant suppliers in the second auctions, then Proposition 2 holds with the following change: In part (i)(d), the buyer's ex ante and updated beliefs are given by $\tilde{f}(x) = \frac{1}{\alpha} f\left(\frac{x+\Delta}{\alpha}\right)$ and $\tilde{f}(x|\frac{b+\Delta}{\alpha}; b_1, b_2, \ldots, b_{n-1}) = \frac{1}{\alpha} f\left(\frac{x+\Delta}{\alpha}\Big|\frac{b+\Delta}{\alpha}; \frac{b_1+\Delta}{\alpha}, \frac{b_2+\Delta}{\alpha}, \ldots, \frac{b_{n-1}+\Delta}{\alpha}\right)$ with $\Delta \in \mathbb{R}$ and $\alpha \in \mathbb{R}^+$, respectively.*

5. Conclusions

In this paper we have addressed a fundamental question in sourcing that has received surprisingly little attention in the literature to date: How does a buyer who already has a set of suppliers she has worked with in the past balance the cost of recruiting even more suppliers against the (uncertain) savings that she would accrue by doing so? The key question we address is whether it is better for the buyer to immediately recruit more suppliers and auction off the contract or whether it would be better to first run a test auction to gain a better understanding of the incumbent suppliers' costs for the contract and use that information to assess how many more suppliers to recruit.

It is natural for buyers to want to test the market by contracting out a portion of their total purchase quantity. Therefore, we believe our paper makes a practical contribution to the sourcing literature by providing a framework for characterizing buyers' test auction strategies. For brevity we will not recapitulate our findings, but instead highlight a model extension and discuss the applicability of our paper.

We analyzed how the buyer can incorporate a reserve price into her procurement strategy. Although the auction literature commonly uses the notion of an outside option cost (needed

when setting a reserve price), such a cost is typically assumed to be exogenous, e.g., typically the cost of in-house production. In contrast, we endogenize this cost — namely, the cost of recruiting new suppliers and running a new auction. In practice, many buyers we have dealt with (including the buyer described in the Introduction) need to procure items that they are *not* capable of producing themselves. The buyer relies on outside suppliers for production, a fact which she must carefully consider when setting, say, a reserve price. To our knowledge, this is the first paper to inform this decision by endogenizing the cost of recruiting new suppliers and re-running an auction. That said, one could relax the assumption that the buyer must purchase her demand Q from suppliers. This would change some details of the analysis: for example, letting the finite non-transaction cost per unit be denoted C^{NT}, one can show that the reserve price of the second auction becomes the minimum of the clearing price of the test auction and $\psi^{-1}(C^{NT})$. However, the qualitative insights are unchanged, and for brevity we omit further detail on this extension.

Given that costly new supplier recruitment is a key aspect of global supply chain management, we believe our work is applicable in a wide variety of industrial settings, and can spur future studies on using market test processes (such as test auctions) to help better manage recruitment costs. However, there are situations where our proposed test strategy may not be appropriate. For example, if the firm needs to buy tooling for the supplier, and the tooling is very expensive, the firm may prefer to only strike a contract with a single supplier. Also, if the buyer is auctioning a short-term contract and suppliers' costs and/or capacities change quickly, this may limit the relevance of information gleaned in a test auction, pushing the buyer to the no-test approach. The authors have also seen cases where the response time required by the buyer firm's customer (e.g., an OEM) is very short and the buyer firm's sourcing team does not have time to do even one round of new supplier recruitment. Of course, in this situation, the firm had no choice but to award the contract to one of the existing suppliers. The implication is that although market test processes may be useful in many settings, they will not be appropriate for all settings.

References

Ausubel, L.M., P. Cramton. 2006. Dynamic auctions in procurement. N. Dimitri, G. Piga, G. Spagnolo, eds., *Handbook of Procurement*. Cambridge University Press, 220–246.

Chaturvedi, A., V. Martínez-de-Albéniz. 2011. Optimal procurement design in the presence of supply risk. *Manufacturing & Service Operations Management* **13**(2) 227–243.

32

Beil, Chen, Duenyas, and See: *When to Deploy Test Auctions in Sourcing*
Article submitted to *Manufacturing & Service Operations Management*; manuscript no. MSOM-14-313

Chen, F. 2007. Auctioning supply contracts. *Management Science* **53**(10) 1562–1576.

de Boer, L., G. van Dijkhuizen, J. Telgen. 2000. A basis for modelling the costs of supplier selection: The
economic tender quantity. *Journal of the Operational Research Society* **51**(10) 1128–1135.

Elmaghraby, W.J. 2000. Supply contract competition and sourcing policies. *Manufacturing & Service
Operations Management* **2**(4) 350–371.

Federgruen, A., N. Yang. 2009. Optimal supply diversification under general supply risks. *Operations
Research* **57**(6) 1451–1468.

Hartline, J., D. Hoy, S. Taggart. 2014. Price of anarchy for auction revenue. Working Paper.

Lebrun, B. 1999. First price auctions in the asymmetric n bidder case. *International Economic Review* **40**(1)
125–142.

Mares, V., J.M. Swinkels. 2014. On the analysis of asymmetric first price auctions. *Journal of Economic
Theory* **152** 1–40.

Marshall, R.C., M.J. Meurer, J.F. Richard, W. Stromquist. 1994. Numerical analysis of asymmetric first
price auctions. *Games and Economic Behavior* **7**(2) 193–220.

Maskin, E.S., J.G. Riley. 2000. Equilibrium in sealed high bid auctions. *Review of Economic Studies* **67**(3)
439–454.

McAfee, R.P., J. McMillan. 1988. Search mechanisms. *Journal of Economic Theory* **44**(1) 99–123.

Myerson, R.B. 1981. Optimal auction design. *Mathematics of Operations Research* **6**(1) 58–73.

Peleg, B., H.L. Lee, W.H. Hausman. 2002. Short-term e-procurement strategies versus long-term contracts.
Production and Operations Management **11**(4) 458–479.

Tomlin, B. 2006. On the value of mitigation and contingency strategies for managing supply chain disruption
risks. *Management Science* **52**(5) 639–657.

Wan, Z., D.R. Beil. 2009. RFQ auctions with supplier qualification screening. *Operations Research* **57**(4)
934–949.

Wan, Z., D.R. Beil, E. Katok. 2012. When does it pay to delay supplier qualification? Theory and experi-
ments. *Management Science* **58**(11) 2057–2075.

Yang, Z., G. Aydin, V. Babich, D.R. Beil. 2012. Using a dual-sourcing option in the presence of asymmet-
ric information about supplier reliability: Competition vs. diversification. *Manufacturing & Service
Operations Management* **14**(2) 202–217.