Management Science

Split-Award Auctions: Insights from Theory and Experiments

Aadhaar Chaturvedi, Elena Katok, Damian R. Beil

Split-Award Auctions: Insights from Theory and Experiments

Aadhaar Chaturvedi,[a] Elena Katok,[b] Damian R. Beil[c]

[a] Department of Business Administration, University of Namur, B-5000 Namur, Belgium; [b] Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080; [c] Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109
Contact: aadhaar.chaturvedi@unamur.be (AC); ekatok@utdallas.edu, http://orcid.org/0000-0001-9099-3824 (EK); dbeil@umich.edu (DRB)

Abstract. We investigate procurement in a setting in which the buyer is bound by sourcing rules. Sourcing rules may limit the minimum and maximum amounts of business that can be awarded to a single supplier or dictate the minimum number of suppliers who are awarded business, thus necessitating split awards. The buyer announces the splits before the auction, and suppliers bid accordingly. We consider two auction formats: the sealed-bid first-price auction, and a version of the open-bid descending-price auction. We characterize the suppliers' symmetric equilibrium bidding strategy for both formats and find that the two formats yield the same expected buyer's cost. We characterize the cost of multisourcing, showing among other things that it is always costly for the buyer to split its award among more suppliers if the suppliers' costs are regularly distributed, but that doing so can actually reduce the buyer's expected auction payment if the suppliers' costs are not regularly distributed. The results from controlled laboratory experiments, involving human subjects, indicate that expected cost equivalence fails when costs are regularly distributed because suppliers bid more aggressively in the sealed-bid auction. However, for split-award auctions with nonregularly distributed costs, the sealed-bid prices are actually higher than predicted by theory. We explain these mismatches between observations and theory through a behavioral model based on bidders' aversion to anticipated regret. The experimental results indicate that the theory does a good job of predicting the relationship between the buyer's average cost and the award splits, as well as the cost of multisourcing. Importantly, the experiments confirm that when suppliers' costs come from a nonregular distribution, it may be to the buyer's advantage to diversify the supply base more than is strictly necessitated by sourcing rules.

1. Introduction

Getting a low cost of input is often thought of as the primary goal of a procurement auction. However, in reality, the quest for low input costs must be tempered by other, both short- and long-range considerations of the firm. Rather than setting its procurement managers loose to minimize purchase costs, many firms establish guidelines for procurement managers that ensure that such short- and long-run considerations are addressed. These "sourcing rules" establish things like the maximum order size that can be allocated to any one supplier (for resilience to supply disruptions) and the minimum number of suppliers who must be awarded business (to foster the long-run health of the suppliers and ensure viable sourcing options in the future). But multisourcing also increases the total purchase cost, as well as administrative costs associated with managing multiple suppliers. Therefore, sourcing rules often also specify a minimum amount of business

that can be awarded to any active supplier. Hohner et al. (2003) and Bichler et al. (2006) discuss these and other sourcing rules used in practice for the procurement of materials at Mars, Incorporated.

One straightforward and transparent approach that buyers often use in practice, called a split-award auction, is to announce, prior to the procurement auction, that the best bidder will be allocated the highest share $x\%$ of the contract, the second-best bidder will be allocated the second-highest share $y\%$ of the contract, and so on. (Of course, the allocation that the buyer announces must be consistent with the firm's sourcing rules.) Such split-award auctions show up in various settings, ranging from advertising services to packaging.[1] For example, a buyer one of the authors interacted with used them for buying printed labels for canned food. The buyer decided that three printing suppliers would be awarded volume over a three-year horizon. The suppliers understood that each year there would

be six to 12 label changes (reflecting a holiday season, sporting event sponsorship, etc.), and when a change happened, they would have to compete for volume in a three-way split-award auction. Changeover costs (e.g., spending time fine-tuning a new label's color fidelity for branding) made suppliers highly motivated to win as much volume as possible. To ensure that each supplier covered its changeover costs and made profit, the buyer used a minimum split of around 20%, choosing the first and second splits to try to generate as much pricing competition as possible; it was common for the buyer to use splits such as 50–25–25 or 50–30–20. Moreover, announcing the splits prior to the procurement auction made it relatively easy for the labels buyer to explain the format to the printing suppliers—the lower you bid, the higher volume you stand to win.

Such split-award auctions are the focus of our paper. They are different from standard winner-take-all auctions in terms of both the bidding strategies of the suppliers and the auction design perspective of the buyer—the two issues that we examine in this paper.

We consider two auction design decisions that buyers make for split-award auctions: the bidding format (sealed-bid versus open-bid) and the allocation rules. Sealed-bid and open-bid are the canonical formats used in a variety of industries (Beall et al. 2003; Jap 2003, 2007). The sealed-bid format we analyze is a generalization of the sealed-bid winner-take-all auction; each supplier participating in the auction bids a single price per unit, and the lowest bid receives the largest share, the second-lowest bid receives the second-largest share, etc. Thus, it represents the one-shot (each supplier simultaneously submits their best and final bid without knowing what the other suppliers have bid) auction format that is the classical way of awarding contracts. The open-bid format that we analyze for a split-award setting is a generalization of the simple open-descending format for the winner-take-all auction (see Chen et al. 2015). The open-bid auction represents the dynamic format in which suppliers can respond to competing bidders by staying in the auction at successively lower bid prices.

For procurement managers that organize split-award auctions, *getting the splits right is critical to balancing the strategic need for splitting the award and keeping the costs under control*. Given sourcing rules about the minimum number of suppliers and maximum percentage share awarded to the lowest bidders, the procurement manager still has latitude in determining the number of winning suppliers and the share percentage awarded to each. It is not obvious up front how a procurement manager can decide on the most cost-efficient split among the many different possible splits. For instance, would it be most cost efficient to award the maximal allowed percentages to the lowest bids despite the fact that doing so lessens competition between low-cost bidders? Or would it be better to increase competition among low-cost bidders by widening the gaps between the percentages awarded to them, even though doing so means shifting some allocation percentages to higher-cost bidders?

The theoretical contribution of this paper is to characterize the buyer's expected cost, from the bidding equilibrium, for both auction formats, and then analyze the problem that procurement managers face in choosing the allocation that minimizes the expected procurement cost while satisfying the sourcing rules. We define a *greedy* allocation as one that satisfies the sourcing rules and for which a positive quantity cannot be transferred from the allocation of a higher-bidding supplier to the allocation of a lower-bidding supplier without violating the sourcing rules. Intuitively, in a greedy allocation, the buyer allocates the maximum possible business sequentially, starting from the lowest-bidding supplier and moving toward higher-bidding suppliers, such that the sourcing rules are satisfied. A managerially insightful and theoretically novel result that we derive in this paper is that, as along as the underlying cost distribution of suppliers is well behaved (i.e., regular), it is most cost efficient for the buyer to award splits greedily.[2] Many well-known distributions are regular—e.g., uniform, normal, and exponential distributions (Bagnoli and Bergstrom 2005).

However, other common cost distributions that would arise naturally in practice, such as multimodal distributions, need not be regular.[3] As an easy example, one can think of a two-type cost distribution. This arises in practice when suppliers, qualitatively speaking, have costs that are either on the low end or on the high end. An automotive firm for which one of the authors helped source parts buys screw machine parts. These complex parts, milled out of bar stock, can be produced on single-spindle machines or a multispindle machine where parallel operations can be carried out on multiple parts as the spindles rotate to each machining station. If the supplier has a multispindle machine available that is well suited to the part being sourced, then its per-unit cost would tend to be lower compared to the case where the supplier would intend to use a single-spindle machine. The buyer might not know a priori which type of machine any one supplier intends to use to produce the part (the use of single or multispindle machine might depend on spare capacity available on those machines), but her prior would be that the distribution of supplier costs has weight built up around higher- and lower-cost types, and this would be reflected in the cost distribution function having two "modes." Surprisingly, for nonregular distributions, we show that sometimes the optimal allocation is not greedy—for instance, the buyer might want to

allocate its business to more than the minimum number of suppliers required by the sourcing rules. In other words, sometimes the buyer can minimize its payment costs by splitting its contract among more suppliers. This differs from the otherwise obvious-sounding intuition that splitting the award achieves diversification but comes at the cost of higher payment.

We would like to use our theoretical insights to improve procurement auction design; to this end, we test several of our theoretical predictions using a controlled laboratory experiment with human subjects incentivized with money. The purpose of this test is twofold: first, we wish to identify any systematic behavioral deviations that exist in this setting; second, we aim to extend the theoretical model to account for systematic behavior we observe. We find that bidding behavior in open-bid auctions is quite close to theoretical predictions. For sealed-bid auctions, previous experimental work has repeatedly shown systematic behavioral deviations with regularly distributed costs (see Kagel and Roth 1995), with explanations such as regret—particularly loser's regret—organizing the data (Engelbrecht-Wiggans and Katok 2008). However, in our context, when we consider nonregular distributions and split awards, we have reason to expect that winner's regret can become more salient, which might balance out loser's regret and thus lead to less aggressive bidding. We find that in sealed-bid split-award auctions, participants bid more aggressively than they should under the risk-neutral Nash equilibrium for regular cost distributions, but by contrast, when the cost distribution is nonregular, the participants in the sealed-bid split-award auctions bid more conservatively than the theoretical predictions of the risk-neutral Nash equilibrium. We develop a new model based on aversion to anticipated regret in the split-award setting that successfully organizes our data.

Consistent with our base model, when the cost distribution is well behaved, the buyer's average cost monotonically increases as the allocation becomes less weighted toward lower-cost bids. We also find, consistent with the theory, that the buyer's average cost could be lower with nongreedy split as compared to greedy split when the cost distribution is not well behaved. Thus, our experimental results validate the managerially significant insights derived from the model—i.e., greedy allocations are more cost efficient for the buyer when underlying cost distribution is regular; however, a nongreedy split could be optimal for nonregular cost distribution. We also find that the model does a good job of predicting the cost of multisourcing for the buyer. Because sourcing rules have to balance the costs and benefits of multisourcing, having an analytical model that is able to accurately predict the cost of multisourcing can help buyers design effective sourcing rules.

1.1. Literature Review

Existing literature has investigated the use of multisourcing in auctions to reduce the total procurement cost for the buyer due to many factors. Klotz and Chatterjee (1995) have analyzed split awards when bidders face entry cost and are risk averse. Unlike them, we analyze arbitrary allocation rules, not just a two-way split, and we show that nongreedy splits can be optimal even without appealing to entry costs and risk aversion. Dasgupta and Spulber (1990) investigate split awards when suppliers face convex production cost. However, in our setting, suppliers face constant marginal production costs; our buyer uses split awards to satisfy sourcing rules, subject to which she chooses the cost-minimizing allocation, and we show that more evenly spread splits can be used as a tool to reduce purchase costs in our setting, without appealing to convexity in supplier production costs. Thus, we contribute to the extant split-award literature by showing that more evenly spread splits can reduce the buyer's sourcing costs, without invoking entry costs, risk aversion, or cost convexity.

Tunca and Wu (2009) focus on bounding the optimality loss imparted by using a two-stage procurement process versus a single-stage optimal mechanism (that takes into account supplier production cost convexity). We, on the other hand, focus on a single-stage auction event where the award splits are chosen by the buyer subject to sourcing rules, and characterize the effect of multisourcing on the buyer's expected cost. Chaturvedi and Martínez-de-Albéniz (2011) and Chaturvedi et al. (2014) find the optimal mechanism that multisources to address concerns of supply risk and supply base maintenance, respectively. Here, we abstract away from specific factors (like risk or supply base maintenance) for multisourcing by taking them into account through the buyer's sourcing rules. Subject to these sourcing rules, we then analyze the optimal allocation splits for sealed-bid and open-bid auction formats.

Also related to our work is the literature that has investigated the bidding equilibrium in split-award auctions when the splits are exogenously specified. Anton and Yao (1992) show that auctioning a two-way split (which is decided after auction) among just two bidders can result in coordinated bids when the two participating suppliers can submit multiple bids. Unlike their paper, we consider n bidders facing splits announced up front as part of the auction format, and we then go on to show how the buyer can determine these splits given the allocation rules it faces.

Our paper also relates to Bichler et al. (2014), but unlike them, we analyze an arbitrary split, not just two-way splits. Moreover, as explained above, we study the buyer's problem in terms of how to design allocation rules given sourcing constraints and test the

model that helps estimate the cost of multisourcing, whereas Bichler et al. (2014) focuses on comparing bidding behavior and the resulting performance under two auction formats. Finally, from their experimental study, Bichler et al. (2014) conjecture risk aversion as a major driver for overly aggressive bidding. By contrast, we find that bidders might also overbid (above the Nash-equilibrium predictions) when costs are not regularly distributed, which is inconsistent with the risk-aversion explanation.

Testing auction theory using controlled laboratory experiments has a long tradition (see Kagel and Roth 1995 for an overview of early work and Kagel and Levin 2015 for an overview of more recent work in economics). Much of the early work focused on testing revenue equivalence among the four basic forward auction formats (the sealed-bid first-price, Dutch, English, and the sealed-bid second-price). The findings are that the bidding in the sealed-bid first-price auction is more aggressive than the risk-neutral Nash equilibrium, and the bidding in the Dutch auction is not independent of the speed of the Dutch clock (see Katok and Kwasnica 2008). Thus, generally revenue equivalence (in our case cost equivalence) fails in the laboratory. Elmaghraby et al. (2012) report a similar finding in open-bid reverse auctions with rank feedback, and Haruvy and Katok (2013) observe the same thing in buyer-determined reverse auctions. To the best of our knowledge, all previous experimental studies of auctions used a regular cost distribution (usually Uniform). Our paper is the first to consider a nonregular cost distribution and to investigate the effect of multisourcing with more than two splits in the laboratory, while comparing sealed- and open-bid formats in that setting.

2. Model

We model a buyer that faces a unit (normalized) one-time demand for a standardized homogeneous and divisible product. It can buy this quantity from the $n \geq 2$ qualified suppliers in its supply base. To discover the best available price for the required product, the buyer invites the suppliers to competitively bid for its business. As is common in the literature (e.g., Chen 2007), for each supplier i, the cost to produce q units is given by $q \cdot c_i$, where c_i is the supplier's privately known (known only to the supplier) per-unit cost, which remains constant for producing the quantity q. This captures situations where variable costs are the dominant cost drivers. This arises in a variety of settings—e.g., plastic injection molding (where buyers typically purchase the tooling, so machine time, resin, and electricity are the primary cost drivers at the supplier), or labor-intensive work like simple assembly. With this setup, we will show that greater split awards can actually help the buyer

reduce costs, without appealing to notions of cost convexity that clearly favor the use of split awards (e.g., Dasgupta and Spulber 1990). We let $\mathbf{c} = (c_1, c_2, \ldots, c_n)$ denote the vector of the suppliers' per-unit cost of production. We assume that the costs are identically, independently, and continuously distributed in the interval $[\underline{c}, \bar{c}]$, according to cumulative distribution F (with density f and $\bar{F} = 1 - F$). Finally, we assume that F is common knowledge, the suppliers are risk-neutral profit maximizers, and that the buyer is risk neutral and seeks to minimize its expected cost subject to sourcing rules.

The buyer faces *sourcing rules* that address multiple operational concerns in procurement, such as supply risk, maintaining the supply base, or controlling the administrative cost of purchasing from multiple suppliers. These sourcing rules can be formally characterized as follows:

1. No one supplier can win more than fraction $0 \leq A \leq 1$ of the business, to avoid too much dependence on any one supplier.

2. There must be a minimum number of suppliers, $M \leq n$, that are awarded business. Recognizing that suppliers who do not win any business may disengage from the supply base, buyers wishing to maintain competition for future bidding events may require that at least a handful of suppliers win business in any given bidding event, to keep the suppliers from abandoning the supply base in search of greener pastures.

3. Any supplier awarded business should win at least $0 < B \leq A$ of the business, to avoid administrative inefficiencies of working with very small contracts.

To leverage supplier bid competition and also ensure that its sourcing rules are followed, the buyer organizes a split-award auction. As is common in practice, to keep the auction procedure transparent and straightforward the buyer announces, before the auction, the percentage of its business that it would allocate to the suppliers as a function of the rank of their bid. We let Q_1 denote the highest fraction of the total business that the buyer would award to the lowest bid, Q_2 denotes the second-highest fraction of business that would be awarded to the second-lowest bid, and so on up to Q_n. We denote by $\mathbf{Q} = (Q_1, Q_2, \ldots, Q_n)$, the vector of these allocations such that $Q_i \geq Q_j$ for any $i < j$ and $\sum_{i=1}^{n} Q_i = 1$. Thus, if the buyer announces that it will award 70% of its business to the lowest bidder and 30% to the second-lowest, then $Q_1 = 0.7$, $Q_2 = 0.3$, and $Q_3 = \cdots = Q_n = 0$.

The buyer decides to announce the vector of allocations $\mathbf{Q}$ that minimizes its expected purchase cost such that $\mathbf{Q}$ is consistent with the sourcing rules that the buyer faces. We let $C_{\text{buyer}}(\mathbf{Q})$ denote the expected purchase cost of the buyer if it announces the vector of

allocations $\mathbf{Q}$.[4] Overall, the buyer's problem of deciding the optimal $\mathbf{Q}$, given the sourcing rules, can be characterized as follows:

$$\min_{Q} \ C_{\text{buyer}}(\mathbf{Q}) \tag{1a}$$

$$\text{s.t. } Q_i \geq Q_j, \quad \forall\, i < j,$$

$$\sum_{i=1}^{n} Q_i = 1,$$

$$Q_i \in \{0, [B, A]\}, \quad \forall\, i,$$

$$\sum_{i=1}^{n} z_i \geq M, \quad \text{where } z_i = 1_{Q_i>0}, \forall\, i. \tag{1b}$$

Objective function (1a) characterizes the buyer's objective, whereas constraints (1b) characterize the sourcing constraints that $\mathbf{Q}$ must satisfy. The sourcing constraints (1b) are *consistent*—i.e., give a nonempty set of feasible $\mathbf{Q}$s if the following assumptions are satisfied:

$$n \cdot A \geq 1,$$
$$B \cdot \max(M, \lceil 1/A \rceil) \leq 1. \tag{2}$$

Note that $\lceil 1/A \rceil$ gives the nearest integer greater than $1/A$. The first condition ensures that the constraint on the maximum fraction of business A that can be given to a supplier does not restrain the buyer from satisfying its unit demand. The second condition ensures that the constraint on the minimum amount of business given to a supplier (that gets a nonzero allocation) does not result in the buyer having to procure more than its demand. As an example, values of $A = 40\%$ and $B = 35\%$ would imply that three suppliers would each get at least 35% of the buyer's business, which is not consistent. In this paper, we assume that the conditions (2) are always satisfied.

We will analyze sealed- and open-bid auctions. In Section 3, we investigate the bidding equilibrium in both auction formats for a given vector of allocations $\mathbf{Q}$. We assume that the buyer must transact for her full quantity, so her reservation price is set to the upper bound of the cost distribution, $\bar{c}$. For both auction formats, Section 4 characterizes the buyer's problem of deciding the optimal splits, $\mathbf{Q}$, that minimizes the buyer's expected cost, C_{buyer}, subject to sourcing constraints (1b).

3. Analysis of Sealed and Open Split-Award Auctions

In this section, we analyze the suppliers' bidding equilibrium in both the sealed- and open-bid auction formats for a given allocation structure $\mathbf{Q}$, and formulate the corresponding expected cost for the buyer. We first analyze the sealed-bid auction and then the open-bid auction.

3.1. Sealed-Bid Auction

The sealed-bid auction is implemented in the following way. The buyer announces the allocation vector, $\mathbf{Q} = Q_1, \ldots, Q_n$, where Q_i represents the allocation to the ith-ranked bid. If two or more suppliers bid the same, then each of the tied suppliers is awarded the average of the allocations associated with the tied ranks. The buyer also announces that the price paid per unit of allocation would be the bid quoted by the supplier. Each supplier then submits its bid to the buyer. After collecting all of the bids, the buyer makes the allocations and payments. The expected payoff function for supplier i when it bids b_i can be expressed as

$$\Pi_i(b_i, c_i, \mathbf{Q}) = (b_i - c_i)H(b_i, \mathbf{b}_{-i}, \mathbf{Q}),$$

where $b_i - c_i$ is the profit margin, $H(b_i, \mathbf{b}_{-i}, \mathbf{Q})$ represents the expected allocation to supplier i, and $\mathbf{b}_{-i}$ represents the vector of bids submitted by all of the other suppliers. To find the equilibrium bidding strategy, we adapt the canonical, winner-take-all sealed-bid auction equilibrium analysis (e.g., Krishna (2010)). Assume that the bidding strategy of all suppliers, except for supplier i, is $b_j = \beta(c_j, \mathbf{Q})$ defined on the domain $[\underline{c}, \bar{c}]$. For now, we assume that $\beta(c, \mathbf{Q})$ is a continuously differentiable and increasing function of c (we later verify that these assumptions are indeed true in equilibrium). The assumption on continuous cost distribution together with the assumptions on $\beta(\cdot, \cdot)$ implies that $H(b_i, \mathbf{b}_{-i}, \mathbf{Q})$ can be expressed as $H(\beta^{-1}(b_i, \mathbf{Q}), \mathbf{Q})$, where $\beta^{-1}(\beta(c, \mathbf{Q}), \mathbf{Q}) = c$.

Define

$$H(x, \mathbf{Q}) \equiv \frac{Q_1 \bar{F}(x)^{n-1} + Q_2 \binom{n-1}{1} \bar{F}(x)^{n-2}F(x) + \cdots}{+ Q_{n-1}\binom{n-1}{n-2}\bar{F}(x)F^{n-2}(x) + Q_n F^{n-1}(x)}. \tag{3}$$

Substituting $x = \beta^{-1}(b_i, \mathbf{Q})$ in Equation (3) gives the expected allocation of supplier i. Hence, supplier i's expected payoff can be characterized as

$$\Pi_i(b_i, c_i, \mathbf{Q}) = (b_i - c_i)H(\beta^{-1}(b_i, \mathbf{Q}), \mathbf{Q}). \tag{4}$$

Thus, one can differentiate Equation (4) with respect to b_i to characterize the best response of supplier i given that all other suppliers use a symmetric bidding strategy, β. Then, by assuming that supplier i's best response is also β, one can characterize the symmetric equilibrium strategy β. If the strategy β satisfies all of the assumptions (i.e., it is continuously differentiable, increasing, and the best response of supplier i given all other suppliers' best response is β), then it does indeed formulate the symmetric bidding equilibrium strategy. The following proposition (all results are formally proved in the appendix) characterizes the equilibrium.

Proposition 1. *For any allocation $Q_1 \geq Q_2 \geq \cdots \geq Q_n$, the symmetric equilibrium bid function in a sealed first price auction is given by*

$$\beta(c, \mathbf{Q}) = \bar{c} \cdot \frac{H(\bar{c}, \mathbf{Q})}{H(c, \mathbf{Q})} - \frac{1}{H(c, \mathbf{Q})} \int_{x=c}^{\bar{c}} x \, dH(x, \mathbf{Q}). \quad (5)$$

Note that for a winner-take-all auction, the allocation vector would be $Q_1 = 1$ and $Q_2 = \cdots = Q_n = 0$; accordingly, the equilibrium bid function (5) gives $\beta(c) = (1/\bar{F}(c)^{n-1}) \int_{x=c}^{\bar{c}} x(n-1)\bar{F}(x)^{n-2} f(x) \, dx$, which is indeed the equilibrium bid function for a sealed-bid winner-take-all auction—i.e., each supplier, conditional on it being the lowest-cost supplier, bids the expected cost of the lowest-cost supplier among the other $n-1$ suppliers (see Krishna 2010).

Using the equilibrium bids, we can characterize the buyer's expected cost in the sealed-bid auction. Even though the equilibrium bid function of the suppliers appears complicated, it turns out that the buyer's expected cost can be simplified to a rather clean expression. The following proposition does exactly that. We define $\mu_m \equiv E_c[C_{m:n}]$, the expected mth order statistic from n draws. For notational convenience, for any $m > n$, we take $\mu_m = \bar{c}$.

Proposition 2. *The buyer's expected cost in the sealed-bid auction is*

$$\begin{aligned}
C_{\text{buyer}}(\mathbf{Q}) = {} & \mu_2 Q_1 + (2\mu_3 - \mu_2)Q_2 + \cdots \\
& + (m\mu_{m+1} - (m-1)\mu_m)Q_m + \cdots \\
& + ((n-1)\mu_n - (n-2)\mu_{n-1})Q_{n-1} \\
& + (n\bar{c} - (n-1)\mu_n)Q_n. \qquad (6)
\end{aligned}$$

Thus, the buyer's problem for deciding the optimal $\mathbf{Q}$ for the sealed-bid split-award auction can be characterized as: $\min_{\mathbf{Q}} C_{\text{buyer}}(\mathbf{Q})$ such that $\mathbf{Q}$ satisfies the constraints (1b). In Section 4, we find the optimal allocation vector $\mathbf{Q}$ for the sealed-bid auction by solving this problem.

But before that, we first rearrange the terms of the buyer's expected cost in Equation (6) as follows:

$$\begin{aligned}
C_{\text{buyer}}(\mathbf{Q}) = {} & nQ_n\bar{c} + (n-1)(Q_{n-1} - Q_n)\mu_n + \cdots \\
& + m(Q_m - Q_{m+1})\mu_{m+1} + \cdots + (Q_1 - Q_2)\mu_2. \\
& \qquad (7)
\end{aligned}$$

Equation (7) implies that if the buyer gave Q_n amount of business to each of the n bidders at price $\bar{c}$ and then gave $Q_{n-1} - Q_n$ amount of business to $n-1$ bidders at a price equal to the cost of the highest bidder, and so on, then the buyer's expected cost would match Equation (6). Below, we describe how this outcome can be implemented through an open descending-price-clock auction.

3.2. Open Descending Auction

For any allocation structure $\mathbf{Q}$, we implement the open auction as a descending-price-clock auction. In this auction format, the price clock starts at price $\bar{c}$. At the start of the price clock, the buyer allocates Q_n amount of business to each of the n suppliers, and for this quantity pays them a per-unit price of $\bar{c}$. The price clock then begins to move down. The suppliers can drop out of the auction at any time. Dropping out of the auction does not give any additional allocation or payment to the supplier who drops out, beyond what it has already received. However, a supplier that drops out does result in each of the suppliers who remain in the auction getting allocated some nonnegative quantity for a per-unit payment equal to the price at which the supplier dropped out of the auction. Specifically, the first supplier to drop out results in each of the remaining $n-1$ suppliers getting an allocation of $Q_{n-1} - Q_n$ at a per-unit payment equal to the price at which that first supplier dropped out. More generally, the dropping out of the $(n-m+1)$th supplier (for $2 \leq m \leq n$) results in the remaining $m-1$ suppliers getting allocated $Q_{m-1} - Q_m$ amount of business at the auction price at which the $(n-m+1)$th supplier dropped out. The auction stops when the second-to-last (the $(n-1)$th) supplier drops out of the auction. Note that the overall allocation awarded to the $(n-m+1)$th supplier to drop out is Q_m, and the last remaining supplier gets an overall allocation of Q_1.

In such an auction, for any allocation $\mathbf{Q}$ such that $Q_1 \geq Q_2 \geq \cdots \geq Q_n$, a supplier i finds it optimal to drop out when the price clock reaches its marginal cost c_i. If supplier i drops out any sooner (at $b > c_i$) then it only loses the opportunity to get an allocation at a profitable price had some other supplier dropped out between b and c_i. If supplier i drops out later (at $b < c_i$), then it only increases the likelihood of getting an allocation at a loss-making price that happens if some other supplier drops out at a price between c_i and b. Thus, in an open descending auction, the equilibrium strategy for all suppliers (except the lowest-cost, since the auction stops when the second-to-last supplier drops out) is to drop out when the price clock reaches their respective per-unit cost.[5]

Since bidders find it optimal to drop out at their true costs, it is easy to see that the above auction results in the same expected cost as the sealed-bid format from the previous section. This follows naturally given the well-known revenue equivalence theorem; indeed, Wambach (2002) formally extends the notion of revenue equivalence (in a forward auction context) to split-award auctions that award the largest share to the bidder with the highest bid, the second-largest share to the bidder with the second-highest bid, and so on. For us, the implication of cost equivalence is that the buyer's problem of deciding the optimal splits $\mathbf{Q}$ for

the open-descending split-award auction also remains the same as for the sealed-bid auction.

4. Optimal Splits

In this section, we solve for the buyer's problem of deciding the optimal splits, $\mathbf{Q}$, for both the sealed-bid and open descending auction formats. Namely, we embed the buyer's expected cost, $C_{\text{buyer}}(\mathbf{Q})$, as characterized by Equation (6), into the objective (1a). The decision variables $Q_1, \ldots, Q_n$ are real numbers not less than 0, and $C_{\text{buyer}}(\mathbf{Q})$ in Equation (6) is linear in $Q_1, \ldots, Q_n$; hence, the math program (1) formulates a constrained fractional knapsack problem and thus would give corner solutions. However, it is not obvious whether the solution is also greedy (as defined in the introduction). As an example, for parameter values $B = 10\%$, $M = 3$, and $A = 50\%$, a greedy allocation would imply allocating 50%, 40%, and 10% of the business to the lowest-, second-lowest-, and third-lowest-bidding suppliers, respectively, and giving a zero allocation to all of the remaining suppliers. In the following lemma we show that allocating greedily is similar to solving an optimization problem.

Lemma 1. *Maximizing $\sum_{i=1}^n Q_i^2$ such that the sourcing constraints (1b) are satisfied would give a unique solution—namely, the greedy allocation.*

Note that $\sum_{i=1}^n Q_i^2$ is the same as the Herfindahl–Hirschman index (HHI) used to describe market concentration. Thus, allocating greedily is the same as maximizing the HHI of the allocations such that sourcing constraints (1b) are satisfied.

Indeed, a greedy allocation is always (i.e., for all sourcing constraints) optimal if and only if the coefficients of the objective function given in (6) are increasing—i.e., the coefficient of Q_1 is less than that of Q_2, and so on. The following lemma characterizes the necessary and sufficient conditions for the greedy allocation to be optimal.

Lemma 2. *Allocating greedily to the lowest-bidding supplier is always (i.e., for all sourcing constraints) optimal. Allocating greedily to all of the other suppliers is always (i.e., for all sourcing constraints) optimal if and only if*

$$m\mu_{m+1} - (m-1)\mu_m \geq (m-1)\mu_m - (m-2)\mu_{m-1},$$
$$\forall\, 2 < m < n, \quad (8a)$$

$$and \quad n\bar{c} - (n-1)\mu_n \geq (n-1)\mu_n - (n-2)\mu_{n-1}. \quad (8b)$$

Lemma 2 implies that the buyer would never find it optimal, as an example, to go from an 80–20 split to a 70–30 split, provided that the sourcing constraints are met in both cases. Similarly, the buyer would never benefit from splitting a sole award into more splits, provided that sole sourcing does not violate the sourcing constraints. However, what about splitting the

award fraction among more suppliers who do not bid the lowest, when the sourcing constraints require the buyer to procure from at least two suppliers? In the presence of such sourcing constraints, the optimality of the greedy allocation would depend on whether the conditions (8) hold or not.

Let us build some intuition into why conditions (8) might not hold (we will use an example that we will return to later in the experiments). Suppose the buyer has four bidders and is comparing $Q_1 = 1/3$, $Q_2 = 1/3$, $Q_3 = 1/3$, $Q_4 = 0$ split versus $Q_1 = 1/3$, $Q_2 = 1/3$, $Q_3 = 1/6$, $Q_4 = 1/6$ split. Compared to the former greedy allocation, the latter nongreedy allocation has advantages: it encourages competition among the lowest-cost suppliers. However, it has disadvantages in that it sacrifices competition among the higher-cost suppliers—now, even the worst-cost supplier receives some allocation. Interestingly, it turns out that the advantage of nongreedy can outweigh the disadvantage. To see how this might happen, consider the following bi-modal probability density function of suppliers' cost (in the top part of Figure 1, we show the probability density function (p.d.f.) characterized below):

$$f(x) = \begin{cases} 500x & \text{if } 0 \leq x \leq 0.01, \\ 5 - 24.95(x - 0.01) & \text{if } 0.01 \leq x \leq 0.21, \\ 0.01 & \text{if } 0.21 \leq x \leq 0.97, \\ 0.01 + 1035.78(x - 0.97) & \text{if } 0.97 \leq x \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$
$$(9)$$

The distribution characterized in Equation (9) represents a continuous version of the familiar two-type distribution described in the introduction: "low-cost" suppliers and "high-cost" suppliers, where the buyer does

Figure 1. (Color online) (Top) The p.d.f. $f(x)$ Characterized by Equation (9); (Bottom) the Equilibrium Bid for Greedy and Nongreedy Allocations When Cost Are Distributed According to Equation (9)
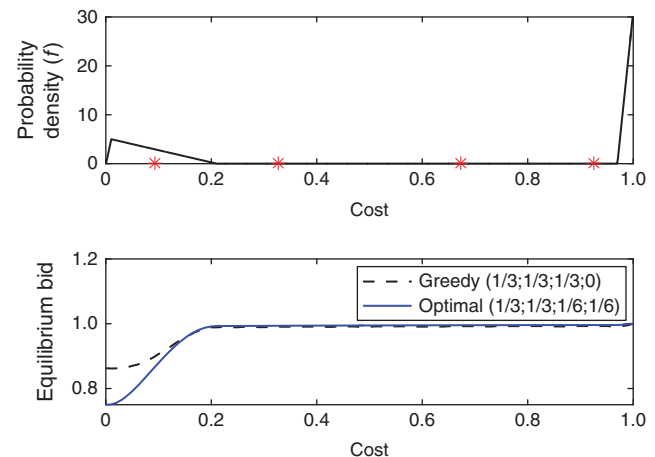
Note. The first four expected order statistics are marked with an asterisk.

not know each supplier's type. Below, we explain why, for such a multimodal distribution, allocating greedily might not be optimal.

For the distribution characterized by Equation (9) and $n = 4$ suppliers and sourcing constraint $A = 1/3$, the optimal allocation according to math program (1) would be nongreedy—i.e., $Q_1 = Q_2 = 1/3$ and $Q_3 = Q_4 = 1/6$. On the other hand, the greedy allocation would be $Q_1 = Q_2 = Q_3 = 1/3$ and $Q_4 = 0$. The intuition is that with the nongreedy allocation, competition at the low end of the cost distribution is more fierce, and this is not offset by less competition at the high end of the cost distribution. The bottom part of Figure 1 depicts the equilibrium bids for both optimal and greedy splits when the cost distribution is bimodal (according to Equation (9)). One can see that for the nongreedy allocation, the low-cost bidders bid significantly lower than they would under the greedy allocation. The reason is that under the greedy allocation the low-cost bidders simply want to avoid coming in last, but they have no incentive to be the first- or second-best bidder like they would with the nongreedy allocation. The resulting difference in the buyer's expected cost between the greedy allocation and the optimal allocation would be 3.81%. Moreover, we see that in presence of sourcing constraints, it might be optimal for the buyer to source from four suppliers instead of the minimum required three suppliers—i.e., the buyer can reduce its expected purchasing cost by diversifying more. Indeed, other distributions can be thought of for which a nongreedy allocation is optimal. In Online Appendix A, we provide another such example in which the buyer can reduce its expected cost by splitting its business among four suppliers, even though the sourcing constraints require awarding business to just two suppliers (e.g., when $A = 55\%$, $B = 0$, and $M = 0$).

The above discussion provided insight on why nongreedy can outperform greedy. The greedy allocation reduces competition among lower-cost bidders, encouraging lower-cost bidders to inflate their bids. Intuitively, this is particularly problematic for the buyer in cases where there is a sizable cost gap between lower- and higher-cost bidders (as can happen with a multimodal distribution), which is why nongreedy allocations can be optimal in such cases. What matters is how competition heats up among lower-cost bidders when moving to a nongreedy allocation, and whether or not this offsets allocating more units to higher-cost bidders. The importance of Lemma 2 is underscored by this discussion, precisely because it helps us understand the conditions for when greedy will still be optimal. What matters is that the gaps between the order statistics of the underlying cost distribution are well behaved. It turns out that a familiar condition is all we need to guarantee that the gaps between the order statistics are well behaved; thus, allocating greedily is

optimal. In particular, all we need is that the underlying distribution is regular. The next theorem formalizes this result.

Theorem 1. *Shifting a positive amount from a low bidder's allocation to a high bidder's allocation (i.e., shifting an $\epsilon > 0$ from Q_i to Q_j for any $i < j$), with everything else being the same, decreases the Herfindahl–Hirschman index (HHI) of the allocations, and for any regular distribution (a continuous distribution for which $c + F(c)/f(c)$ is increasing) will increase the buyer's expected cost.*

Distributions that have a log-concave density satisfy the regularity condition, including uniform, exponential, normal, and power-function distributions (see Bagnoli and Bergstrom 2005).

By definition, an allocation is greedy if a positive quantity cannot be transferred from Q_j to Q_i for any $j > i$ without violating the sourcing constraints (1b). The next corollary follows from Theorem 1.

Corollary 1. *For any regular cost distribution, the buyer finds it optimal to announce a greedy allocation—i.e., maximize the Herfindahl–Hirschman index (HHI) of the allocations such that the sourcing constraints (1b) are satisfied.*

In Figure 2, we evaluate the buyer's expected cost with four participating suppliers when it allocates 1/3 of its business to the two lowest-cost suppliers and progressively changes the allocation to the third-lowest- and the fourth-lowest-cost suppliers from 1/3 to 1/6 and 0 to 1/6, respectively. Thus, we evaluate the buyer's expected cost as it diversifies more (and hence decreases the HHI of allocations). Consistent with Theorem 1, we find that the buyer's expected

Figure 2. (Color online) Expected Buyer Cost as Allocation Is Transferred from Lower-Cost Supplier to Higher-Cost Supplier

Notes. $n = 4$, $Q_1 = 1/3$, $Q_2 = 1/3$, and Q_3 and Q_4 are changed from 1/3 to 1/6 and 0 to 1/6, respectively. Costs are uniformly distributed in unit interval for the regular distribution and are distributed according to Equation (9) for the nonregular distribution.

cost increases as it diversifies more when the underlying costs are uniformly distributed. However, we see that the buyer's expected cost decreases as it diversifies more when the underlying costs follow a nonregular distribution (characterized by Equation (9)). Thus, we see that a buyer might decrease its expected purchasing cost by diversifying more than what is strictly required by its sourcing rules when the underlying cost distribution is not regularly distributed.

Moreover, note that Theorem 1 only provides a sufficient condition for optimality of the greedy allocation. Thus, cost distributions that do not satisfy the regularity condition of Theorem 1 can still result in optimality of the greedy allocation. As an example, consider an arc-sine cost distribution that has a density function defined as $f(c) = 1/\pi\sqrt{c(1-c)}$ and cumulative distribution function (c.d.f.) $F(c) = (2/\pi)\sin^{-1}(\sqrt{c})$ in the interval $[0,1]$. It can be easily established that this distribution does not satisfy the regularity condition. However, for $n = 3$ suppliers, this distribution does satisfy conditions (8), thus resulting in the greedy allocation being optimal.[6]

The results of Theorem 1 and Lemma 1 also allow us to investigate the sensitivity of the buyer's expected cost as the sourcing rules (parameters A, B, and M defined in Section 2) are relaxed. For regular cost distributions, we find that the buyer's expected cost is convex decreasing in A and is convex increasing in B and M. We provide a more formal statement of this result and the related proof in Online Appendix B. Intuitively, this result implies that marginal increases in diversification at low levels of diversification will cost the buyer less as compared to marginal increases in diversification at higher levels of diversification. The online companion provides exact expressions for the rate of change in the buyer's expected cost (C^*_{buyer}) as the sourcing rules (A, B, and M) are changed (in Equations (A3)–(A5) of Online Appendix B). Thus, a buyer could compute how much additional diversification would cost as it changes the sourcing rules.

5. Research Hypotheses, Experimental Design, and Results

This section reports on laboratory experiments designed to test theoretical predictions developed about split-award auctions in the preceding sections. For this, we conduct split-award auctions in a controlled laboratory environment and compare the results obtained in these experiments to those predicted by the theory. We begin by formulating specific research hypotheses. Then, we describe the experiment we designed to test these hypotheses. Lastly, we report experimental results.

5.1. Research Hypotheses

Our model makes predictions about bidding behavior and the buyer's resulting cost under split awards, with sealed-bid and open-bid auction formats, and regular and nonregular underlying cost distributions. It also makes predictions about how the optimality of greedy allocation is affected by the distribution of bidders' costs. Therefore, we designed our study to test all of these predictions.

The first set of hypotheses test the predictions in Propositions 1 and 2, regarding the bidding behavior in the sealed-bid auction with split awards, the average cost of the buyer that results from this behavior, and the resulting auction efficiency. The hypothesis pertains to settings with regular and nonregular cost distributions; however, our tests of the hypothesis focus on the former (uniform distribution).

Hypothesis 1A (H1A). *Bidding behavior will follow, on average, equilibrium predictions of Proposition 1. The average cost to the buyer will not be significantly different from the buyer's cost prediction that follows from Proposition 2; all auctions will be 100% efficient.*

The first part of the hypothesis, regarding bidding behavior and the cost of the buyer, is a strong test of Proposition 1 because it requires the bidding behavior to match, on average, the risk-neutral Nash equilibrium (RNNE) (see Online Appendix C for the sealed-bid first-price bidding equilibrium function for uniformly distributed costs). The second part of the hypothesis tests the efficiency of allocation. Efficiency can be measured in different ways, but the one most commonly used is *allocational efficiency*, the proportion of efficient allocations. To measure allocational efficiency, we code an allocation as efficient whenever no bidder with lower cost is allocated a market share that is smaller than the market share allocated to any bidder with a higher cost. It is possible for the bidding behavior to be different from Proposition 1, rejecting the first part of H1A, while the auction still remains efficient.

Equilibrium predictions about bidding behavior and resulting buyer costs are based on the assumption that bidders are fully rational and risk-neutral, and more importantly, are not affected by any behavioral biases. These assumptions have been tested and rejected in the prior literature for first-price sealed-bid auctions under single sourcing but not for auctions with split awards. The survey of Kagel and Roth (1995) summarizes experimental economics literature that reports overly aggressive bidding in sealed-bid first-price (single-source) auctions, as well as bidding that closely follows the dominant strategy in open-bid (single-source) auctions. The alternative hypothesis below is based on our knowledge from this prior experimental work that bidding in sealed-bid first-price auctions tends to be more aggressive than the RNNE prediction.

Hypothesis 1B (H1B). *Bidding will be, on average, more aggressive than the equilibrium predictions of Proposition 1. The average cost to the buyer will be lower than the buyer's cost prediction that follows from Proposition 2.*

The second hypothesis tests the buyer cost equivalence between the sealed- and open-bid formats. We again state two versions of the hypothesis, the first based on our model in Section 3.2, and an alternative hypothesis based on our knowledge from prior experimental work that reports overly aggressive bidding in sealed-bid auctions.

Hypothesis 2A (H2A). *Average buyer's cost will not be significantly different under the open-bid and sealed-bid formats.*

Hypothesis 2B (H2B). *Average buyer's cost will be significantly higher under the open-bid format than under the sealed-bid format.*

We are the first to conduct laboratory tests of sealed-bid auctions with nonregular cost distributions, and because deviations from RNNE bidding may be affected by the cost distribution, we formulate the next hypothesis to specifically test whether the cost distribution affects bidding behavior and the buyer's cost.

Hypothesis 3 (H3). *Any systematic deviations from RNNE (specifically, overly aggressive bidding) will be observed in regular and nonregular distributions.*

Next, we test the prediction of Theorem 1. This is a qualitative test rather than a test about point predictions. According to Theorem 1, shifting some positive amount from the allocation of a low bidder to the allocation of a high bidder, with everything else remaining unchanged, would decrease the Herfindahl–Hirschman index (HHI) of allocation, and for any regular cost distribution should increase the buyer's average cost.

Hypothesis 4 (H4) (Optimality of Greedy Allocation). *For regular cost distributions, the average cost to the buyer will always decrease as the allocation becomes more greedy. For a nonregular cost distribution, a more greedy allocation may increase the cost to the buyer.*

A useful feature of our model is that it predicts the cost of multisourcing. For example, we can compute the predicted cost of multisourcing in our study by comparing the predicted cost of the buyer for any given set of parameters. More specifically, the cost of multisourcing is the cost that the buyer incurs by spreading its award more. We formulate our final hypothesis to test the cost of multisourcing predicted by the theory.

Hypothesis 5 (H5) (The Cost of Multisourcing). *Pairwise differences in average buyer cost will not be different from those predicted by the model (Proposition 2).*

Proposition 2 characterizes the buyer's expected cost as a function of splits. Therefore, H5 provides an indirect test of how well our model predicts the buyer's cost of multisourcing.

5.2. Experimental Design, Implementation, and Protocol

Our study includes 12 experimental treatments—all between-subjects. In all treatments, $n = 4$ suppliers compete for a contract to provide units of a commodity to a computerized buyer seeking 100 units in total (in a few treatments, the number of units is 102 to make all splits integer). Suppliers' costs are privately known. In nine of the treatments, we use a regular cost distribution; costs are distributed according to the uniform distribution from 0 to 100, $c_i \sim U(0, 100)$. In three treatments, we use a nonregular cost distribution with the p.d.f. described by Equation (9), scaled to the interval $[0, 100]$.

We present our design in Figure 3.[7] For the regular-distribution treatments, we conducted treatments with eight different split awards that vary in their market concentration (Herfindahl–Hirschman) index, going from 100% to 27.8%. We use these eight treatments to test H1A and H1B, and this design is presented in Figure 3(a). We have one open-bid treatment that uses the 40–35–25–0 split award, and comparing this treatment with the analogous sealed-bid treatment provides a test of H2A and H2B. This design is shown in Figure 3(b). The test of H3 is a 3×2 design presented in Figure 3(c) that manipulates the cost distribution (regular and nonregular) and three levels of market concentration (HHI = 100%, 33.3%, and 27.8%). And finally, the test of H4 is a 2×2 design in Figure 3(d), in which we manipulate the cost distribution (regular versus nonregular) and allocation (greedy versus nongreedy).

In total, 560 participants were included in our study. We randomly assigned participants to treatments. Each human subject participated in one treatment only. We conducted all sessions at a public university in the United States, in a computer laboratory dedicated to research. Our participants were students, mostly master-level, primarily business and engineering majors. We recruited them through SONA, an online recruitment system, offering the earning of cash as the only incentive to participate.

On arrival at the laboratory, the participants were seated at computer terminals in isolated cubicles. We handed out written instructions (see Online Appendix D for samples) to participants. After they read the instructions, we then read the instructions aloud before starting the auctions, to ensure common knowledge about the game's rules. Each session included 8–12 participants who competed in a series of 40 auctions. For each auction, we randomly rematched participants in each session, into two-three groups of four

Figure 3. Experimental Design to Test Hypotheses 1–4

(a) Design to test Hypothesis 1

Distribution: Regular
Format: Sealed-bid

HHI = 100%	HHI = 66.5%	HHI = 50%	HHI = 39.5%	HHI = 37.5%	HHI = 34.5%	HHI = 33.3%	HHI = 27.8%
100–0–0–0	80–15–5–0	50–50–0–0	50–35–15–0	50–25–25–0	40–35–25–0	34–34–34–0	34–34–17–17

(b) Design to test Hypothesis 2

Regular sealed-bid	40–35–25–0
Regular open-bid	40–35–25–0

(c) Design to test Hypothesis 3

	HHI = 100%	HHI = 33.3%	HHI = 27.8%
Regular sealed-bid	100–0–0–0	34–34–34–0	34–34–17–17
Nonregular sealed-bid	100–0–0–0	34–34–34–0	34–34–17–17

(d) Design to test Hypothesis 4

	Greedy allocation	Nongreedy allocation
	34–34–34–0	34–34–17–17
	34–34–34–0	34–34–17–17

bidders each. Typically, at least 24 participants were in the laboratory at the same time, and participants did not know the session size. In sealed-bid format treatments, each participant placed a single per-unit bid. After all of the bids were received, market shares were allocated according to the splits in the treatment.

In the open-bid format treatment (which included the 40–35–25 split), the per-unit price started at 100 and automatically decreased. Bidders could drop out of the auction by clicking a button on their screens. After the first bidder dropped out of the auction, each of the remaining three bidders were allocated 25 units at the price at which the first bidder dropped out. After the second bidder dropped out of the auction, the remaining two bidders were allocated an additional 10 units at the price at which the second bidder dropped out. And finally, after the third bidder dropped out of the auction, the remaining bidder was allocated the additional five units at the price at which the third bidder dropped out.

We programmed the experimental interface using the z-Tree system (Fischbacher 2007). At the end of each session, we computed cash earnings for each participant by multiplying the total earnings from all rounds by a predetermined exchange rate and adding it to a $5 participation fee. Participants were paid their earnings in private and in cash at the end of the session. Average earnings, including the show-up fee, were $25.

5.3. Results
5.3.1. Buyer Cost, Efficiency, and Individual Bidding Behavior with Regular Cost Distribution (Hypothesis 1A and 1B).
In Table 1, we display data from all sealed-bid regular-distribution treatments, comparing average buyer costs to their theoretical benchmarks. Average buyer cost is below the RNNE benchmark for all regular-distribution treatments, and the differences are statistically significant ($p < 0.05$ using a two-sided t-test). Here and in the rest of the results section, we use session average (the amount averaged over all auctions and all periods, for a given session) as the unit of analysis because sessions are independent, and report two-sided p-values from a t-test.[8]

Hypotheses 1A and 1B depend on the extent of the individual bidding behavior matching the equilibrium prediction. Figure 4 shows the scatter plots of bids for each split. We find that in regular-distribution treatments, most bids are between the cost and the equilibrium bid—i.e., bidders bid overaggressively.[9]

Table 2 reports allocational efficiency, which is always significantly below 100%.

Based on our regular-distribution sealed-bid treatments data, we can reject all aspects of H1A. The data are consistent with H1B. We conclude that bidding behavior in auctions with split awards is qualitatively similar to what has been observed in sealed-bid auctions without split awards, when cost distribution is uniform.

Table 1. Comparison of Average Buyer Cost and the RNNE Benchmarks Under the Sealed-Bid Format for Regular-Distribution Treatments

Splits	100–0–0–0	80–15–5–0	50–50–0–0	50–35–15–0	50–25–25–0	40–35–25–0	34–34–34–0	34–34–17–17
Optimal	40.73	51.38	61.82	67.62	70.67	75.27	80.85	87.68
Observed	32.89	42.79	50.87	56.45	60.48	62.83	74.46	82.26

Figure 4. Bids as a Function of Cost for Sealed-Bid Treatments with Regular Distribution
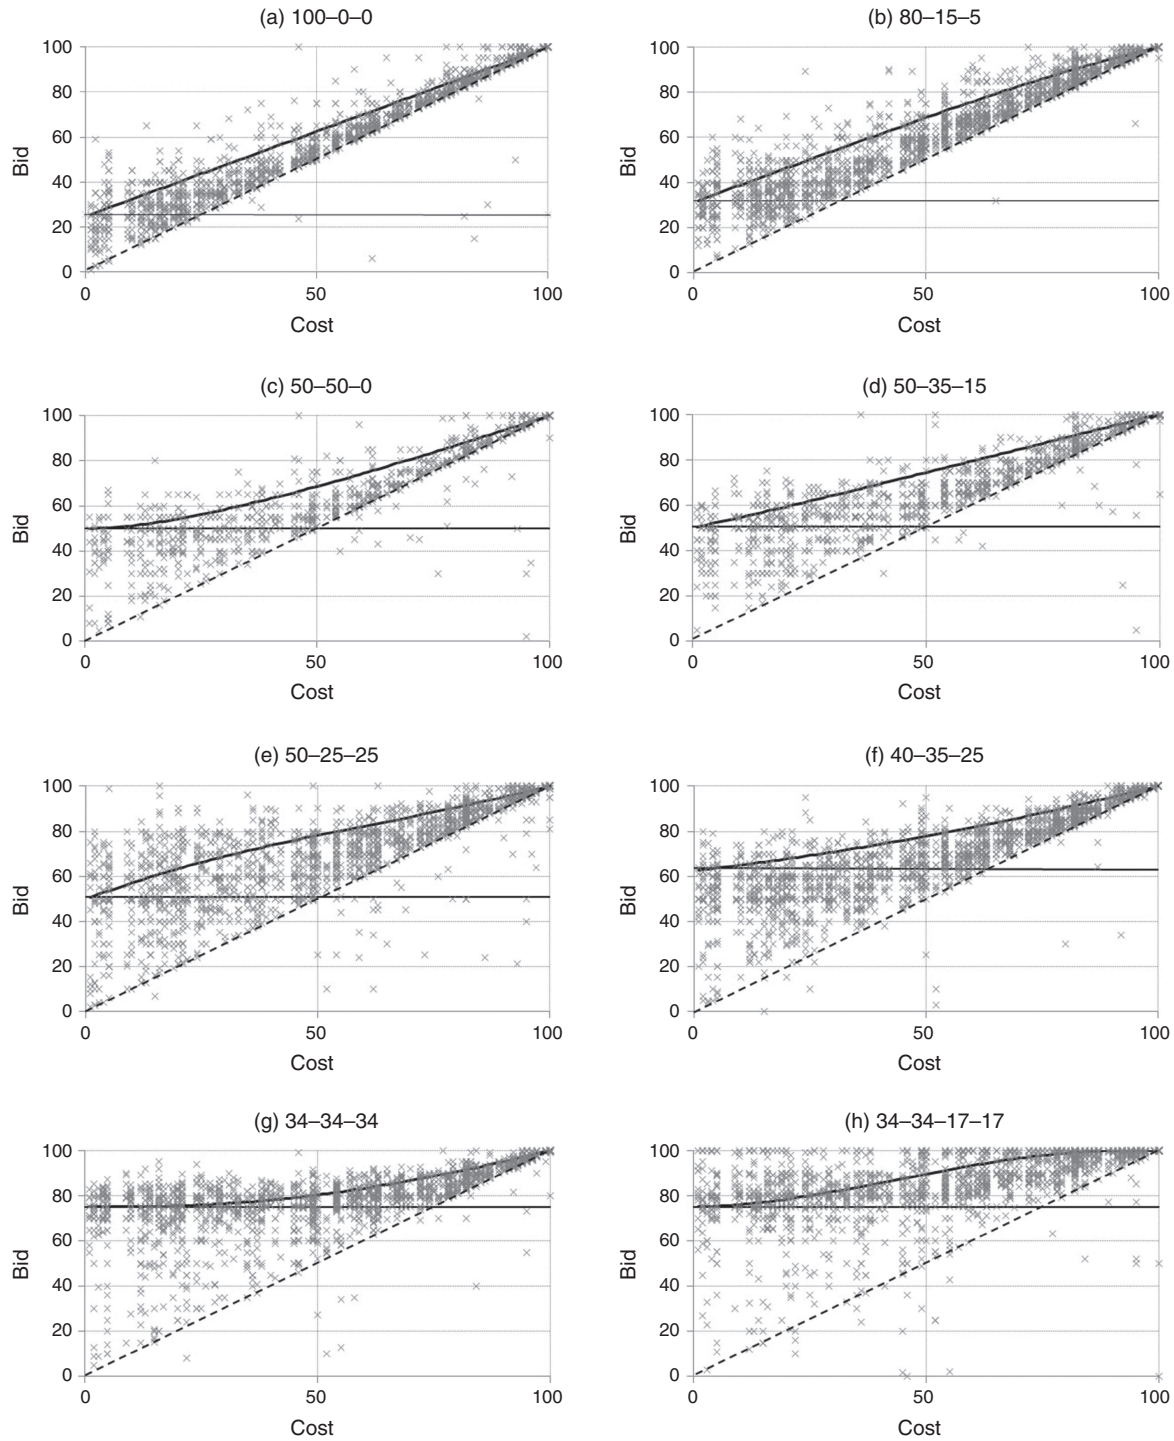
Note. RNNE is marked by a solid line, and the 45° line is marked by a dashed line.

Table 2. Proportion of Efficient Allocations in the Sealed-Bid Treatments

Splits	100–0–0–0	80–15–5–0	50–50–0–0	50–35–15–0	50–25–25–0	40–35–25–0	34–34–34–0	34–34–17–17
Allocation efficiency	0.865**	0.615**	0.890**	0.503**	0.729**	0.438**	0.696**	0.516**
Standard error	0.012	0.033	0.011	0.029	0.020	0.030	0.016	0.044

Note. Testing H_o that the proportion = 100%.
  **$p < 0.01$.

Table 3. Comparison of Average Buyer Cost for the 40–35–25–0 Split Under the Sealed- and Open-Bid Formats, and the RNNE Benchmark

	Sealed bid	Open bid
Optimal	75.27	75.27
Human	62.83	76.52

Table 4. Comparison of Average Buyer Cost and the RNNE Benchmarks Under the Sealed-Bid Format for Nonregular-Distribution Treatments

Splits	100–0–0–0	34–34–34–0	34–34–17–17
Optimal	27.76	92.86	88.75
Human	15.05	93.95	91.31

5.3.2. Buyer Cost Equivalence (Hypotheses 2A and 2B). We test buyer cost equivalence by comparing data from the sealed- and open-bid treatments with 40–35–25–0 split. Table 3 summarizes the test of buyer cost equivalence in those two treatments.

Under the sealed-bid format, the average buyer's cost for the 40–35–25 split is significantly below the RNNE, but for the same split, the average buyer's cost is only weakly different from the RNNE prediction when the open-bid format is used ($p = 0.0887$), and directionally, the average buyer's cost is slightly above predicted, rather than below. So the average buyer's cost is significantly lower under the sealed-bid format than under the open-bid format ($p < 0.001$), contrary to H2A and consistent with H2B. Comparing the two treatments in Figure 5, we find that overall bidding is much closer to the RNNE equilibrium under the open-bid format than under the sealed-bid format.

5.3.3. The Effect of Cost Distribution on Bidding Behavior and Buyer's Cost (Hypothesis 3). We test H3 with a 3×2 design that compares behavior in treatments with three different splits (100–0–0–0, 34–34–34–0, and 34–34–17–17) with regular and nonregular

cost distributions. Table 4 displays the average buyer costs, optimal and observed, in the treatments with nonregular cost distribution. Corresponding regular-distribution treatments are presented in Table 1. The observed buyer's cost in the 100–0–0–0 nonregular-distribution treatment is significantly below optimal, as it is in the 100–0–0–0 regular-distribution treatment. However, this is not the case for the 34–34–34–0 treatment, for which the observed buyer cost is not significantly different than optimal ($p = 0.213$), or the 34–34–17–17 treatment, for which the observed buyer's cost is actually above optimal ($p = 0.018$). We conclude that we can reject H3.

Figure 6 shows scatter plots of bids for nonregular-distribution treatments. We find that while in the 100–0–0–0 treatment there is a good deal of bidding activity below the bidding equilibrium, there is very little bidding activity in this area in the other two treatments (namely, 34–34–34–0 and 34–34–17–17).

5.3.4. Optimality of Greedy Allocation (Hypothesis 4). In this subsection, we test H4 with a 2×2 design that varies the cost distribution (regular and nonregular) and allocation (greedy and nongreedy). The greedy

**Figure 5.** Bids as a Function of Cost for Treatments with 40–35–25 Split
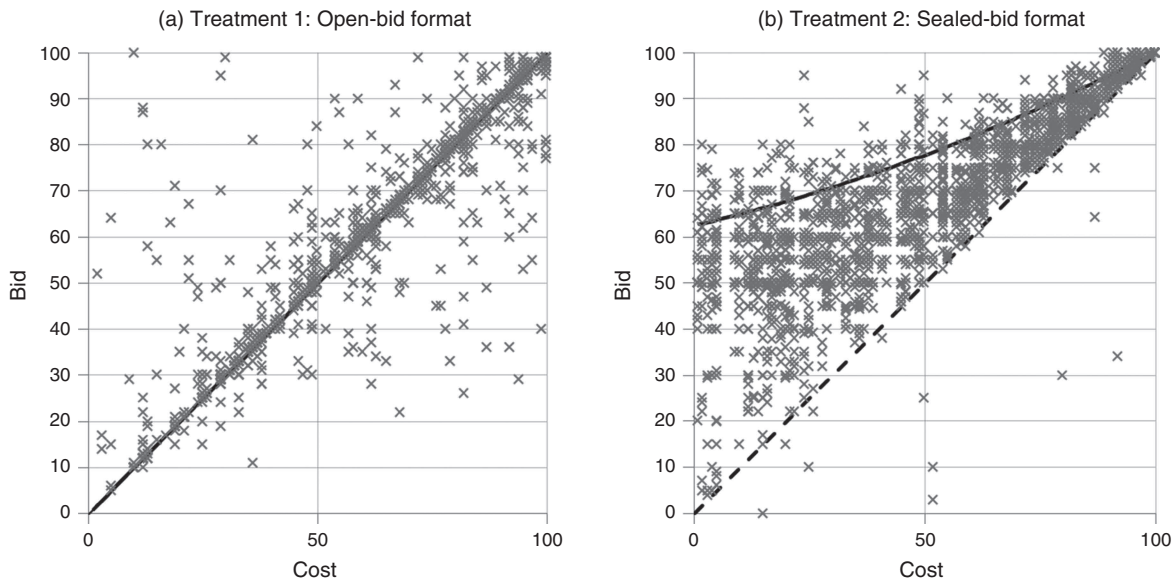
Note. RNNE is marked by a solid line, and the 45° line is marked by a dashed line.

Figure 6. Bids as a Function of Cost for Treatments with Nonregular Cost Distribution

Note. RNNE is marked by a solid line, and the 45° line is marked by a dashed line.

Figure 7. Average Buyer's Cost over the 40 Rounds of the Experiment
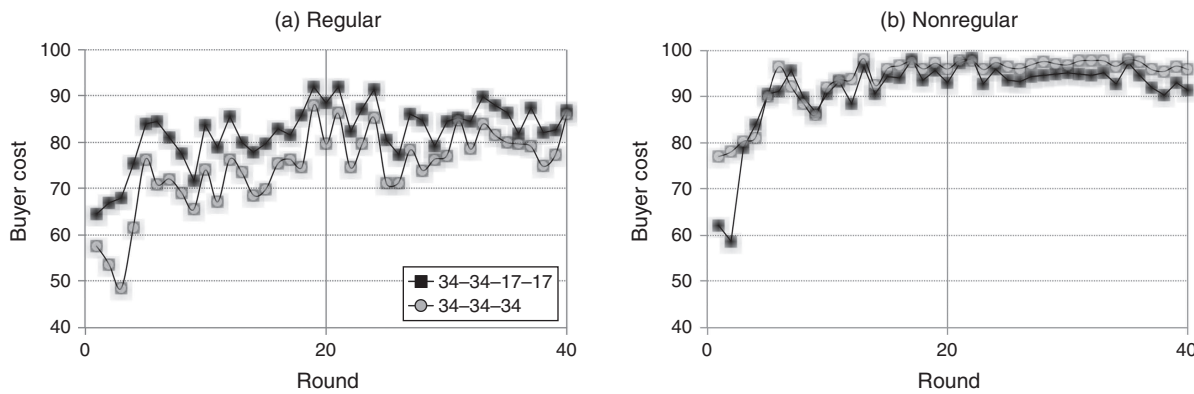
allocation is 34–34–34–0 and the nongreedy allocation is 34–34–17–17. Given the parameters in our experiment, the buyer's cost should be lower with the greedy allocation when the distribution is regular and with nongreedy allocation when it is nonregular. In the regular-distribution treatments, going from 34–34–34–0 to 34–34–17–17 significantly increases average buyer's cost from 74.46 to 82.26 ($p = 0.003$). In the nonregular-distribution treatments, going from 34–34–34–0 to 34–34–17–17 significantly decreases average buyer's cost from 93.95 to 91.31 ($p = 0.031$). We find support for H4.

Figure 7 plots the buyer's average cost for the two allocations for the regular distribution and for the nonregular distribution. Bidders in all four treatments learn to bid higher as they gain experience, but differences in average buyer's cost continue to be significant at the end of the session (p-values < 0.05).

5.3.5. The Effect of Multisourcing on the Buyer's Cost (Hypothesis 5). To directly test H5 (the extent to which the model is able to accurately predict the effect of multisourcing on the buyer's cost), we summarize pairwise differences in average buyer cost, and their

standard errors, for regular and nonregular-distribution treatments, in Table 5. For comparison, we also include in square brackets pairwise differences predicted by the model. Average pairwise differences are generally smaller than predicted, although most differences are not statistically significant. For the two nonregular-distribution treatments, the difference in average buyer's cost is 2.64, which is marginally smaller than the predicted difference of 3.98 ($p = 0.048$). So, although we can reject H5, we note that the model is fairly accurate overall in predicting the effect of multisourcing on the buyer's cost.

6. Discussion: Regret-Based Explanation

In the previous section, we observed that bidding is systematically below the risk-neutral Nash equilibrium (see Equation (5)) in all uniform distribution treatments and in the winner-take-all treatment for the nonregular distribution. However, we also observed that bidding was quite close to (slightly above) the risk-neutral equilibrium prediction in the multisourcing treatments for the nonregular distribution. In this section, we propose a simple behavioral model, based

Table 5. Pairwise Differences in Average Buyer Cost and Their Standard Errors

	Regular							Nonregular	
Treatment	80–15–5	50–50–0	50–35–15	50–25–25	40–35–25	34–34–34	34–34–17–17	34–34–17–17	34–34–34
100–0–0	9.90 (1.322) [10.65]	17.98 (3.495) [21.09]	23.66* (1.301) [26.89]	27.72 (2.607) [29.93]	29.93* (1.357) [34.54]	41.56 (1.700) [42.04]	49.37 (1.213) [49.04]	76.26* (1.817) [60.82]	78.90* (1.901) [64.57]
80–15–5		8.08* (3.504) [19.29]	13.77 (1.316) [16.24]	17.83 (2.620) [19.29]	20.04* (1.381) [23.89]	31.67 (1.719) [31.39]	39.47 (1.239) [38.40]		
50–50–0			5.69 (3.496) [5.80]	9.75 (4.198) [8.84]	11.96 (3.517) [13.44]	23.58 (3.663) [20.95]	31.39 (3.464) [27.95]		
50–35–15				4.06 (0.938) [3.05]	6.27 (1.360) [7.65]	17.90 (1.703) [15.15]	25.70* (1.216) [22.15]		
50–25–25					2.21 (2.638) [4.60]	13.84 (2.830) [11.35]	21.64 (2.566) [18.35]		
40–35–25						11.63 (2.024) [7.50]	19.43* (1.276) [14.51]		
34–34–34							7.80 (1.636) [7.00]	2.64 (0.949) [3.748]	

Note. Standard errors are in parentheses, and predicted differences (from Tables 1 and 4) are in square brackets.
  *$p < 0.05$.

on aversion to anticipated regret (see Engelbrecht-Wiggans and Katok 2007 and 2008), that qualitatively organizes our data.

In this model, the expected utility of a supplier may be affected by two kinds of regret. If a supplier does not win some allocation Q_l that would have been profitable, he experiences *loser's regret*. If a supplier wins allocation Q_w that he would have preferred to win at a higher price, he experiences *winner's regret*. Hence, the expected utility of a supplier with a marginal cost of c_i who submits a bid b_i, when every other supplier follows a bidding strategy $\beta_{\text{regret}}(c, \mathbf{Q})$, is given by

$$\Pi_i(b_i, c_i, \mathbf{Q}) = (b_i - c_i) H(\beta_{\text{regret}}^{-1}(b_i, \mathbf{Q}), \mathbf{Q})$$
$$- R_l(b_i, c_i, \beta_{\text{regret}}^{-1}(\cdot), \mathbf{Q})$$
$$- R_W(b_i, c_i, \beta_{\text{regret}}^{-1}(\cdot), \mathbf{Q}). \quad (10)$$

Here, $R_l(\cdot)$ and $R_w(\cdot)$ denote the expected loser's regret and expected winner's regret, respectively. Loser's and winner's regret can be characterized in many ways for the multisourcing case. For instance, the supplier could experience loser's and winner's regret on every allocation Q_m (for $m = 1, \ldots, n$). If supplier i fails to profitably win an allocation Q_m, then loser's regret would reduce its utility by $Q_m \cdot L \cdot (\beta_{\text{regret}}(c_m, \mathbf{Q}) - c_i)$, where $\beta_{\text{regret}}(c_m, \mathbf{Q})$ is the bid of the supplier who wins Q_m, and where parameter L is the weight on loser's regret. The anticipated loser's regret could be the maximum of the expected loser's regret on each

allocation. On the other hand, if a supplier i wins Q_m, then winner's regret would reduce the supplier's utility by $Q_m \cdot W \cdot (\beta_{\text{regret}}(c_{m+1}, \mathbf{Q}) - b_i)$, where $\beta_{\text{regret}}(c_{m+1})$ is the bid of the supplier who wins Q_{m+1}, and where parameter W is the weight on winner's regret. The anticipated winner's regret could be the maximum of expected winner's regret on each allocation. In Online Appendix F, we fully write down the expected loser's regret, $R_l(\cdot)$, and the expected winner's regret, $R_w(\cdot)$, for such a model.

One can also capture loser's regret and winner's regret through a simplified version of the above model wherein a supplier experiences loser's and winner's regret only on specific allocations (rather than on every possible allocation). Our goal is to develop a parsimonious behavioral model that can organize our data, so we assume that the supplier experiences loser's regret only on some allocation Q_l and winner's regret only on some allocation Q_w. Denoting by f_m the density function of the mth smallest order statistic out of $n - 1$ draws, we characterize this simple model of regret as

$$\Pi_i(b_i, c_i, \mathbf{Q}) = (b_i - c_i) H(\beta_{\text{regret}}^{-1}(b_i, \mathbf{Q}), \mathbf{Q})$$
$$- L \cdot \int_{x = \beta_{\text{regret}}^{-1}(c_i, Q)}^{\beta_{\text{regret}}^{-1}(b_i, Q)} Q_l \cdot (\beta_{\text{regret}}(x, \mathbf{Q}) - c_i) f_l(x) \, dx$$
$$- W \cdot \int_{y = \beta_{\text{regret}}^{-1}(b_i, Q)}^{\beta_{\text{regret}}^{-1}(\bar{c}, Q)} Q_w \cdot (\beta_{\text{regret}}(y, \mathbf{Q}) - b_i) f_w(y) \, dy. \quad (11)$$

Table 6. Average Observed and Predicted Buyer Cost Under the Risk-Neutral and Regret-Averse Models

	Uniform distribution								Nonregular		
Split	100–0–0	80–15–5	50–50–0	50–35–15	50–25–25	40–35–25	34–34–34	34–34–17	100–0–0	34–34–34	34–34–17–17
Risk-neutral	40.73	51.38	61.82	67.62	70.67	75.27	80.85	87.68	27.76	92.86	88.74
Observed	32.89	42.79	50.87	56.45	60.48	62.83	74.46	82.26	15.05	93.95	91.31
Regret	32.41	41.38	52.85	58.59	65.75	69.60	73.23	73.90	15.07	94.70	90.93

Note. For the uniform distribution, the regret parameters are $L = 0.97$ and $W = 0$, and for the nonregular distribution, the regret parameters are $L = 0.76$ and $W = 8.62$.

Here, Q_l is the allocation for which $Q_l − Q_{l+1}$ is maximum (for $l = 1, \ldots, n − 1$). In case of multiple Q_l's satisfying this condition, we take the Q_l with maximum l. Intuitively, Q_l is our focal quantity for loser's regret because winning an allocation smaller than Q_l causes the bidder to miss out on a significant amount of quantity. We denote by Q_w the allocation for which $Q_w − Q_{w+1}$ is minimum (for $w = 1, \ldots, n − 1$). In the case of multiple Q_w's satisfying this condition, we take the Q_w with minimum w. Intuitively, Q_w is our focal quantity for winner's regret because bidding higher and winning the next smaller allocation Q_{w+1} would minimally impact quantity but increase the profit margin per item. This is a useful model because it is parsimonious, tractable, and captures the qualitative features in our data.

Differentiating Equation (11) with regard to b_i and setting it to zero and substituting $b_i = \beta_{\text{regret}}(c_i, \mathbf{Q})$ (to obtain a symmetric bidding strategy) gives the following first-order condition:

$$\frac{\partial \beta_{\text{regret}}(c, \mathbf{Q})}{\partial c} = \left(-(\beta_{\text{regret}}(c, \mathbf{Q}) - c) \cdot \frac{dH(c, \mathbf{Q})}{dc} \right.$$
$$\left. + Q_l \cdot L(\beta_{\text{regret}}(c, \mathbf{Q}) - c) f_l(c) \right)$$
$$\cdot (H(c, \mathbf{Q}) + Q_w \cdot W \cdot \bar{F}(c))^{-1}. \quad (12)$$

Solving the above differential equation with the boundary condition of $\beta_{\text{regret}}(\bar{c}, \mathbf{Q}) = \bar{c}$ gives the equilibrium bidding function when bidders anticipate regret.

In Table 6, we compare the observed average cost of the buyer with the buyer's expected cost predicted by the risk-neutral model and by the regret model with the regret parameters that we fitted to our data. We fit regret parameters separately for the uniform treatments (within sample) and nonregular-distribution treatments (also within sample), and use the same parameters (L and W) for deriving equilibrium bids for all splits corresponding to a distribution.

Regret parameters that fit our data best are $L = 0.97$ and $W = 0$ for the uniform distribution and $L = 0.76$ and $W = 8.62$ for the nonregular distribution. So, consistent with previous work of Engelbrecht-Wiggans and Katok (2008), loser's regret is driving overly aggressive bidding with the uniform cost distribution. With the nonregular distribution, loser's regret is not affected much,

but winner's regret becomes very large. This is intuitive because in the two multisourcing treatments, low-cost bidders (in fact, bidders with cost below about 97) are almost certain to win market share of 34, so the anticipated regret of being paid less for the same market share is quite salient.

Finally, the regret-based model also explains the lack of expected cost equivalence between the open-bid and sealed-bid in the 40–35–25 treatment, because in the clock auction implementation of the open-bid mechanism, bidders have a dominant strategy to drop out at their cost. If they follow this strategy, they experience neither type of regret.

7. Conclusion

To incorporate sourcing rules into a simple-to-communicate and easy-to-understand auction, the buyer announces upfront the percentage of business that it would allocate to each supplier, depending on the rank of this supplier's bid. We study, analytically and experimentally, two types of these split-award auctions: the sealed-bid first-price auction, and a version of the open descending-price auction.

Many potential splits could satisfy the sourcing rules, and it is not obvious which would be the most cost-efficient split that should be announced before the auction. For instance: Should the buyer award the maximal allowed percentages to the lowest bids? Or would it be better to increase competition among low-cost bidders by widening the gaps between the percentages awarded to them, even though doing so means shifting some allocation percentages to higher-cost bidders? A managerially insightful result we find is that if the underlying cost distribution of the suppliers is well behaved (i.e., it is regular), then the most cost-efficient split (among all splits that satisfy the sourcing rules) would allocate greedily from the lowest to the highest bidder. However, when the underlying cost distribution is nonregular (e.g., bimodal), the most cost-efficient split might not allocate greedily—i.e., it might allocate more business to a higher bidder than the minimum amount required by the sourcing rules. These insights derived from the model are validated by our experimental results. Although the small but significant cost savings from nongreedy splitting (for nonregular distribution) might indicate robustness of the

greedy allocation, the interesting aspect of nongreedy splitting is that it aligns the buyer's cost reduction and diversification goals when the cost distribution is non-regular. Moreover, the experimental results show that our model does a good job of predicting the cost of multisourcing for the buyer. Knowing the costs and benefits of multisourcing is important for buyers looking to design effective sourcing rules.

Besides announcing splits ex ante, the buyer could use a more complex mechanism where it could announce only the sourcing rules up front and then announce the splits and payments ex post, based on the bids, such that they minimize the buyer's expected cost. Applying Myerson (1981), one can conclude that for regular distributions, the ex post allocations and the buyer's expected cost would remain the same in such a mechanism as compared to the simpler auction formats that we have discussed in this paper. For the nonregular distribution, we find that such an auction mechanism might give different allocations (and would result in lower expected cost for the buyer) as compared to announcing splits up front. However, we find that our main insight on the buyer awarding non-greedily for nonregular distributions still holds true for such a mechanism. For instance, for the nonregular distribution given in Equation (9), and for the allocation constraint that no supplier gets more than 1/3 of the buyer's business, one can show that the optimal mechanism would allocate 1/3, 1/3, 1/6, 1/6 (i.e., allocate nongreedily) if the two lowest bids are below 0.21 and the other two bids are above 0.21.

Even with up-front announcement of splits, one could investigate other auction formats besides the two that we investigate. For instance, an open reverse English auction can be implemented without a clock—bidders simply bid the auction price down using bid decrements of their choosing. Theoretically, such an implementation may admit a collusive equilibrium (such as Fugger et al. 2016), in which case, both theoretical and empirical properties of such a mechanism may be a promising direction for future research.

Finally, our main theoretical result is that a buyer can reduce its purchase cost by awarding business nongreedily to ex ante symmetric suppliers when the underlying cost distribution is not well behaved. In fact, if bidders are ex ante asymmetric, we can show (Online Appendix G) that allocating nongreedily can be optimal even if the underlying cost distributions are well behaved (i.e., regular). This underscores our paper's insight that nongreedy allocations can be optimal for a buyer, and presents an additional direction that may be promising for further experimental studies.

Acknowledgments

Appendix. Proofs

Proof of Proposition 1. This proof is shown in the following four steps: (1) We assume that all suppliers, except supplier i, follow an increasing and continuously differentiable bidding strategy β, and then characterize supplier i's first-order condition to find its surplus-maximizing bid b_i. (2) We then solve the first-order condition assuming that supplier i also uses bidding strategy $b_i = \beta(c_i)$, which gives us a closed-form solution to β. (3) We then verify that β is increasing and continuously differentiable. (4) Finally, we verify that if all suppliers except i follow β, then β also maximizes supplier i's surplus.

Step 1. To find the optimal bid b_i of supplier i, given that all other suppliers are bidding with an increasing and continuously differentiable bidding strategy β, we first differentiate Equation (4) with respect to b_i:

$$\frac{\partial \Pi(b_i, c_i, \cdot)}{\partial b_i} = H(\beta^{-1}(b_i, \cdot), \cdot) + (b_i - c_i)\frac{\partial H(\beta^{-1}(b_i, \cdot), \cdot)}{\partial b_i}. \quad \text{(A.1)}$$

Since $\partial \Pi(b_i, c_i, \cdot)/\partial b_i > 0$ at $b_i = c_i$ for any $\underline{c} \leq c_i < \bar{c}$, we know that the bid that maximizes supplier i's surplus will also be a solution to the first-order condition $\partial \Pi(b_i, c_i, \cdot)/\partial b_i = 0$.

Step 2. We now characterize β by substituting $b_i = \beta(c_i)$ in the first-order condition (A.1).

$H(c_i, \cdot) = -(\beta(c_i, \cdot) - c_i)(\partial H(c_i, \cdot)/\partial c_i) \cdot 1/(\partial \beta(c_i, \cdot)/\partial c_i)$. We can therefore characterize the bidding function through the following differential equation: $(\partial \beta(c_i, \cdot)/\partial c_i) \cdot H(c_i, \cdot) + (\beta(c_i, \cdot) - c_i)(\partial H(c_i, \cdot)/\partial c_i) = 0$, which can be rewritten as

$$\frac{\partial(H(c_i, \cdot) \cdot \beta(c_i, \cdot))}{\partial c_i} = c_i \frac{\partial H(c_i, \cdot)}{\partial c_i}. \quad \text{(A.2)}$$

At $c_i = \bar{c}$, it is an equilibrium strategy for supplier i to bid $\bar{c}$—i.e., $\beta(\bar{c}, \cdot) = \bar{c}$. Thus, integrating Equation (A.2) in the limits c_i to $\bar{c}$ gives us the solution

$$\beta(c_i, \cdot) = \bar{c} \cdot \frac{H(\bar{c}, \cdot)}{H(c_i, \cdot)} - \frac{1}{H(c_i, \cdot)}\int_{x=c_i}^{\bar{c}} x \cdot \frac{\partial H(x, \cdot)}{\partial x} \cdot dx. \quad \text{(A.3)}$$

Step 3. We had assumed that $\beta(c, \cdot)$ is continuously differentiable and increasing in c. Indeed, $\beta(c, \cdot)$ as characterized by Equation (A.3) is continuously differentiable. Next, we show that $\beta(c, \cdot)$ in Equation (A.3) is also increasing. Differentiating $\beta(c, \cdot)$ with respect to c in Equation (A.3), we get

$$\frac{\partial \beta(c, \cdot)}{\partial c} = -\frac{\beta(c, \cdot)}{H(c, \cdot)} \cdot \frac{\partial H(c, \cdot)}{\partial c} + \frac{c}{H(c, \cdot)} \cdot \frac{\partial H(c, \cdot)}{\partial c}$$
$$= -\frac{\beta(c, \cdot) - c}{H(c, \cdot)} \cdot \frac{\partial H(c, \cdot)}{\partial c}.$$

To determine the sign of $\partial \beta(c, \cdot)/\partial c$, we need to determine the sign of $\partial H(c, \cdot)/\partial c$. For this, we use the expression of $H(x, \cdot)$ from Equation (3). Thus,

$$\begin{aligned}
\frac{\partial H(c, \mathbf{Q})}{\partial c} = &-\binom{n-1}{1}\bar{F}(c)^{n-2}f(c) \cdot (Q_1 - Q_2) \\
&- 2\binom{n-1}{2}\bar{F}(c)^{n-3}F(c)f(c) \cdot (Q_2 - Q_3) \\
&- 3\binom{n-1}{3}\bar{F}(c)^{n-4}F(c)^2 f(c) \cdot (Q_3 - Q_4) - \cdots \\
&- (m-1)\binom{n-1}{m-1}\bar{F}(c)^{n-m}F(c)^{m-2}f(c) \\
&\cdot (Q_{m-1} - Q_m) - \cdots \\
&- (n-1)\binom{n-1}{n-1}F(c)^{n-2} \cdot (Q_{n-1} - Q_n).
\end{aligned}$$

For $Q_1 \geq Q_2 \geq \cdots \geq Q_n$, we get $\partial H(c, \cdot)/\partial c \leq 0$; therefore, $\beta(c, \cdot)$ is increasing in c.

Step 4. Finally, we show that $\beta(c, \cdot)$ forms a symmetric equilibrium—i.e., if all other $n-1$ suppliers bid with strategy β, then β would also maximize supplier i's surplus. Suppose that supplier i with marginal cost c chooses to bid $b = \beta(z, \cdot)$—i.e., it emulates a supplier with marginal cost z. Then its payoff, from Equation (4), is $\Pi(b, c, \cdot) = (\beta(z, \cdot) - c) \cdot H(z, \cdot)$. Substituting the value $\beta(z, \cdot)$, we get the supplier's payoff if it emulates a supplier of cost z as $\Pi(b, c, \cdot) = \bar{c}H(\bar{c}, \cdot) - \int_{x=z}^{\bar{c}} x(\partial H(x, \cdot)/\partial x) \cdot dx - cH(z, \cdot) = (z - c)H(z, \cdot) + \int_{x=z}^{\bar{c}} H(x, \cdot)\,dx$. If, instead, the supplier with cost c does bid $\beta(c, \cdot)$, then its payoff is $\Pi(\beta(c, \cdot), c, \cdot) = \int_{x=c}^{\bar{c}} H(x, \cdot)\,dx$. Since $H(c, \cdot)$ decreases in c, for any $z \neq c$, we get $\Pi(\beta(c, \cdot), c, \cdot) - \Pi(\beta(z, \cdot), c, \cdot) = (c - z)H(z, \cdot) + \int_{x=c}^{z} H(x, \cdot)\,dx \geq 0$. Therefore, β is indeed a symmetric equilibrium bidding strategy. □

Proof of Proposition 2. Using Equation (3), buyer's expected cost can be characterized as $C_{\text{buyer}} = n \cdot \int_{x=0}^{\bar{c}} \beta(x, \mathbf{Q}) \cdot H(x, \mathbf{Q})f(x)\,dx$. Substituting the value of β from Equation (5), we get $C_{\text{buyer}} = n\bar{c}Q_n - n\int_{x=0}^{\bar{c}} \int_{y=x}^{\bar{c}} y\,dH(y, \mathbf{Q})f(x)\,dx$. Integrating by parts, we get $C_{\text{buyer}} = n\int_{x=0}^{\bar{c}}(xH(x, \mathbf{Q}) + \int_{y=x}^{\bar{c}} H(y, \mathbf{Q})\,dy) \cdot f(x)\,dx$. Now, $n\int_{x=0}^{\bar{c}} xH(x, \mathbf{Q})f(x)\,dx = Q_1\mu_1 + Q_2\mu_2 + \cdots + Q_n\mu_n$. We next resolve $n\int_{x=0}^{\bar{c}} \int_{y=x}^{\bar{c}} H(y, \mathbf{Q})\,dy f(x)\,dx$ term by term. For this, take any mth term between 1 and $n-1$ of $H(y, \mathbf{Q})$—i.e., take $Q_m \binom{n-1}{m-1}\bar{F}^{n-m}(y)F^{m-1}(y)$. We then characterize

$$n \cdot \binom{n-1}{m-1} \int_{x=0}^{\bar{c}} \int_{y=x}^{\bar{c}} \bar{F}^{n-m}(y)F^{m-1}(y)\,dy f(x)\,dx$$
$$= m \cdot \binom{n}{m} \int_{x=0}^{\bar{c}} \int_{y=x}^{\bar{c}} \bar{F}^{n-m}(y)F^{m-1}(y)\,dy f(x)\,dx$$
$$= m \cdot \binom{n}{m} \int_{x=0}^{\bar{c}} \bar{F}^{n-m}(x)F^m(x)\,dx$$

(integrating by parts). Pearson (1902) showed that $\mu_{m+1} - \mu_m = \binom{n}{m}\int_{x=0}^{\bar{c}} \bar{F}^{n-m}(x)F^m(x)\,dx$ for any $m = 1, \ldots, n-1$. For $m = n$, we get $\bar{c} - \mu_n = \int_{x=0}^{\bar{c}} F^n(x)\,dx$. Thus, we have that

$$n \int_{x=0}^{\bar{c}} \int_{y=x}^{\bar{c}} H(y, \mathbf{Q})\,dy f(x)\,dx$$
$$= Q_1(\mu_2 - \mu_1) + 2Q_2(\mu_3 - \mu_2) + \cdots + m(\mu_{m+1} - \mu_m)Q_m + \cdots$$
$$+ (n-1)Q_{n-1}(\mu_n - \mu_{n-1}) + nQ_n(\bar{c} - \mu_n).$$

Finally, adding this term back, we get

$$C_{\text{buyer}} = \mu_2 Q_1 + (2\mu_3 - \mu_2)Q_2 + \cdots + (m\mu_{m+1} - (m-1)\mu_m)Q_m$$
$$+ \cdots + ((n-1)\mu_n - (n-2)\mu_{n-1})Q_{n-1}$$
$$+ (n\bar{c} - (n-1)\mu_n)Q_n. \quad \square$$

Proof of Lemma 1. We first show that maximizing $\sum Q_i^2$ such that sourcing constraints (1b) are satisfied would give a unique solution. We show this result by contradiction. Consider any two sets of allocations $\mathbf{Q}$ and $\mathbf{Q}'$ that maximize $\sum Q_i^2$ and satisfy the sourcing constraints (1b). Since one of the constraint is $\sum Q_i = \sum Q_i' = 1$, one can always find some $1 \leq l < m \leq n$ such that $Q_l > Q_l'$ and $Q_m' > Q_m$ (or the reverse (i.e., $Q_l < Q_l'$ and $Q_m' < Q_m$)). Assuming the former—i.e., $Q_l > Q_l'$ and $Q_m' > Q_m$—implies that an $0 < \epsilon < \min(Q_l - Q_l', Q_m' - Q_m)$ can be found such that $Q_l > Q_l' + \epsilon$ and $Q_m' - \epsilon > Q_m$. Moreover, $Q_l' + \epsilon$ and $Q_m' - \epsilon$ would also satisfy the sourcing

constraints. Now, $(Q_l' + \epsilon)^2 + (Q_m' - \epsilon)^2 = Q_l'^2 + Q_m'^2 + 2\epsilon(\epsilon + Q_l' - Q_m') > Q_l'^2 + Q_m'^2$. Thus, $\mathbf{Q}'$ cannot be an optimal solution. Following the same steps, one can show that for $Q_l < Q_l'$ and $Q_m' < Q_m$, $\mathbf{Q}$ cannot be optimal. Hence, maximizing $\sum Q_i^2$ such that constraints of Equation (1b) are satisfied would give a unique solution. Next, we show that this unique solution, denoted by $\mathbf{Q}$, is also greedy. We show this result also by contradiction—i.e., if $\mathbf{Q}$ is not greedy, then an allocation $\mathbf{Q}'$ can be found that satisfies the constraints of Equation (1b) and $\sum Q_i'^2 \geq \sum Q_i^2$. By definition of greedy allocation, if $\mathbf{Q}$ is not greedy, then for some $\epsilon > 0$ and some $1 \leq l < m \leq n$, the allocation vector $\mathbf{Q}'$ would satisfy sourcing constraints if $Q_l' = Q_l + \epsilon$, $Q_m' = Q_m - \epsilon$, and $Q_j' = Q_j$ for all $j = 1, \ldots, n$ excluding $j = l$ and excluding $j = m$. Now, $\sum Q_i'^2 = \sum Q_i^2 + 2\epsilon(\epsilon + Q_l - Q_m) > \sum Q_i^2$. □

Proof of Lemma 2. Since the decision variables $Q_1, \ldots, Q_n$ are real numbers not less than zero, and C_{buyer} in Equation (6) is linear in $Q_1, \ldots, Q_n$, hence, the math program (1) formulates a constrained fractional knapsack problem. Therefore, the optimal allocations are always greedy (for all sourcing constraints) if and only if the coefficients of Q_is in Equation (6) are increasing in i. This implies that the optimal allocation is always greedy to all but the lowest-bidding supplier if and only if $m\mu_{m+1} - (m-1)\mu_m \geq (m-1)\mu_m - (m-2)\mu_{m-1}$ is true for all $m > 2$ and if $n\bar{c} - (n-1)\mu_n \geq (n-1)\mu_n - (n-2)\mu_{n-1}$. For the lowest-bidding supplier, it can be shown that $\mu_2 \leq m\mu_{m+1} - (m-1)\mu_m$ for all $m \geq 2$ since $\mu_2 - \mu_{m+1} \leq 0 \leq (m-1)(\mu_{m+1} - \mu_m)$ for all $m \geq 2$. Also, $\mu_2 \leq n\bar{c} - (n-1)\mu_n$ since $\mu_2 - \bar{c} \leq 0 \leq (n-1)(\bar{c} - \mu_n)$ for any $n > 1$ suppliers participating in the auction. Hence, allocating greedily to the lowest-bidding supplier is always optimal. □

Proof of Theorem 1. Rearranging the terms of Equation (6) gives $C_{\text{buyer}}(\mathbf{Q}) = \sum_{j=1}^{n}(\mu_j Q_j + \sum_{i=j}^{n}(\mu_{i+1} - \mu_i)Q_i)$. For a given vector of costs $\mathbf{c} = c_1, \ldots, c_n$, we define

$$q(x, \mathbf{c}) \equiv \begin{cases} Q_i & \text{if } C_{i:n} \leq x \leq C_{i+1:n} \text{ for all } i = 1, \ldots, n-1, \\ Q_n & \text{if } C_{n:n} \leq x \leq \bar{c}, \end{cases}$$

where $C_{j:n}$ represents the jth-lowest cost from the sample $\mathbf{c}$. One can then express $C_{\text{buyer}}(\mathbf{Q}) = \mathbb{E}_c(\sum_{j=1}^{n}(C_{j:n}Q_j + \int_{x=C_{j:n}}^{\bar{c}} q(x, \mathbf{c})dx)) = \mathbb{E}_c(\sum_{i=1}^{n}(c_i q(c_i, \mathbf{c}) + \int_{x=c_i}^{\bar{c}} q(x, \mathbf{c})dx))$. Integrating by parts gives $\mathbb{E}_c(\int_{x=c_i}^{\bar{c}} q(x, \mathbf{c})dx) = \mathbb{E}_c(q(c_i, \mathbf{c})F(c_i)/f(c_i))$. Thus, $C_{\text{buyer}}(\mathbf{Q}) = \mathbb{E}_c(\sum_{i=1}^{n}(c_i + F(c_i)/f(c_i))q(c_i, \mathbf{c}))$. Thus, for a regular distribution (i.e., $c + F(c)/f(c)$ increasing in c), shifting an $\epsilon > 0$ from Q_i to Q_j for any $i < j$, with everything else being the same, would increase $C_{\text{buyer}}(\mathbf{Q})$. Finally, shifting $\epsilon > 0$ from Q_i to Q_j for any $i < j$, with everything else being the same, would increase the Herfindahl–Hirschman index. □

Endnotes

[1] Multiunit forward auctions for wireless spectrum have used set-asides and spectrum caps (Cramton et al. 2011), which help entrants win spectrum while limiting the amount of spectrum that incumbents can control in a geographic area. Although similar in spirit to sourcing rules in industrial procurement, politics can play an important role in how such set-asides and caps are deployed by governments.

[2] Specifically, a regular distribution is defined as a continuous distribution for which $c + F(c)/f(c)$ is increasing in c over the entire domain (for all c).

[3] We get nonmonotonicity of $c + F(c)/f(c)$ (hence, nonregularity) when the probability density f contains one or more portions that has a "U-shape."

[4] In Section 3, we explicitly characterize $C_{\text{buyer}}(\mathbf{Q})$ for the sealed-bid and open-bid auctions.

[5] Note that this descending-price-clock auction is similar in spirit to the *clinching* auction mechanism used for the auctioning of multiple units in an ascending-price-clock auction (as discussed in Ausubel 2004).

[6] For $n = 3$, an arc-sine distribution gives $\mu_1 = 0.196$, $\mu_2 = 0.5$, $\mu_3 = 0.804$. Thus, the coefficients of Q_3, Q_2 in Equation (6) are $3 - 2\mu_3 = 1.3921$ and $2\mu_3 - \mu_2 = 1.1079$, respectively.

[7] Sample sizes: all treatments included 4 sessions of 12 participants (48 participants), with the following exceptions. Sealed-Bid 40–35–25–0 and 80–15–5–0 treatments each included 4 sessions of 12 (60 participants), the 50–35–15–0 split treatment included 4 sessions of 8 (32 participants), and the 50–50–0–0 split treatment included 3 sessions of 8 (24 participants).

[8] We used the data from periods 1–40 in the analysis we report. We also repeated the same analysis for periods 21–40 only, to check whether learning causes any difference in our conclusions, and we found that none of the significance levels were affected.

[9] To further highlight overaggressive bidding, we performed a Tobit regression with bids as dependent variable and with splits and cost as independent variables. In Online Appendix E, we present results of this Tobit regression.

References

Anton JJ, Yao DA (1992) Coordination in split award auctions. *Quart. J. Econom.* 107(2):681–707.

Ausubel LM (2004) An efficient ascending-bid auction for multiple objects. *Amer. Econom. Rev.* 94(5):1452–1475.

Bagnoli M, Bergstrom T (2005) Log-concave probability and its applications. *Econom. Theory* 26(2):445–469.

Beall S, Carter C, Carter PL, Germer T, Hendrick T, Jap S, Kaufmann L, Maciejewski D, Monczka R, Petersen K (2003) The role of reverse auctions in strategic sourcing. Report, Center for Advanced Procurement Studies Research, Tempe, AZ.

Bichler M, Guler K, Mayer S (2014) Split-award procurement auctions—Can Bayesian equilibrium strategies predict human bidding behavior in multi-object auctions? *Production Oper. Management* 24(6):1012–1027.

Bichler M, Davenport A, Hohner G, Kalagnanam J (2006) Industrial procurement auctions. *Combinatorial Auctions* (MIT Press, Cambridge, MA).

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

Chaturvedi A, Beil DR, Martínez-de-Albéniz V (2014) Split-award auctions for supplier retention. *Management Sci.* 60(7):1719–1737.

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

Chen W, Dawande M, Gupta S, Janakiraman G (2015) Optimal descending mechanisms for constrained procurement. *Production Oper. Management* 24(12):1955–1965.

Cramton P, Kwerel E, Rosston G, Skrzypacz A (2011) Using spectrum auctions to enhance competition in wireless services. *J. Law Econom.* 54(4):167–188.

Dasgupta S, Spulber SF (1990) Managing procurement auctions. *Inform. Econom. Policy* 4(1):5–29.

Elmaghraby WJ, Katok E, Santamaría N (2012) A laboratory investigation of rank feedback in procurement auctions. *Manufacturing Service Oper. Management* 14(1):128–144.

Engelbrecht-Wiggans R, Katok E (2007) Regret in auction: Theory and evidence. *Econom. Theory* 33(1):81–101.

Engelbrecht-Wiggans R, Katok E (2008) Regret and feedback information in first-price sealed-bid auctions. *Management Sci.* 54(4):808–819.

Fischbacher U (2007) z-Tree: Zurich toolbox for ready-made economic experiments. *Experiment. Econom.* 10(2):171–178.

Fugger N, Katok E, Wambach A (2016) Collusion in dynamic buyer-determined reverse auctions. *Management Sci.* 62(2):518–533.

Haruvy E, Katok E (2013) Increasing revenue by decreasing information in procurement auctions. *Production Oper. Management* 22(1):19–35.

Hohner G, John R, Ng E, Davenport AJ, Kalagnanam JR, Lee HS, An C (2003) Combinatorial and quantity-discount procurement auctions benefit mars, incorporated and its suppliers. *Interfaces* 33(1):23–35.

Jap SD (2003) An exploratory study of the introduction of online reverse auctions. *J. Marketing* 67(3):96–107.

Jap SD (2007) The impact of online reverse auction design on buyer-supplier relationships. *J. Marketing* 71(1):146–159.

Kagel JH, Levin D (2015) Auctions: A survey of experimental research. Kagel JH, Roth AE, eds. *The Handbook of Experimental Economics*, Vol. 2 (Princeton University Press, Princeton, NJ), 563–637.

Kagel JH, Roth AE, eds. (1995) *The Handbook of Experimental Economics* (Princeton University Press, Princeton, NJ).

Katok E, Kwasnica AM (2008) Time is money: The effect of clock speed on seller's revenue in Dutch auctions. *Experiment. Econom.* 11(4):344–357.

Klotz DE, Chatterjee K (1995) Variable split awards in a single-stage procurement model. *Group Decision Negotiation* 4(4):295–310.

Krishna V (2010) *Auction Theory* (Academic Press, Burlington, MA).

Myerson R (1981) Optimal auction design. *Math. Oper. Res.* 6(1):58–73.

Pearson K (1902) Note on Francis Galton's problem. *Biometrika* 1(4):390–99.

Tunca TI, Wu Q (2009) Multiple sourcing and procurement process selection with bidding events. *Management Sci.* 55(5):763–780.

Wambach A (2002) A simple result for the revenue in share auctions. *Econom. Lett.* 75(3):405–408.