M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 14, No. 1, Winter 2012, pp. 128–144
ISSN 1523-4614 (print) | ISSN 1526-5498 (online)

informs
http://dx.doi.org/10.1287/msom.1110.0354

A Laboratory Investigation of Rank Feedback in Procurement Auctions

Wedad J. Elmaghraby
Robert H. Smith School of Business, University of Maryland, College Park, Maryland 20742,
welmaghr@rhsmith.umd.edu

Elena Katok
Smeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802,
ekatok@psu.edu

Natalia Santamaría
Departamento de Ingeniería Industrial, Universidad de Los Andes, Bogotá D.C., Colombia,
n-santam@uniandes.edu.co

A popular procurement auction format is one in which bidders compete during a live auction event but observe only the rank of their own bid and not the price bids of their competitors. We investigate the performance of auctions with rank feedback in a simple setting for which analytical benchmarks are readily available. We test these benchmarks in the laboratory by comparing the performance of auctions with rank-based feedback to auctions with full-price feedback as well as to auctions with no price feedback (sealed-bid auctions). When bidders are risk-neutral expected-profit maximizers, the buyer's expected costs should be the same under rank and full-price feedback; moreover, expected buyer costs should be the same as in a sealed-bid auction. However, when we test this theoretical equality in a controlled laboratory setting we find that, consistent with practitioners' beliefs but contrary to our model, rank feedback results in *lower* average prices than full-price feedback. We identify two behavioral reasons for the difference. The first explanation is based on the similarity of the bidders' problem in a sealed-bid first-price auction and an open-bid auction with rank feedback. The second explanation incorporates the use of jump bids motivated by bidder impatience.

Key words: procurement auctions; bidding behavior; laboratory experiments; rank; feedback
History: Received: March 12, 2010; accepted: July 27, 2011. Published online in *Articles in Advance* December 2, 2011.

1. Introduction

In the last decade electronic auctions became part of the procurement toolkit; their ability to reduce transaction costs and lower purchasing costs through improved competition are two of the main reasons managers have included them into their procurement arsenal. In the early 1990s, almost all auctions were open-bid auctions with full-price feedback: Suppliers (bidders) could watch a computer screen that tracked the past and current bids from each bidder and watch as they and their opponents dynamically updated their bids throughout the duration of the auction. The early preference for full-price feedback auctions stemmed from the belief that the psychology of an open auction induces suppliers to bid more aggressively.

However, as the use of auctions expanded, the wisdom of this belief has been called into question. For example, as bidder pools started to include both domestic and foreign suppliers, practitioners noted that domestic suppliers frequently stopped bidding early on in the auction above their historic bids/contract prices when they saw the low bids submitted by their foreign competitors under full-price feedback. Empirical support for this argument can be found in Carter et al. (2004); they find that a full disclosure auction could be unsuccessful if the initial bid was sufficiently low to stifle competitive bidding by the other suppliers. This reluctance of domestic suppliers to bid competitively frustrated buyers because (i) the final award selection is typically based on the consideration of multiple price and nonprice attributes; (ii) in many procurement settings, only a subset of these attributes is explicitly brought into the auction for competitive bidding and feedback, and the remaining attributes are excluded from the auction feedback process but play a role in the final award selection; and (iii) although domestic suppliers typically have higher costs than their foreign competitors, they frequently dominate their foreign opponents along noncost dimensions.

To overcome the potential anticompetitive effect of full-price feedback, auction designers have begun to adopt rank feedback format; analyzing the effect of rank feedback is the main contribution of our work. Under rank feedback, bidders can see the ordinal ranking of their bid (e.g., if they were the lowest bidder they would have a rank of 1) but not the actual bids of their competitors. The use of rank feedback is increasing and now overshadows the use of full disclosure feedback (Elmaghraby 2007). From our discussion with a senior category manager at Ariba (one of the leading business-to-business auction companies) we learned that more than 70% of the auctions at Ariba are of the rank feedback format. Rank feedback is now used in settings in which the information contained in the rank feedback is not always sufficient for a bidder to determine whether he is currently winning. Again, from our discussions with the senior category manager at Ariba, we learned that for the more than 500 categories in which Ariba has experience, only *one* category is awarded based on bid prices alone. For direct materials, buyers are willing to identify supplier-specific transformational factors, provide auction feedback based on the transformed bids, and make awards based on these transformed bids. For indirect materials and services, however, the buyers prefer to base the auction feedback on price only and bring nonprice terms into consideration for the final selection only after the conclusion of the auction. In addition to transformational bids, some auctions designers have adopted richer auction bid formats to allow suppliers to compete along multiple dimensions (Koppius and van Heck 2003, Gallien and Wein 2005, Chen-Ritzo et al. 2005); however, these multidimensional auctions are more the exception than the rule (Jap 2002, Teich et al. 2004, Bichler and Kalagnanam 2005).

The use of rank feedback when (i) the final award is based on multiple attributes and (ii) not all of these attributes are reflected in the rank feedback poses an interesting bidding challenge.

Consider a setting in which (i) the buyer evaluates suppliers along multiple dimensions but only a subset of these dimensions is used in providing rank feedback and (ii) each supplier is fully informed of the postauction evaluation and selection method. If a supplier can see his opponents' bids (full feedback), then he knows for certain whether his bid is winning or losing and can, if necessary, calculate by exactly how much the losing bid should be decreased in order to become a winning bid (provided that is a profitable move). However, rank feedback does not allow bidders the same luxury; rank feedback introduces ambiguity around the winning (and losing) status of a bid precisely because it does not allow bidders to see their opponents' bids. It is possible for a supplier to be at rank 1 but not be winning when other nonprice factors are taken into consideration. Conversely, it is possible for a supplier at rank 2 or 3 to be winning when all factors are considered. Anecdotal evidence from rank feedback auctions in Carter et al. (2004) suggests that the ambiguity surrounding opponents' bids, and the winning status of a supplier's own bid, stimulates more aggressive bidding behavior, particularly from bidders in second or third place. In addition to anecdotal evidence that rank feedback reduces procurement costs, rank feedback is believed to be less harmful to the buyer-supplier business relationship because it provides more information privacy for the suppliers (Jap 2007, Carter and Stevens 2007). For example, Mithas and Jones (2007) find that rank feedback is typically used when the buyer is concerned with market-price confidentiality.

Although the use of rank feedback is rampant and generally thought by practitioners to decrease the buyer's expected total cost, we are the first to systematically analyze the effect of rank feedback on procurement costs.

We present a parsimonious model, designed specifically to explore the effect of rank feedback while abstracting away many other confounding features. We consider a model in which (i) the buyer faces two suppliers who may differ in their cost and quality; (ii) the "cost" component of a supplier's type, which is always the supplier's private information, is meant to capture all of the attributes that the supplier is able to explicitly include in the formal auction bid and on which he can receive feedback; and (iii) the "quality" component of the supplier's type reflects all the other attributes that are important to the buyer.

We first explore a setting in which suppliers have highly dispersed costs and each supplier's quality is common knowledge. For this setting, we show that when bidders are risk-neutral expected-profit maximizers, the buyer's expected costs should be the same under rank and full-price feedback; moreover, expected buyer costs should be the same as in a sealed-bid auction. However, when we test this theoretical equality in a controlled laboratory setting, we find that, consistent with practitioners' beliefs but contrary to our model, rank feedback results in *lower* average prices than full-price feedback.

With our parsimonious model and matching experimental framework, we can attribute certain features of the observed bidding behavior under rank feedback to a feature that we term *the sealed-bid effect* and we can link the regularities we observe back to existing literature. We also identify factors that cannot be explained by the sealed-bid effect and show that bidder impatience is one explanation for this behavior. We then test the robustness of our results and provide evidence that the competitive effects of rank feedback are undiminished when either (i) the supplier's costs significantly overlap and/or (ii) an opponent's quality is unknown.

In §2 we present a review of the relevant literature, followed by our theoretical model (§3 and §4), the experiments (§5), some extensions (§6), and the concluding remarks (§7).

2. Literature Review

Bichler and Steinberg (2007), Elmaghraby (2007), and Rothkopf (2007) provide comparisons between and discussions of the state of the art in the theory and practice of procurement auctions.

Our work is similar in spirit to Haruvy and Katok (2010), who study the effect of price visibility in procurement auctions. Haruvy and Katok (2010) compare a buyer's cost under a sealed-bid auction (RFP) and an open-bid dynamic auction. They find that the RFP format generates higher buyer surplus than does the open-bid format due to overly aggressive bidding. In our setting, we compare the effect of price visibility in two open-bid auction formats: rank versus full-price feedback. Consistent with Haruvy and Katok (2010), we find that the format with less price visibility (rank feedback) consistently generates the lowest procurement cost for the buyer.

Compte and Jehiel (2007), Engelbrecht-Wiggans et al. (2007), and Kostamis et al. (2009) consider auctions that are conducted on price, but in which the buyers award contracts based on a combination of price and nonprice attributes. Engelbrecht-Wiggans et al. (2007) analyze a model (which they also test in the laboratory) to compare a price-based mechanism, in which the bid taker commits to awarding the contract to the winner of the auction, to a buyer-determined mechanism, in which the bid taker is free to select the bid that maximizes his surplus (which is not necessarily the lowest bid from the auction because the surplus includes nonprice attributes). Their results show that the buyer-determined mechanism is better for situations with many bidders and a high correlation between the supplier's nonprice attributes and his cost. In situations with few bidders and less dependence between the supplier's value and cost, the price-based mechanism performs better. Whereas Engelbrecht-Wiggans et al. (2007) attempt to understand how different selection mechanisms—in particular, price-only based versus price and quality based (buyer determined) mechanisms—affect bidder behavior, we are trying to understand how different *information feedback* formats affect bidder behavior given the same selection method. In the auctions we consider here, the buyer always selects the supplier based on the lowest *quality-adjusted* bid price; what differs across auctions is the interim feedback bidders receive as they bid in an open-bid auction.

Compte and Jehiel (2007) compare dynamic and static (forward) auction formats with respect to the bidders' incentives for information acquisition and with respect to the revenue they generate for the bid taker. Under dynamic auctions, it is assumed that bidders can gain more information about their opponents as the auction progresses (as opposed to no information acquisition opportunities under sealed-bid (static) auctions). The authors conclude that dynamic procedures generate higher revenues than do their static counterparts, the main reason being the ability of bidders to incorporate the information into their bidding strategy; the better informed the bidders are, the higher the revenues for the bid taker (given that the number of competitors is not too small). In contrast, in our experiments we find that the opportunity offered by full-price feedback to gather information about an opponent's costs actually reduces the buyer's welfare by increasing her average costs.

Kostamis et al. (2009) consider a setting in which a buyer faces suppliers who each have a distinct, exogenously determined quality level, represented in the buyer's objective function as a quality adjustment cost. The buyer wishes to select a supplier that minimizes her total cost of procurement, i.e., price plus nonprice quality adjustment cost. Kostamis et al. (2009) compare the performance of a first-price sealed-bid auction, in which the suppliers' quality levels are hidden, with an open descending auction, whereby suppliers can observe their competitors' total cost bids. This form of feedback allows a supplier to determine the minimum level of adjustment needed to his own total cost bid to displace the current winner (adjustments to total cost bid are made by changing the bid price). The authors conclude that the superior (lower cost) auction format critically depends on the realizations of the nonprice attributes and suppliers' beliefs in each format.

In theory, if bidders are risk neutral, the expected revenues are the same in the four major (forward) auction formats (the sealed-bid first price, the sealed-bid second-price, the open-bid ascending (English), and the clock descending (Dutch); Vickrey 1961). However, if bidders are risk averse, that equivalence holds only between the sealed-bid first price and Dutch and the sealed-bid second price and English. Virtually all laboratory work to date that deals with revenue equivalence in auctions deals with forward auctions (see Kagel 1995 for a review). Generally, laboratory tests reject all versions of revenue equivalence. Sealed-bid first-price revenues were reported to be higher than Dutch revenues (Cox et al. 1982), but later Lucking-Reiley (1999) as well as Katok and Kwasnica (2008) found that prices in the Dutch auction critically depend on the speed of the clock. In addition, several models of bidder impatience have been offered (Carare and Rothkopf 2005, Katok and Kwasnica 2008). Similarly, there is no support for the revenue equivalence between the sealed-bid second-price and English auctions because although both formats have the same dominant bidding strategy, bidders in English auctions tend to follow it, whereas bidders in sealed-bid second-price auctions tend to place bids above their valuations (Kagel 1995).

The literature stream comparing revenues in sealed-bid first-price and English auctions, however, is most relevant to our study. Cox et al. (1988) were the first to report that bidding in sealed-bid first-price auctions is more aggressive than it should be in equilibrium, and thus the revenue is higher when the sealed-bid first-price auction is used than when the English auction is used. Cox et al. (1988) show that qualitatively, this difference is consistent with risk aversion. Equivalently, in a procurement setting, Holt (1980) shows that when bidders are risk averse, the expected procurement cost in equilibrium is lower in sealed-bid first-price auctions than in their open-descending counterparts. But as Kagel (1995, p. 525) points out, in sealed-bid first-price auctions, "risk aversion is one element, but far from the only element, generating bidding above the [risk-neutral Nash equilibrium]."

A number of studies show that risk aversion does not fully explain the results obtained in many auction-like controlled laboratory experiments settings (Kagel et al. 1987, Kagel and Levin 1993, Cason 1995, Isaac and James 2000). More recent studies propose other explanations, such as aversion to regret (Engelbrecht-Wiggans and Katok 2007, 2008; Filiz-Ozbay and Ozbay 2007); learning (Ockenfels and Selten 2005, Neugebauer and Selten 2006); and simply reacting to errors (Goeree et al. 2002). Although the precise explanation for overly aggressive bidding in sealed-bid first-price auctions appears to be elusive, and is beyond the scope of this paper, that bidders tend to bid more competitively in sealed-bid than in open-bid auctions appears to be quite robust and general. We show that this regularity may apply to a wider set of auctions than just sealed-bid first price; the "sealed-bid effect" we identify applies also to dynamic auctions in which bidders are not certain whether they are winning or losing the auction.

3. The Model

The main purpose of our analysis is to define a parsimonious setting specially designed to provide clear analytical benchmarks for the optimal bidding behavior in auctions with rank feedback. We study a setting in which a buyer wishes to procure a service (or good) from one of two potential suppliers, each of whom could either be a high-type (H-type) or low-type (L-type) supplier (the assumption of two bidders does not limit the results of our model; we adopt it to simplify the exposition and experimental setting). A supplier is defined by his quality as well as the distribution of his production cost: We consider a setting in which an H-type bidder (privately) draws his cost, c_H, from a Uniform distribution between $[h, 100+h]$ and has a (known) quality *discount* factor given by $q_H = h$. An L-type bidder (privately) draws his cost, c_L, from a Uniform distribution between $[0, 100]$ and has a (known) quality discount factor given by $q_L = 0$. We assume that the identity of both suppliers is known to all *before* the start of the auction (in our laboratory experiments, we displayed the opponent's identity prominently on the bidding screens). The buyer may find herself procuring from a symmetric supply base (both suppliers are either H-types or L-types) or from an asymmetric supply base (one supplier is an H-type, the other an L-type). It is important to note that the suppliers' cost support as well as the quality discounts were specifically selected to render H-type and L-type suppliers effectively equivalent. That is, in the quality-adjusted space, L-type and H-type suppliers are symmetric and hence can be viewed as drawing their quality-adjusted costs from the same distribution $U[0, 100]$.

The buyer runs an open-descending auction and at the conclusion of the auction, the bidder with the lowest quality-adjusted bid, $b_i - q_i$, wins the auction and earns $b_i - c_i$, $i = \{L, H\}$. If several bidders placed identical price bids, there is an additional tie-breaking rule; for example, the bid placed earlier receives the lowest rank. Hence, an H-type's quality-adjusted bid, given that its last bid in the auction was b_H, is given by $b_H - h$; an L-type's quality-adjusted bid is equal to its last submitted bid b_L (it receives zero quality discount).

We consider two different price feedback mechanisms for the open-descending auction. Under *full-price* feedback, both bidders observe all price bids placed by their competitor; under *rank* feedback, each bidder can observe only the *rank* of her own price bid but not the actual bids of her competitor. It is important to emphasize that, consistent with practice, under both full and rank feedback formats, suppliers are given feedback on their actual (price) bid position, not their quality-adjusted bid position. We assume that the auction event has a fixed starting time and a soft close, meaning that bidding continues until some prespecified time and the auction ends when there are no bids placed for some prespecified amount of time. This type of ending rule is common in procurement auctions because it precludes the possibility of sniping (placing a bid in the final second of the auction, which is generally considered an undesirable element for the bid taker because it hinders competition; Roth and Ockenfels 2002).

When the supply base is asymmetric, rank information is insufficient for a bidder to determine whether she is winning or losing the auction at any given price. The reason is that if the cost supports of different bidders overlap, then a bidder whose bid is currently at rank 1 is not always able to determine whether her bid is winning. For example, consider an asymmetric supplier setting where $h = q_H = 100$, $c_H = 105$, and $c_L = 40$. Suppose that the H-type currently has a bid of 110 and the L-type bidder currently has a bid of 50. Under the full-price feedback, the L-type bidder knows she is losing the auction at her current bid of 50 because her quality-adjusted bid of 50 is above the H-type's quality-adjusted bid of $110 - 100 = 10$. Therefore, she has the dominant strategy to continue to bid down, and the H-type bidder has the dominant strategy not to decrease her bid any further (because she is currently winning the auction). In contrast, under the rank feedback mechanism, the L-type would be in rank 1 (based on her price bid) and would not be able to determine whether she was currently winning or losing the auction. Similarly, the H-type bidder would not know for certain that she is winning the auction; if her opponent's bid is greater than 10, then she is the current winner, but rank feedback does not allow the H-type bidder to ascertain her opponent's exact bid. In summary, under full-price feedback, suppliers can perfectly assess whether they are winning, and if not, by how much they must lower their current bid in order to win the auction. In contrast, the minimum bid amount necessary to win the contract may be impossible to calculate under a rank format because a bidder is unable to see her opponent's bid.

As a first step toward gaining analytical insights on rank feedback and its specific impact on procurement costs, we will analyze two cost/quality settings. In our base model (§4), $h = 100$; i.e., the cost support of c_H and c_L do not overlap. In our extensions (§6), we consider a setting in which $h = 10$; that is, the cost distributions significantly overlap.

4. Nonoverlapping Cost Support

In this section we present the equilibrium bidding strategies for both full and rank feedback when $h = 100$. Recall that the buyer may either face a symmetric supplier base (two H-types or two L-types) or an asymmetric supplier base. We start the analysis with the following remarks:

Remark 1. For an L-type it is never optimal to stop bidding above her cost while being in rank 2.

Because an L-type who is in rank 2 is certain to lose the auction regardless of her opponent's type, she should decrease her bid.

Remark 2. An L-type bidder will never stop bidding at a level above 100.

To see why, suppose an L-type bidder is considering stopping at a bid above 100 and allowing the auction to end. This L-type bidder must be in rank 1, for otherwise she would certainly lose the auction, regardless of the type of her opponent (see Remark 1). If the opponent is H-type, the L-type bidder will lose for certain; if, however, the opponent is an L-type, one of the L-types will find itself in a rank of 2 and not allow the auction to end with a bid above 100 (Remark 1).

Remark 3. An H-type will never bid below 100.

Finally, we assume that bidders will always reduce their bids by the smallest allowable bid decrement; this bid strategy never exposes the bidder to the possibility of regret of overshooting that can occur under jump bidding. We next characterize the equilibrium bidding strategies under full and rank feedback.

Claim 1. *Under full-price feedback, bidders know whether they are winning the auction at any given price, irrespective of the supplier base composition (symmetric or asymmetric). Therefore, losing bidders have the weakly dominant strategy of continuing to bid down until they either bid down to their costs or become winning bidders.*

It is interesting to note that the only difference between the full-price feedback setting and the standard auction without quality consideration is that in the asymmetric setting, both bidder types have to adjust down their own price bid by 100 to determine whether they are currently winning.

Under a rank format, bidders receive feedback on the rank of their bid but cannot see their opponent's bid. The equilibrium bidding behavior depends on whether a bidder is in a symmetric or asymmetric market setting.

Claim 2. *Under rank feedback, if a bidder is in a symmetric market setting, i.e., facing an opponent of same type, the bidding strategy under rank feedback reduces to that in a full feedback auction. If a bidder is in an asymmetric market setting, bidding in a rank feedback auction is strategically equivalent to a first-price sealed-bid auction.*

In a *symmetric* market setting, a bidder knows that if she is in rank 1 she is winning and if she is in rank 2 she is losing. So whenever a bidder is in rank 2, she should bid down, one bid decrement at a time, until either her rank changes to 1 or she bids down to her cost. In an *asymmetric* setting, a bidder's rank will never change because the cost supports do not overlap: an H-type is always in rank 2 and an L-type is always in rank 1. This makes the bidding problem equivalent to bidding in a first-price sealed-bid auction because each bidder will be paid her bid if she should win. The corresponding equilibrium bids for an H-type with cost c_H and an L-type with cost c_L

for the *asymmetric* market setting are (see Appendix A for the calculations)

$$b_H = \frac{c_H + 200}{2} \quad \text{and} \quad b_L = \frac{c_L + 100}{2}. \tag{1}$$

Vickrey (1961) shows that there is expected revenue equivalence in forward auctions between the sealed-bid first-price format (SB) and the open-ascending (English) format. This equivalence, but now with respect to the expected buyer cost, also holds in reverse auctions. By construction, our setting is equivalent to the standard auction setting in which suppliers have independent and identically distributed quality-adjusted costs. As a result, the average quality-adjusted buyer costs in our setting are theoretically the same in the full-price feedback format as in the sealed-bid first-price format.

Proposition 1. *When bidders are risk-neutral expected-profit maximizers, the buyer's expected quality-adjusted cost is the same under full and rank feedback formats, when bidders know the type of their competitor with certainty and the cost support of different type bidders does not overlap. Furthermore, this is the same expected quality-adjusted cost the buyer would expected under a first-price sealed-bid auction.*

Proof. The bidding strategies under rank and full-price feedback are identical under symmetric market settings. Furthermore, the bidding strategy under rank feedback in an asymmetric setting is equivalent to the bidding strategy under the sealed-bid format. Finally, there is expected quality-adjusted cost equivalence between the sealed-bid and the full feedback formats. □

The above proof hinges on the critical result that there is expected quality-adjusted cost equivalence between the sealed-bid and the full feedback formats. This equivalence is a direct consequence of the risk-neutral expected-profit maximization assumption, whereas the equivalence between full and rank based feedback in symmetric markets is a direct consequence of the assumption of incremental bidding in auctions with rank feedback, when market settings are known to be symmetric.

If we relax the expected-profit maximization assumption, then we have a different ranking of the expected quality-adjusted cost for the three auction formats:

Proposition 2. *When bidders are risk averse or regret averse, the buyer's expected quality-adjusted cost is higher under full than under rank feedback formats in asymmetric markets, when bidders know the type of their competitor with certainty and the cost support of different type bidders does not overlap. Furthermore, the expected quality-adjusted cost of the buyer in a first-price sealed-bid auction is lower than in an auction with full feedback.*

Proof. The bidding strategy under rank feedback in an asymmetric setting is equivalent to the bidding strategy under the sealed-bid format. Holt (1980) showed that risk-averse bidders bid more aggressively in sealed-bid first-price reverse auctions than in open-bid reverse auctions with full-price feedback. Engelbrecht-Wiggans and Katok (2007) showed that there is a formulation of regret-sensitive bidders that is mathematically equivalent to constant relative risk aversion. Therefore, either risk aversion or bidder regret implies that average quality-adjusted buyer cost will be lower in asymmetric markets in auctions with rank feedback than in auctions with full-price feedback. □

Theory provides very clear predictions for the full and rank feedback as well as sealed-bid formats, and different behavioral assumptions lead to different predictions in asymmetric markets. We next describe our main laboratory experiment that tests our theoretical predictions and the assumptions that underlie them.

5. Experimental Design, Hypotheses, and Results

5.1. Experimental Design and Protocol

We designed our initial experiment to follow the simplified model in the previous section. Two suppliers compete to provide a contract to a computerized buyer. Suppliers are randomly matched in groups of two and compete in 30 auctions. Each supplier has the same type (H or L) for the entire session. Half of the suppliers in each session are L-types and half are H-types. The distribution of cost c_i of supplier i depends on his type. For L-type suppliers costs are $c_L \sim U(0, 100)$ (rounded to the nearest integer) and their quality adjustment factor is $q_L = 0$. The costs of H-types suppliers are $c_H \sim U(100, 200)$ (rounded to the nearest integer) and their quality adjustment factor is $q_H = 100$. The information regarding the cost distributions and the quality adjustment factors is known by all suppliers. In all formats, the bidder with the lowest quality-adjusted bid at the end of the auction, $b_i - q_i$, wins the auction and earns $b_i - c_i$; in the event of a tie, the supplier who submitted the bid first wins.

Participants were students, mostly undergraduates, from a variety of majors. We recruited them using a computerized recruitment system, and earning cash was the only incentive we offered. Participants arrived at the computer laboratory at a prespecified date and time, were seated at computer terminals, and asked to read written instructions. After all participants had a chance to read the instructions, we read the instructions to them aloud to ensure common knowledge about the rules of the game. After we read the instructions we started the auctions. The experiment software used in all sessions was z-Tree (Fischbacher 2007). At the beginning of each auction each supplier received a new cost draw from the appropriate distribution. The auction lasted 60 seconds, during which time bidders could enter bids and observe feedback. The software did not accept bids above the reservation price but did not put any restrictions on bids below cost. After the first 60 seconds, if no new bids were placed for 10 seconds, the auction ended. But as long as new bids continued to be made, the auction continued. After the auction ended both bidders learned the auction outcome, including the winner, the profits, and the buyer quality-adjusted cost. After the 30 auctions were completed, earnings from all auction were converted to U.S. dollars at a preannounced conversion rate and paid to participants in private and in cash. All sessions were conducted at the Laboratory of Economic Management and Auctions at the Penn State University's Smeal College of Business in 2008 and 2009. Average earnings, including a $5 participation fee, were $25 per participant. Sessions lasted an hour and half on average. The instructions, which contain snapshots of the screen seen by the suppliers, can be found in the online supplement.

We also conducted a baseline treatment with a sealed-bid first-price auction in which each bidder simply submitted a single bid. The rest of the procedure was the same. In total, the initial study included three treatments. The full feedback treatment (Full) included 10 sessions, the rank feedback treatment (Rank) included 8 sessions, and the sealed-bid (SB) treatment included 4 sessions. The second study included three additional treatments, all with rank feedback: one in which bidders were not informed of the type of their competitor, one in which the H-type and L-type bidders had overlapping cost supports, and the third in which the cost supports overlapped and the type of competitors was unknown. Each of those treatments included four sessions. Each session included 8 participants for a total of 272 participants. Quality-adjusted cost draws and the matching protocol (how participants were matched in pairs each round) were identical in each session.[1]

[1] Of the 10 sessions in the Full treatment, we conducted 6 with the uniform reservation price of 200 and the other 4 with the individual reservation price of 200 for H-type bidders and 100 for L-type bidders. Of the eight Rank sessions, four used uniform reserve and the other four used individual reserves. Reservation prices had no effect on the bidding behavior or the auction outcomes, so we pooled the data with uniform and individual reserve for the analysis we report in the paper. We also repeated all analysis using only uniform reserve data, and none of the results were affected. All other treatments were conducted with uniform reserve.

5.2. Research Hypotheses

In this section we formulate three research hypotheses that reflect our theoretical predictions as well as predict auction performance if our theory's underlying assumptions are relaxed.

The first hypothesis deals with average quality-adjusted buyer cost under the assumption that bidders are risk-neutral expected-profit maximizers.

Hypothesis 1 (Buyer Costs Under Risk Neutrality). *Average quality-adjusted buyer cost under Full, Rank, and SB treatments should be the same and not significantly different from* 67.67:

$$Full = Rank = SB \approx 67.67.$$

Given the actual draws in our experiment, theoretical predictions under decremental bidding and under the sealed bid strategy are slightly different from 67.67 (see numbers in square brackets in Table 1).

The fact that expected quality-adjusted costs should be approximately the same in all three treatments critically hinges on the observation that expected quality-adjusted cost is the same in the open-bid and sealed-bid first-price auction. However, as we mentioned in §2, experimental economics literature provides a great deal of evidence that in fact bidding in sealed-bid first-price auctions is generally more competitive than in open-bid auctions (see Kagel and Roth 1995 for a review of the experimental economics literature on auctions). Of course if the sealed-bid first-price bidding is more competitive in our setting, the equality in Hypothesis 1 will not hold.

If we relax the assumption that bidders are merely maximizing their expected profit and are instead, for example, either risk averse or sensitive to regret, then the sealed-bid first-price bidding strategy should result in lower quality-adjusted prices. Note that although bidders in the Full treatment can bid incrementally in the quality-adjusted space and never have to rely on the sealed-bid first-price strategy, bidders in the Rank treatment should follow the sealed-bid first-price strategy in asymmetric auctions, but not in symmetric auctions.

Hypothesis 2 (Buyer Costs Without Risk Neutrality). *In asymmetric markets, average quality-adjusted costs should be higher in Full than in Rank and not different in Rank than in SB*:

$$Full > Rank = SB.$$

A behavioral assumption that affects performance in symmetric markets is that bidders follow the incremental bidding strategy—they lower their bids one bid decrement at a time. The incremental approach is reasonable under full feedback, where bidders can perfectly gauge the minimum amount by which they must decrease their bid in order to move into a winning position. The incremental approach is a stronger assumption under rank feedback. If we relax the incremental bidding assumption in auctions with rank feedback, the average quality-adjusted costs in symmetric auctions will no longer be the same as in auctions with full feedback. Bidders in auctions with rank feedback may fail to follow incremental bidding strategy due to overreaction and impatience because in this setting the effect of one's action is not immediately observed. It has been shown (see, for example, Sterman 1989) that when people are unable to observe the effect of their actions, they tend to overreact.[2] When there is a delay between the action and its outcome, decision makers are impatient to see the result of their action and sometimes overreact by performing the action again. In auctions with rank-based feedbacks, bidders are exactly in this situation—they lower their bid, but if their rank does not change as a result, they may become impatient.

There is also evidence of bidder impatience in auctions: Katok and Kwasnica (2008) show that prices in slow Dutch (forward-descending clock) auctions are higher than in fast Dutch auctions because impatient bidders stop the clock too soon when this clock is slow. Kwasnica and Katok (2007) show that bidders are also impatient in forward-ascending auctions because they place jump bids and tend to place more and larger jump bids when this behavior results in the ability to complete more auctions.

Our third hypothesis deals with the effect of bidder impatience on the behavior in symmetric auctions with rank feedback. When two bidders of the same type are competing in an auction with rank feedback, the best they can do is to bid incrementally. This incremental bidding would result in bidding down exactly to the higher of the two costs, and the winner of the auction will be paid one bid increment less than this amount. However, impatient bidders do not bid incrementally—they place jump bids. With full-price feedback, jump bidding does not affect the outcome of the auction significantly (Kwasnica and Katok 2007 show this). Bid decrements tend to decrease as the auction price approaches the final price, and because bidders know the current low price, auction winners do not "jump over" the opponent's final bid. But under rank feedback, bidders do not know how much lower they need to bid in order to outbid the competitor. They may overreact and increase their jump bidding decrement when they do not observe a change in their rank. We formulate the next hypothesis to formally express how impatience affects behavior in symmetric auctions under rank feedback.

[2] Sterman (1989) analyzes the bullwhip effect in the beergame, but his point is broader, and he applies it to a much wider set of settings (see Sterman 2000).

Hypothesis 3 (Buyer Costs With Jump Bidding). *In symmetric markets, average quality-adjusted buyer cost should be higher in Full than in Rank*:

$$Full > Rank.$$

Average bid increments in symmetric auctions should be higher in Rank than in Full.

5.3. Results of Experiments

We begin by comparing average quality-adjusted buyer costs to theoretical benchmarks under risk neutrality.

Table 1 summarizes average prices and their standard deviations, computed using a session as a unit of analysis in all three treatments in the experiment, as well as theoretical risk-neutral predictions. Figure 1 compares cumulative distributions of quality-adjusted costs in symmetric auctions (a) and asymmetric auctions (b), in the three treatments (Full, Rank, and SB). We present average quality-adjusted costs and bid decrements by session in Appendix D.

According to the Wilcoxon Rank-Sum Test (here and in subsequent analysis we use the session average as a unit of analysis), our first result largely rejects Hypothesis 1.

Result 1. Average quality-adjusted buyer costs in the Full treatment are not significantly different from theoretical predictions, but in the SB and Rank treatments they are significantly lower than predicted.

We use the Wilcoxon Signed-Rank Test to compare average quality-adjusted prices in symmetric and asymmetric markets separately. Average quality-adjusted costs in the Full treatment are significantly higher than they are in the SB treatment ($p = 0.0047$). In asymmetric auctions, quality-adjusted costs are lower in the Rank treatment than in the Full treatment ($p = 0.0004$), and the Rank costs are not different

Table 1 Summary of the Average Prices, Bid Decrements (Standard Deviations in Parentheses) and Theoretical Predictions [in Square Brackets]

	Prices			Bid decrements	
Treatment	Overall	Symmetric	Asymmetric	Symmetric	Asymmetric
Full	67.11	69.47	65.65	6.35**	5.82**
No. of obs. = 10	(3.06)	(3.93)	(3.07)	(1.26)	(1.14)
	[67.60]	[71.24]	[65.34]	[1]	[1]
Rank	58.73**	63.68**	55.66**	12.61**	21.86
No. of obs. = 8	(3.59)	(3.03)	(4.62)	(3.59)	(4.65)
	[69.04]	[71.24]	[67.67]	[1]	[≫ 1]
SB	57.49*	58.21*	57.04*		
	(2.20)	(2.86)	(1.94)	N/A	N/A
No. of obs. = 4	[68.26]	[69.20]	[67.67]		

Note. H_o: Data = theoretical prediction.
$^*p < 0.10$; $^{**}p < 0.05$ (two-sided).

Figure 1 Comparing Cumulative Distributions of Quality-Adjusted Prices

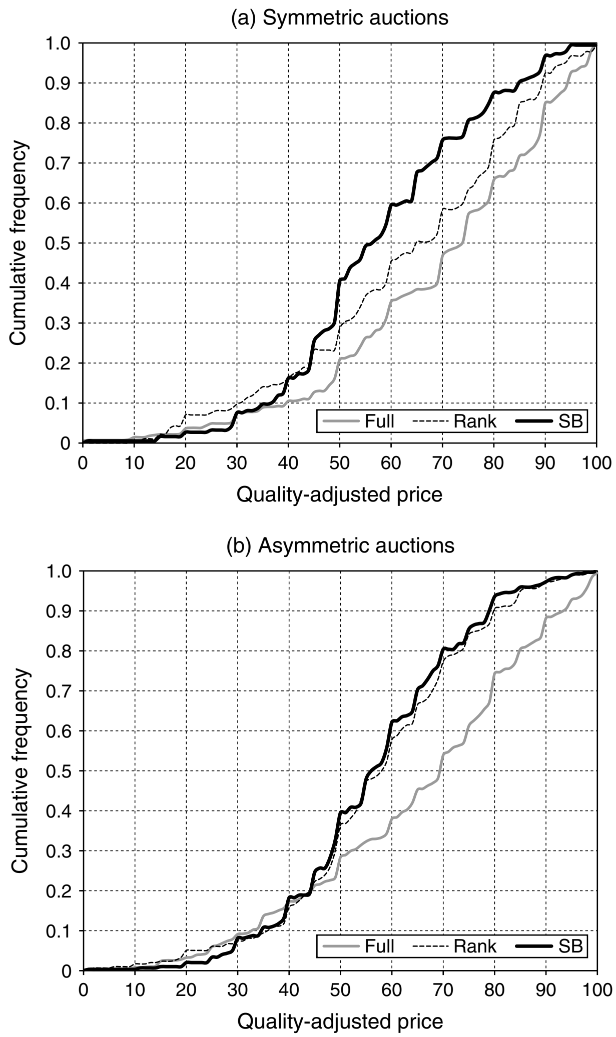

Figure 2 Comparison of Average Bid Decrement Across Treatments in Symmetric Auctions

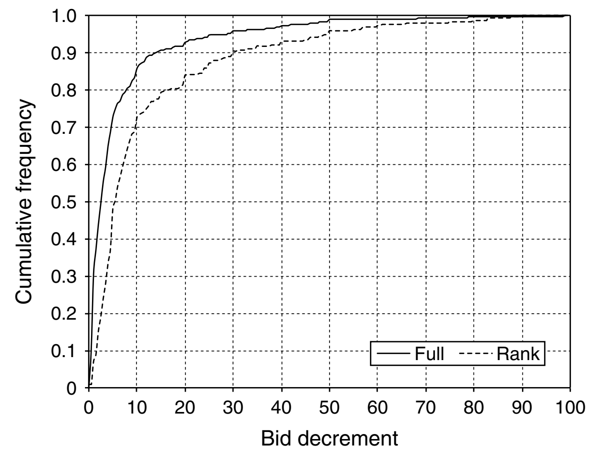

from the SB prices ($p = 0.7341$). In symmetric auctions, quality-adjusted costs are lower in the Rank treatment than in the Full treatment ($p = 0.0059$), and average quality-adjusted costs are higher in the Rank treatment than in the SB treatment ($p = 0.0174$). Therefore, the next two results are consistent with Hypotheses 2 and 3.

Result 2. In asymmetric markets, average quality-adjusted buyer costs in the Full treatment are higher than in the Rank treatment.

Result 3. In symmetric markets, average quality-adjusted buyer costs in the Full treatment are higher than in the Rank treatment.

To check for any learning effects, we also replicated the above analysis excluding the first 10 rounds. In early rounds we observed some obvious errors, such as bids below cost, that largely disappeared in later rounds. Surprisingly, under full feedback, there are the most bids below cost, and they decrease with experience but do not completely go away. Dropping the early rounds made no substantive difference to our conclusions, so we report the analysis based on the full sample. Losses were deducted from participants' earnings; although a few participants lost some money in a few of the rounds, none came close to losing money overall.

Figure 2 compares the distributions of bid decrements in symmetric auctions. It is clear from the figure that, consistent with Hypothesis 3, bid decrements in the Full treatment are much smaller than in the Rank treatment. About 85% of bid decrements in the Full treatment are less than 10, but in the Rank treatment 85% of bid decrements are less than 25. Average bid decrement is smaller in the Full treatment than in the Rank treatment ($p = 0.0004$).

Our results complement the findings of Isaac et al. (2005, 2007). The authors study both real-world data from the FCC (Federal Communications Commission) spectrum auctions (Isaac et al. 2007), as well as from the lab experiments (Isaac et al. 2005); they find that auction formats that allow for jump bids can help increase the bid taker's revenue (alternatively, decrease her procurement costs). In addition, their data suggest that jump bidding arises as a result of bidder impatience rather than effort by bidders to deter competition (signaling). Easley and Tenorio (2004) show that jump bidding can be used strategically to signal high valuation and prevent entry and find evidence of this behavior in Internet auctions, whereas Kwasnica and Katok (2007) rule out signaling as the cause of jump bidding in their laboratory experiment.

6. Extensions: Overlapping Costs and Incomplete Information

The findings we presented thus far are for a setting that we intentionally designed to be parsimonious in order to have tractable theoretical benchmarks and to be able to clearly identify and separate causes of behavioral phenomena. In this section we test the robustness of our findings by considering a setting in which the cost support of two bidder types overlaps as well as a setting in which bidders are not told the type of their competitor (so they do not know whether they are bidding in a symmetric or an asymmetric auction).

6.1. Rank Feedback Under Overlapping Costs

In the previous analysis, the costs c_H and c_L of the H-type and L-type, respectively, do not overlap. For this (extreme) setting, we found that rank feedback results in significantly lower quality-adjusted buyer cost than does the full-price feedback, despite theoretical predictions to the contrary. It is natural to ask whether the performance of rank feedback would suffer if the cost dispersion between H and L types were narrowed (the other extreme case of $h = 0$ would imply that H and L types are equivalent, and hence the buyer would never face an asymmetric setting). In this subsection we explore the performance of rank feedback in settings with overlapping supplier costs; namely we assume $0 < h < 100$. (In the experimental section, we will test the setting with $h = 10$; i.e., $c_H \sim U[10, 110]$ and $c_L \sim U[0, 100]$.)

Using the same arguments as were used for $h = 100$, we know that in symmetric supplier settings, it is a dominant strategy for a supplier to bid down to her cost regardless of the feedback form; hence, average quality-adjusted prices should be the same with the rank and full-price feedback. But the behavior in asymmetric auctions is affected by h—when an L-type is in rank 1 and an H-type is in rank 2, both bidders may potentially have more information about the competitor's cost support than they would have had when $h = 100$. In particular, an L-type can update her beliefs over the H-type opponent's cost as their bids move below 100. An H-type in rank 2 must trade off the desire to decrease his bid (and increase his chances of winning) with the desire to limit this information flow to his competitor (combined with the possibility that he is currently winning with his rank 2 bid).

A consequence of this trade-off is that the bidders do not have a dominant bidding strategy, thereby complicating the search for the equilibrium strategies. Therefore, we step back and calculate myopic best-response (MBR) strategies: The MBR approach is used to determine the *threshold* bid, i.e., the lowest bid that an H-type or L-type with costs c_H and c_L, respectively, is willing to submit given their current information set (namely, whether they face an H or L opponent, their current rank, and the last bid at which the rank changed). The approach of using MBR bidding strategies in settings that do not lend themselves to equilibrium analysis has been used successfully by Parkes and Ungar (2000) and Parkes and Kalagnanam (2005) in analyzing the design of combinatorial auctions and multiattribute procurement auctions, respectively. In fact, Rothkopf (2007) argues that this approach may well be better than equilibrium analysis for modeling bidding strategies in complex real-world auctions.

The full details of our analysis of the MBR strategies for auctions with type information can be found in Appendix B. We use these MBR strategies to simulate the buyer's expected cost; the goal of doing this is to establish if we should expect costs to be lower, the same, or higher under rank with $h \in (0, 100)$ than full feedback. We present results of this simulation in Figure 3—we find that, in general, the costs under rank with $h \in (0, 100)$ should be higher than or equal to the buyer's expected cost under full feedback.

6.2. Rank Feedback Under Unknown Opponent Identity

Up until this point, we have assumed that suppliers know the identity of their opponent. When bidders do not know the type of their competitor, identifying optimal bidding strategies is even more complex because bidders do not always have a dominant bidding strategy. Whereas an H-type in rank 1 should stop bidding, for she is guaranteed to win in both a symmetric and asymmetric market setting, and an L-type in rank 2 should decrease her bid until she is in rank 1 or reaches her cost, the bidding strategy is less clear for an L-type in rank 1 or an H-type in rank 2. For example, an L-type in rank 1 may win (and hence should stop lowering her bid), but only *if* her opponent is an L-type or if her opponent is an H-type *and* the H-type's bid is h units greater than her own. However, the possibility that she may lose against an H-type whose bid is less than h away from hers implies that she would be better off decreasing her bid—and hence an L-type in rank 1 no longer has a dominant bidding strategy. A similar story holds for an H-type in rank 2—she may lose at her current bid if her opponent is another H-type, or she may win *if* her opponent is an L-type and his bid is less than h below hers. Consequently, identifying the equilibrium in these settings becomes an intractable problem.

As above, we identify MBR strategies for the bidders; we then analyze the rank feedback mechanism under the assumption that bidders follow the MBR strategies. As opposed to the previous analysis, best response strategies cannot depend on the state of the market (asymmetric or symmetric). Hence, bidders must determine how low they are willing to bid, i.e., their *threshold* bid, based only on their current rank, the last bid at which rank changed, and their beliefs

Figure 3 Average Cost for the Buyer from 100,000 Simulations of the Myopic Best-Reply Bidding Strategies

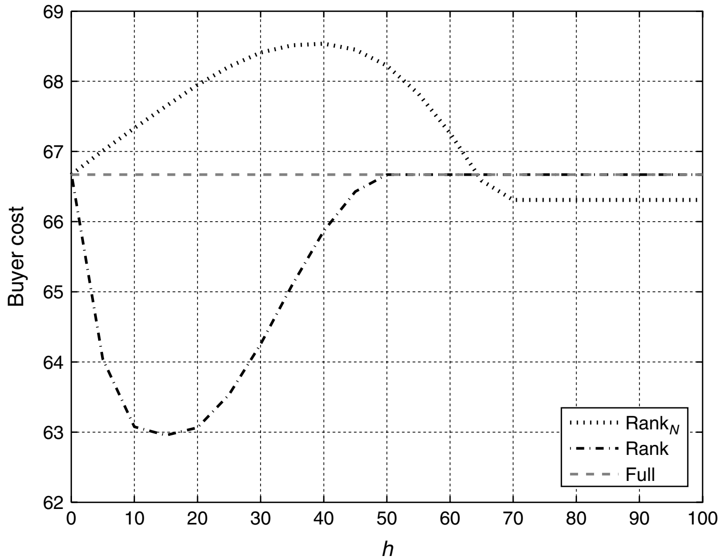

h	Rank$_N$	Rank	Full
0	66.67	66.69	66.67
5	67.01	64.04	66.67
10	67.33	63.08	66.67
15	67.65	62.95	66.67
20	67.95	63.06	66.67
25	68.21	63.53	66.67
30	68.41	64.25	66.67
35	68.51	65.10	66.67
40	68.54	65.87	66.67
45	68.45	66.43	66.67
50	68.22	66.67	66.67
55	67.82	66.67	66.67
60	67.26	66.67	66.67
65	66.59	66.67	66.67
70	66.31	66.67	66.67
75	66.31	66.67	66.67
80	66.31	66.67	66.67
85	66.31	66.67	66.67
90	66.31	66.67	66.67
95	66.31	66.67	66.67
100	66.31	66.67	66.67

over their opponent's type. To evaluate the robustness of the performance of rank feedback when opponent types are unknown, we assume that suppliers know their own type (and cost), know h ($0 < h \leq 100$), and believe their opponent to be an H-type with probability p.

The full details of our analysis of the MBR strategies for auctions without type information can be found in Appendix C.

To give the reader the sense of the effect of cost overlap and type information on the buyer's expected procurement cost, we conducted a simulation with 100,000 replications. In the simulation, the initial probability that the competitor is an H-type is $p = \frac{1}{2}$ and the bid decrement is $\delta = 0.001$. In the setting in which bidders do not know the type of their competitor (Rank$_N$), when auctions are asymmetric, and the threshold of the H-type is above 100, we assume there is no change in the rank, so the L-type bidder bids down to his sealed bid and the H-type bidder bids down to a threshold that is a best reply to the L-type's behavior.

Figure 3 shows the results for the average quality-adjusted cost for the buyer for different values of h; where *Full* refers to a full-price feedback setting with known opponent types, *Rank* refers to a rank feedback setting with known opponent types, and $Rank_N$ refers to a setting with rank feedback and unknown opponent types. Note that in the *Full* condition, the strategy for both types of bidders is to bid down in minimum decrements until they are either winning—in quality-adjusted space—or they reach their cost, regardless of the value of h.

Figure 3 illustrates the following general principles regarding the expected costs of the three auction formats: (i) $Rank < Full < Rank_N$ for $h \leq 50$; (ii) $Rank = Full < Rank_N$ for $50 < h \leq 66\frac{2}{3}$; and (iii) $Rank = Full \approx Rank_N$ for $h > 66\frac{2}{3}$.

6.3. Experimental Results

Next, we report results on three additional experiments that we designed to examine the effect of overlapping cost supports and unknown competitor type on auctions with rank feedback. We first consider the setting where $h = 10$, which we label Rank10 treatment. Our simulation indicates that we should expect the average buyer cost to be significantly lower than in our baseline $h = 100$ treatment presented in the previous section (labeled simply Rank). In the treatment that we label Rank$_N$, we keep $h = 100$ but do not tell bidders the type of their competitor. Finally, in the treatment we label Rank10$_N$, we set $h = 10$ and do not provide opponent's type information. We compare the three new treatments to the Rank and Full treatments in the previous section but do not conduct additional experiments with full-price feedback and overlapping cost support. We refer interested readers to Haruvy and Katok (2010), who specifically examine the effect of providing competitor type information in auctions with overlapping cost support and full-price feedback. They find that when bidders do not know their competitors' type, full-price feedback does not provide sufficient information for bidders to determine their winning status. As a consequence of not knowing their winning status, bidders bid more competitively than they do in auctions in which they do know their winning status. Thus, there is a sealed-bid effect in auctions with full-price feedback, unknown competitor type, and overlapping cost support which is similar to the sealed-bid effect that we document in asymmetric auctions with rank feedback.

Table 2 Summary of the Average Prices, Bid Decrements (Standard Deviations in Parentheses), and MBR Predictions [in Square Brackets] for the Three Additional Rank Treatments

	Prices		Bid decrements	
Treatment	Symmetric	Asymmetric	Symmetric	Asymmetric
Rank10	60.94*	59.11	10.67**	13.16
No. of obs. = 4	(8.88)	(2.97)	(1.17)	(2.09)
	[71.24]	[60.04]	[1]	[1]
Rank_N	62.55*	55.42*	13.46**	15.59
No. of obs. = 4	(4.83)	(2.92)	(1.42)	(1.87)
	[76.32]	[61.39]	[1]	[1]
Rank10_N	64.15*	55.37*	11.94**	12.43
No. of obs. = 4	(3.12)	(2.68)	(2.50)	(4.57)
	[73.41]	[65.04]	[1]	[1]

Note. H_0: Data = theoretical prediction.
$^*p < 0.10$; $^{**}p < 0.05$.

We conducted four independent sessions for each of these three additional treatments. As before, each session included eight participants who were randomly re-matched to compete in 30 auctions. We continue to use the session as a unit of analysis. Table 2 presents the average prices and average bid decrements in the three new treatments.

We start by making three observations about average prices in the three new treatments. First, in symmetric auctions, average prices are lower than in the Full treatment ($p = 0.0339$ for Rank10, $p = 0.0897$ for Rank_N, and $p = 0.0401$ for Rank10_N) and mostly higher (although not always significantly so) than in the SB treatment ($p = 0.2482$ for Rank10, $p = 0.1489$ for Rank_N, and $p = 0.0209$ for Rank10_N). In asymmetric auctions, average prices are all significantly below Full prices ($p < 0.01$ for all comparisons) and not significantly different from SB prices ($p > 0.1$ for all comparisons). Qualitatively, *the effect of rank feedback on average prices does not change when costs overlap and when bidders are unsure about the type of their competitor.*

Second, in symmetric auctions, average prices in the three new treatments are lower than theoretical predictions, a similar ordering to our Rank treatment. In asymmetric auctions average prices are also lower than theoretical predictions (again, as they are in the Rank treatment), with the exception of the Rank10 treatment, in which the MBR strategy itself predicts significantly lower prices.

Third, average prices in the three new treatments are not different from average prices in the Rank treatments, either for symmetric or for asymmetric auctions ($p \gg 0.1$ for all six comparisons). Hence, we can conclude that the main result from our baseline study continues to hold: *rank feedback results in lower prices than does full-price feedback with known type information.* It is worth emphasizing that in more complex (and less tractable) settings, such as a setting with overlapping cost support and no information about the competitor's type, rank feedback may not lower prices relative to full-price feedback. The reason has to do with the sealed-bid effect that Haruvy and Katok (2010) report to be present in auctions with full-price feedback in those complex settings.

Two additional points are worth noting: First, bidders in symmetric auctions in the Rank10 treatment have the dominant bidding strategy, just as they do in the Rank treatment. Lower prices in Rank10 treatment than in the Full treatment are consistent with bidder impatience. That average bid decrements are significantly larger in the Rank10 treatment than in the Full treatment ($p = 0.0047$) but are not different from average bid decrements in the Rank treatment ($p = 0.3082$) again suggests that bidder impatience plays a role and is not affected by cost overlap. Second, average bid decrements in asymmetric auctions are significantly smaller in the Rank10 treatment than in the Rank treatment ($p = 0.0066$), suggesting differences in bidding behavior. With no cost overlap (Rank treatment), bidders in asymmetric auctions cannot expect to gather any meaningful information from bidding feedback and thus should follow the sealed-bid bidding strategy—resulting in very large average bid decrements. But in the presence of significant cost overlap (Rank10 treatment), feedback information is more useful, tempering bidder impatience and yielding a more incremental approach to bidding.

7. Conclusions

Auctions are widely used in procurement because they utilize supplier competition to decrease costs. Procurement auction designers must make a number of design choices, and these choices affect auction performance, often in unexpected ways. One auction format that has recently became popular is the so-called rank feedback auction. Bidders compete on price during a live auction event, but instead of knowing the current bids of their competitors, they are only told the rank of their bid. Ours is the first paper to analyze the effect of rank feedback and note the theoretical similarity between auctions with rank feedback in which bidders have different nonprice attributes and sealed-bid auctions.

We analyze a setting in which bidders differ in quality and the buyer awards the contract to the bidder who submitted the lowest quality-adjusted bid.

We present a simple model for a setting with two types of bidders who have nonoverlapping cost supports and show that when bidders follow the risk-neutral Nash equilibrium in this setting, the average quality-adjusted buyer costs should be the same under the rank feedback as under the full-price feedback. We then extend our model to a setting in which cost support can overlap, but the quality-adjusted cost distributions continue to be identical. We further extend the model to the case in which bidders are not aware of the type of their competitor and must update their beliefs based on auction feedback. We analyze those two extensions by assuming that bidders follow the myopic best-reply bidding strategy and find that in theory large cost overlaps result in lower average quality-adjusted buyer costs when bidders know the type of their competitor, but higher quality-adjusted buyer costs when bidders do not know the type of their competitor.

We test empirical predictions of our models in a controlled laboratory setting. Our laboratory experiment tests the extent to which subjects understand how the feedback mechanism converts the incentives in a dynamic auction to those of a sealed-bid first-price auction. Additionally, our laboratory experiment tests whether there are some unforeseen behavioral elements that cause the mechanism to deviate from the predictions of the sealed bid mechanism. We find that, contrary to the theory, average quality-adjusted prices are lower under rank feedback than under full-price feedback, regardless of the amount of cost overlap or the presence of type information. Interestingly, rank feedback lowers the buyer's cost under *all* supplier scenarios, e.g., whether the suppliers are of the same quality type or not.

We present two explanations for deviations between the analytical predictions and the experimental results. First, when bidders of different types compete in an auction (we call this the asymmetric setting), the rank feedback is uninformative when suppliers' costs do not overlap, and bidders bid down to the same amount as they would in a sealed-bid auction. In our study, bidders in sealed-bid auctions bid more competitively than in open-bid auctions, and bidders in auctions with rank feedback behave in a way that is similar to bidders in sealed-bid auctions. This sealed-bid effect has been documented in a number of studies in experimental economics (Kagel and Roth 1995) and can be explained by regret (Filiz-Ozbay and Ozbay 2007, Engelbrecht-Wiggans and Katok 2008) or by risk aversion (Cox et al. 1988). We are the first to document it in auctions with rank feedback.

But the sealed-bid effect does not fully explain our data. When bidders with the same quality compete, rank feedback is informative, and the bidding behavior should not differ from the bidding behavior in auctions with full-price feedback. However, we find that even when bidders have the same quality (and know that they have the same quality), auctions with rank feedback result in lower quality-adjusted prices. Bidder impatience is one way to explain these data. The main evidence of bidder impatience that we find and present in this study is the difference in the average size of bid decrement under the two feedback formats. Approximately 80% of auctions with full-price feedback have an average bid decrement that is less than 10, compared with 70% for rank auctions; furthermore, more than 20% of auctions with rank feedback have average bid decrements that are larger than 20. Of course the extent to which jump bidding might be due to impatience or signaling in real auctions outside the laboratory remains an important empirical question.

In conclusion, we find that rank feedback is another effective tool to promote jump bidding (and lower procurement costs). This knowledge, when combined with other studies on jump bidding (e.g., Haruvy and Popkowski Leszczyc 2010 and Raviv 2008),[3] can further assist bid takers in selecting the most cost effective auction format.

Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://msom.journal.informs.org/.

Acknowledgments

Wedad Elmaghraby gratefully acknowledges support from the Hewlett Packard Labs grant and the Robert H. Smith School of Business. Elena Katok and Natalia Santamaría gratefully acknowledge support from the Smeal College of Business at the Pennsylvania State University. Elena Katok gratefully acknowledges support from the National Science Foundation [Award 0849054].

Appendix A. Equilibrium Bids for Rank Feedback and High Cost Dispersion ($h = 100$)

Following the notation used in Krishna (2002), let $\Pi_i(b, c)$ be the expected profit of bidder i when she has a cost of c and submits a bid b, and let $\phi_i(b)$ denote the inverse of the equilibrium strategy (assuming it exists). The problem bidder i has to solve to calculate the equilibrium bids can be stated as follows:

$$\Pi_i(b, c) = (b - c)\text{Prob}(b_j > b) \tag{A1}$$

$$= (b - c)(1 - F_j(\phi_j(b))), \tag{A2}$$

$$\frac{d\Pi_i}{db} = 1 - F_j(\phi_j(b)) - bf_j(\phi_j(b))\phi_j'(b) + cf_j(\phi_j(b))\phi_j'(b) = 0, \tag{A3}$$

$$\phi_j'(b) = \frac{1 - F_j(\phi_j(b))}{f_j(\phi_j(b))}\frac{1}{(b - \phi_i(b))}. \tag{A4}$$

The solution of the system of differential equations defined above defines the equilibrium of a reverse "first-price" auction between two bidders. If bidders are symmetric, then

[3] Haruvy and Popkowski Leszczyc (2010) find that the occurrence of jump bidding is negatively correlated with the length of an auction. From a data set collected from open-outcry English car auctions, Raviv (2008) finds that jump bids are a function of the presale estimate of the vehicle's price.

$\phi_j'(b) = \phi_i'(b) = \phi'(b)$ and $F_j(\phi_j(b)) = F_i(\phi_i(b)) = F(\phi(b))$. For an H-type supplier, $F(\phi(b))$ corresponds to a Uniform distribution between 100 and 200. So the system we need to solve reduces to the following differential equation:

$$\phi'(b) = \frac{200 - \phi(b)}{b - \phi(b)}. \tag{A5}$$

The solution to (A5) (using the boundary condition $\phi'(200) = 200$) is as follows:

$$\phi(b) = b + \sqrt{b^2 - 400b + 40{,}000} \tag{A6}$$

$$= 2b - 200. \tag{A7}$$

From (A7) we get that the equilibrium strategy in this case is

$$b(c) = \frac{c + 200}{2}. \tag{A8}$$

In the case of an L-type supplier, $F(\phi(b))$ corresponds to a Uniform distribution between 0 and 100, and similar calculations yield the following equilibrium strategy:

$$b(c) = \frac{c + 100}{2}. \tag{A9}$$

Appendix B. Myopic Best-Reply Bidding Strategies for Rank Feedback Under Overlapping Costs

We know the best-reply bidding strategy of both high and low types if they are in a symmetric setting ((L, L) or (H, H)) is to bid down to cost. We also know that in the asymmetric setting (H, L), if the H-type is in rank 1 and the L-type is in rank 2, the H-type will stop bidding and the L-type must continue lowering his bid as long as his cost is h less than the current bid or he reaches his treshold. Therefore, the only setting that remains is when we are in an (H, L) setting, H is in rank 2, and L is in rank 1.

We are solving for the threshold bid of low and high types under such a setting. Define b_c to be the last bid at which the L-type was in rank 2, c_L to be the cost of the L-type, and c_H the cost of the H-type. To determine the thresholds we use a myopic best-reply approach and assume bidders bid down in decrements.

We start with the threshold for the H-type, which corresponds to the solution of the following problem:

$$\max_{b_H \geq c_H} (b_H - c_H)\text{Prob}(b_H - h < b_L \mid b_L < b), \tag{B1}$$

$$\max_{b_H \geq c_H} (b_H - c_H)(1 - \text{Prob}(b_L \leq b_H - h \mid b_L < b)), \tag{B2}$$

where b_L is the threshold bid of the L-type and where the H-type believes the L-type's bid to be uniformly distributed over $[\alpha_L, b]$. Note that when the H-type is bidding between $[100, 100 + h]$, he does not have additional information regarding the L-type's cost and therefore $b = 100$. The first-order condition yields the following solution:

$$\tilde{b}_H = \frac{c_H + b + h}{2} = \frac{c_H + 100 + h}{2}. \tag{B3}$$

Now we turn to the threshold for the L-type. When bidding starts, the L-type needs to go down from the bid ceiling because she cannot let the auction end while her bid is greater than 100, so she will go down to her own space. As long as the H-type threshold is between 100 and $100 + h$, the L-type will not have additional information regarding the H-type cost and therefore $b_c = 100 + h$. The threshold for the L-type corresponds to the solution to the following problem:

$$\max_{b_L \geq c_L} (b_L - c_L)\text{Prob}(b_H - h > b_L \mid b_H < b_c), \tag{B4}$$

$$\max_{b_L \geq c_L} (b_L - c_L)(1 - \text{Prob}(b_H \leq b_L + h \mid b_H < b_c)), \tag{B5}$$

where b_H is the threshold bid of the H-type and where the L-type believes the H-type's bid to be uniformly distributed over $[\alpha_H, b_c]$. The first-order condition gives

$$\tilde{b}_L = \frac{c_L + b_c - h}{2} = \frac{c_L + 100}{2}. \tag{B6}$$

Summarizing, we have the following results:

$$\tilde{b}_H = \begin{cases} \dfrac{c_H + b + h}{2} & \text{if } b < 100, \\[2ex] \dfrac{c_H + 100 + h}{2} & \text{if } b \geq 100; \end{cases}$$

$$\tilde{b}_L = \begin{cases} \max\left\{\dfrac{c_L + b_c - h}{2}, c_L\right\} & \text{if } b_c < 100, \\[2ex] \dfrac{c_L + 100}{2} & \text{if } b_c \geq 100. \end{cases} \tag{B7}$$

Note that when the type of the competitor is known, and the suppliers are facing an asymmetric setting, the H-type can avoid revealing information regarding her bid to the L-type. Because an L-type cannot let the auction end with a bid greater than 100, at the beginning of the auction she must start bidding down to her threshold immediately. In contrast, knowing her competitor is an L-type, the H-type does not have an incentive to bid down to her threshold immediately (even if she is in rank 2) and will let the L-type lead the way. When $\tilde{b}_H$ is between $[100, 100 + h]$, the L-type will not be able to gather information regarding the bid from the H-type, and the auction ends with both bidders at their respective thresholds, $\tilde{b}_H$ and $\tilde{b}_L$, corresponding to the first-price sealed bids. In this case, the L-type will reach her threshold being in rank 1 and the H-type will reach hers being in rank 2, and they will not see a change in their ranks.

In contrast, when $\tilde{b}_H$ is between $[h, 100]$, suppliers may see a change in their ranks. When $\tilde{b}_H < \tilde{b}_L$, there is guaranteed to be a change in rank at some bid level at or before reaching $\tilde{b}_L$; at that point both bidders will know the current bid of their opponent and will be able to bid down as if they were in a full feedback setting. When $\tilde{b}_L < \tilde{b}_H < 100$, the H-type may bid down to $\tilde{b}_H$ without experiencing a rank change; even if her rank does not change, an H-type's information set is changing as she decreases her bid because she will effectively have a new upper bound on the L-type's cost type (the value of b now corresponding to the H-type's "new" current bid). Hence, she will (possibly multiple times) update her threshold $\tilde{b}_H$ and adopt a strategy to bid down to it. Note that $\tilde{b}_H = b$, when $b = c_H + h$, so an H-type will not bid lower than $c_H + h$ for all $c_H \in [h, 100 + h]$. If when $\tilde{b}_L < \tilde{b}_H < 100$, the bidders do experience a change in their ranks, they will be able to identify their opponent's bid location and bid down as if they were in a full feedback setting.

Appendix C. Myopic Best-Reply Bidding Strategies for Rank Feedback Under Unknown Opponent Identity

When the bidders do not know the type of their competitor, an H-type in rank 1 should stop bidding, because no matter the type of her competitor she will win. An L-type in rank 2 should decrease her bid until she is in rank 1 or reaches her cost, otherwise she will lose for sure, and an L-type in rank 1 should decrease her bid only if her current bid minus h is greater than her cost; note that if her competitor is another L-type she does not need to decrease her bid any further to win, and if her competitor is an H-type, she has a chance of winning only if her current bid minus h is greater than her cost.

To determine the threshold bids for both types, we use myopic best replies and assume bidders bid down in decrements.

The threshold for the H-type depends on whether she is currently bidding above 100 or below. If she is bidding above 100, then she has no information regarding the bid of a potential L-type competitor. If she is bidding below 100, then her current bid is an upper bound of the bid of a potential L-type competitor.

We start with the case where the current bid of the H-type, b, is between h and 100. Let b_L the threshold bid of an L-type competitor, b'_H the threshold bid of an H-type competitor, p_H the conditional probability the competitor is an H-type given b, p_L the conditional probability the competitor is an L-type given b, and p the probability the competitor is an H-type. With those elements, the problem can be formulated as follows:

$$\max_{b_H \geq c_H} (b_H - c_H) \left(\overbrace{p_L \text{Prob}(b_H - h < b_L \mid b_L \leq b)}^{\text{winning versus an } L\text{-type}} + \underbrace{p_H \times \text{Prob}(b_H < b'_H \mid b'_H \leq b)}_{\text{winning versus an } H\text{-type}} \right), \quad \text{(C1)}$$

$$\max_{b_H \geq c_H} (b_H - c_H)\big(p_L(1 - \text{Prob}(b_L \leq b_H - h \mid b_L \leq b)) + p_H(1 - \text{Prob}(b_H \leq b'_H \mid b'_H \leq b))\big), \quad \text{(C2)}$$

$$\max_{b_H \geq c_H} (b_H - c_H)\left(p_L\left(\frac{b - b_H + h}{b - \alpha'_L}\right) + p_H\left(\frac{b - b_H}{b - \alpha'_H}\right)\right). \quad \text{(C3)}$$

Let b_o denote the bid of the opponent, so p_H is defined as follows:

$$\begin{aligned} p_H &= \text{Prob}(H\text{-type} \mid b_o \leq b) \\ &= \frac{\text{Prob}(b_o \leq b \mid H\text{-type})\text{Prob}(H\text{-type})}{\text{Prob}(b_o \leq b)}, && \text{(C4)} \\ &= (\text{Prob}(b_o \leq b \mid H\text{-type})\text{Prob}(H\text{-type})) \\ &\quad \cdot \big(\text{Prob}(b_o \leq b \mid H\text{-type})\text{Prob}(H\text{-type}) \\ &\quad + \text{Prob}(b_o \leq b \mid L\text{-type})\text{Prob}(L\text{-type})\big)^{-1}, && \text{(C5)} \\ &= \frac{((b - \alpha'_H)/(100 + h - \alpha'_H))p}{((b - \alpha'_H)/(100 + h - \alpha'_H))p + (b/100)(1-p)}, && \text{(C6)} \end{aligned}$$

where α'_H is the lower bound of the bid from an H-type, which is equal to h. Equivalently, the lower bound of the bid from an L-type, α'_L is equal to 0. After doing these replacements, the final expressions for p_H and p_L are

$$p_H = \frac{(b-h)p}{b - hp}, \quad \text{(C7)}$$

$$p_L = (1 - p_H) = \frac{b(1-p)}{b - hp}; \quad \text{(C8)}$$

and the first-order condition yields the following threshold for the H-type:

$$\tilde{b}_H^N = \frac{c_H + b + h(1-p)}{2}. \quad \text{(C9)}$$

Similarly, when the current bid from the H-type is greater than 100, the problem is as follows:

$$\max_{b_H \geq c_H} (b_H - c_H) \left(\overbrace{p_L \text{Prob}(b_H - h < b_L \mid b_L \leq 100)}^{\text{winning versus an } L\text{-type}} + \underbrace{p_H \times \text{Prob}(b_H < b'_H \mid b'_H \leq b)}_{\text{winning versus an } H\text{-type}} \right), \quad \text{(C10)}$$

$$\max_{b_H \geq c_H} (b_H - c_H)\left(p_L\left(\frac{100 - b_H + h}{100 - \alpha'_L}\right) + p_H\left(\frac{b - b_H}{b - \alpha_H}\right)\right), \quad \text{(C11)}$$

where p_H and p_L are defined as follows:

$$\begin{aligned} p_H &= \frac{((b - \alpha'_H)/(100 + h - \alpha'_H))p}{((b - \alpha'_H)/(100 + h - \alpha'_H))p + 1 \times (1-p)} && \text{(C12)} \\ &= \frac{(b-h)p}{(b-h)p + 100(1-p)}, && \text{(C13)} \end{aligned}$$

$$p_L = \frac{100(1-p)}{(b-h)p + 100(1-p)}. \quad \text{(C14)}$$

The first-order condition yields the following threshold for the H-type in this case:

$$\tilde{b}_H^N = \frac{c_H + (1-p)(100 + h) + pb}{2}. \quad \text{(C15)}$$

Now let us turn to the threshold for the L-type. Define b_c as the last bid at which the L-type was in rank 2. When $b_c \leq 100$; the threshold for the L-type depends on the value of b_c. When $b_c \geq 100$, the L-type does not have additional information to update the upper bound on a potential L-type competitor's bid; she only knows it is less than 100. When $b_c < 100$, she needs to update her threshold, $\tilde{b}_L^N$, using b_c as the upper bound of the potential L-type. The problem in this case is as follows:

$$\max_{b_L \geq c_L} (b_L - c_L) \left(\overbrace{q_L \text{Prob}(b_L < b'_L \mid b'_L \leq b_c)}^{\text{winning versus an } L\text{-type}} + \underbrace{q_H \times \text{Prob}(b_L < b_H - h \mid b_H \leq b_c)}_{\text{winning versus an } H\text{-type}} \right), \quad \text{(C16)}$$

$$\max_{b_L \geq c_L} (b_L - c_L)\big(q_L(1 - \text{Prob}(b'_L \leq b_c \mid b'_L \leq b_c)) + q_H(1 - \text{Prob}(b_H \leq b_L + h \mid b_H \leq b_c))\big), \quad \text{(C17)}$$

$$\max_{b_H \geq c_H} (b_H - c_H)\left(q_L\left(\frac{b_c - b_H + h}{b_c - \alpha'_L}\right) + q_H\left(\frac{b_c - b_H}{b_c - \alpha'_H}\right)\right), \quad \text{(C18)}$$

where q_H and q_L are the conditional probabilities of the competitor being an H-type and L-type, correspondingly, given b_c. They are given by

$$q_H = \frac{(b_c - h)p}{b_c - hp}, \tag{C19}$$

$$q_L = \frac{b_c(1-p)}{b_c - hp}. \tag{C20}$$

The first-order condition yields the following threshold:

$$\tilde{b}_L^N = \frac{c_L + b_c - hp}{2}. \tag{C21}$$

Similarly, when $b_c \geq 100$, the problem can be formulated as follows:

$$\max_{b_L \geq c_L}(b_L - c_L)\left(\overbrace{q_L \mathrm{Prob}(b_L < b_L' \mid b_L' \leq b_c)}^{\text{winning versus an } L\text{-type}} + \underbrace{q_H \times \mathrm{Prob}(b_L < b_H - h \mid b_H \leq b_c)}_{\text{winning versus an } H\text{-type}}\right), \tag{C22}$$

$$\max_{b_L \geq c_L}(b_L - c_L)(q_L(1 - \mathrm{Prob}(b_L' \leq b_c \mid b_L' \leq b_c)) + q_H(1 - \mathrm{Prob}(b_H \leq b_L + h \mid b_H \leq b_c))), \tag{C23}$$

$$\max_{b_H \geq c_H}(b_H - c_H)\left(q_L\left(\frac{100 - b_H + h}{100 - \alpha_L'}\right) + q_H\left(\frac{b_c - b_H}{b_c - \alpha_H'}\right)\right), \tag{C24}$$

where q_H and q_L are the conditional probabilities of the competitor being an H-type and L-type, correspondingly, given b_c. They are given by

$$q_H = \frac{(b_c - h)p}{(b_c - h)p + 100(1-p)}, \tag{C25}$$

$$q_L = \frac{100(1-p)}{(b_c - h)p + 100(1-p)}. \tag{C26}$$

The first-order condition yields the following threshold:

$$\tilde{b}_L^N = \frac{c_L + 100(1-p) + (b_c - h)p}{2}. \tag{C27}$$

Summarizing, we have the following thresholds:

$$\tilde{b}_H^N = \begin{cases} \dfrac{c_H + b + h(1-p)}{2} & \text{if } b < 100, \\[2ex] \dfrac{c_H + (1-p)(100+h) + pb}{2} & \text{if } b \geq 100. \end{cases}$$

$$\tilde{b}_L^N = \begin{cases} \max\left\{\dfrac{c_L + b_c - hp}{2}, c_L\right\} & \text{if } b_c < 100, \\[2ex] \max\left\{\dfrac{c_L + 100(1-p) + (b_c - h)p}{2}, c_L\right\} & \text{if } b_c \geq 100. \end{cases} \tag{C28}$$

When $p = \frac{1}{2}$, if two L-types compete, they will stop bidding at the maximum cost minus δ; this outcome is equivalent to the case when the type of the competitor is known. If two H-types compete, they will stop bidding when they reach the corresponding thresholds; in contrast, when they know the type of the competitor, they bid down to cost. In general, $\tilde{b}_H^N > \tilde{b}_H$ when $b < 100$ and $\tilde{b}_H^N < \tilde{b}_H$ when $b \geq 100$. When an H-type competes against an L-type, bidding stops when they reach the corresponding thresholds; note that $\tilde{b}_L^N \leq \tilde{b}_L$.

Appendix D. Experimental Results per Session

Tables D.1 and D.2 present average quality-adjusted costs and bid decrements by experimental session.

Table D.1 **Average Quality-Adjusted Prices**

Session	Full	Rank10	Rank_N	Rank10_N	Rank	SB
(a) Symmetric auctions (46 auctions per session)						
1	70.67	63.35	67.64	64.91	58.33	60.39
2	73.91	66.85	64.13	61.26	63.24	56.98
3	64.91	47.80	56.11	68.22	67.17	60.72
4	62.98	65.76	62.33	62.22	64.39	54.76
5	67.61				63.02	
6	69.89				67.76	
7	68.26				61.33	
8	73.17				64.20	
9	68.00					
10	75.30					
Mean	69.47	60.94	62.55	64.15	63.68	58.21
Std. dev.	3.93	8.88	4.83	3.12	3.03	2.86
(b) Asymmetric auctions (74 auctions per session)						
1	70.38	61.35	56.24	55.38	51.57	57.50
2	67.84	61.47	56.92	54.77	61.51	57.69
3	62.74	55.18	51.09	58.91	60.49	58.73
4	61.80	58.45	57.43	52.42	57.57	54.26
5	68.68				51.09	
6	64.20				58.58	
7	64.78				55.27	
8	64.04				49.16	
9	62.91					
10	69.12					
Mean	65.65	59.11	55.42	55.37	55.66	57.04
Std. dev.	3.07	2.97	2.92	2.68	4.62	1.94

Table D.2 **Average Bid Decrement**

Session	Full	Rank10	Rank_N	Rank10_N	Rank
(a) Symmetric auctions (46 auctions per session)					
1	6.16	9.36	12.40	13.48	20.91
2	7.65	10.38	15.31	8.45	10.74
3	4.63	12.20	13.83	11.83	9.84
4	6.86	10.75	12.31	14.00	11.35
5	5.56				12.97
6	7.20				10.71
7	8.43				10.71
8	4.95				13.68
9	6.84				
10	5.20				
Mean	6.35	10.67	13.46	11.94	12.61
Std. dev.	1.26	1.17	1.42	2.50	3.59
(b) Asymmetric auctions (74 auctions per session)					
1	4.48	12.85	17.01	13.24	28.03
2	7.36	14.96	17.38	7.92	18.46
3	5.26	14.48	13.80	10.10	18.58
4	6.15	10.33	14.17	18.44	17.29
5	6.16				22.14
6	3.97				17.20
7	7.49				25.26
8	6.51				27.93
9	5.58				
10	5.26				
Mean	5.82	13.16	15.59	12.43	21.86
Std. dev.	1.14	2.09	1.87	4.57	4.65

References

Bichler, M., J. Kalagnanam. 2005. Configurable offers and winner determination in multi-attribute auctions. *Eur. J. Oper. Res.* **160**(2) 380–394.

Bichler, M., R. Steinberg. 2007. Introduction to the special issue on e-auctions for procurement operations. *Production Oper. Management* **16**(4) 401–403.

Carare, O., M. Rothkopf. 2005. Slow Dutch auctions. *Management Sci.* **51**(3) 365–373.

Carter, C. R., C. K. Stevens. 2007. Electronic reverse auction conguration and its impact on buyer price and supplier perceptions of opportunism: A laboratory experiment. *J. Oper. Management* **25**(5) 1035–1054.

Carter, C. R., L. Kaufmann, S. Beall, P. L. Carter, T. E. Hendrick, K. J. Petersen. 2004. Reserve auctions—Grounded theory from the buyer and supplier perspective. *Transportation Res. Part E: Logist. Transportation Rev.* **4**(3) 229–254.

Cason, T. N. 1995. An experimental investigation of the seller incentives in the EPA's emission trading auction. *Amer. Econom. Rev.* **85**(4) 905–922.

Chen-Ritzo, C.-H., T. P. Harrison, A. M. Kwasnica, D. J. Thomas. 2005. Better, faster, cheaper: An experimental analysis of a multiattribute reverse auction mechanism with restricted information feedback. *Management Sci.* **51**(12) 1753–1762.

Compte, O., P. Jehiel. 2007. Auctions and information acquisition: Sealed bid or dynamic formats? *RAND J. Econom.* **38**(1) 355–372.

Cox, J. C., B. Roberson, V. L. Smith. 1982. Theory and behavior in single object auctions. *Res. Experiment. Econom.* **2** 1–43.

Cox, J. C., V. L. Smith, J. M. Walker. 1988. Theory and individual behavior of first-price auctions. *J. Risk and Uncertainty* **1**(1) 61–99.

Easley, R. F., R. Tenorio. 2004. Jump bidding strategies in Internet auctions. *Management Sci.* **50**(10) 1407–1419.

Elmaghraby, W. 2007. Auctions within e-sourcing events. *Production Oper. Management* **16**(4) 409–422.

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

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

Engelbrecht-Wiggans, R., E. Haruvy, E. Katok. 2007. A comparison of buyer-determined and price-based multiattribute mechanisms. *Marketing Sci.* **26**(5) 629–641.

Filiz-Ozbay, E., E. Y. Ozbay. 2007. Auctions with anticipated regret: Theory and experiment. *Amer. Econom. Rev.* **97**(4) 1407–1418.

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

Gallien, J., L. M. Wein. 2005. A smart market for industrial procurement with capacity constraints. *Management Sci.* **51**(1) 76–91.

Goeree, J. K., C. A. Holt, T. R. Palfrey. 2002. Quantal response equilibrium and overbidding in private-value auctions. *J. Econom. Theory* **104**(1) 247–272.

Haruvy, E., E. Katok. 2010. Increasing revenue by decreasing information in procurement auctions. Working paper, Pennsylvania State University, University Park.

Haruvy, E., P. T. L. Popkowski Leszczyc. 2010. The impact of online auction duration. *Decision Anal.* **7**(1) 99–106.

Holt, C. A. 1980. Competitive bidding for contracts under alternative auction procedures. *J. Political Econom.* **88**(3) 435–445.

Isaac, R. M., D. James. 2000. Just who are you calling risk averse? *J. Risk Uncertainty* **20**(2) 177–87.

Isaac, R. M., T. C. Salmon, A. Zillante. 2005. An experimental test of alternative models of bidding in ascending auctions. *Internat. J. Game Theory* **33**(2) 287–313.

Isaac, R. M., T. C. Salmon, A. Zillante. 2007. A theory of jump bidding in ascending auctions. *J. Econom. Behav. Organ.* **62**(1) 144–164.

Jap, S. D. 2002. Online reverse auctions: Issues, themes, and prospects for the future. *J. Acad. Marketing Sci.* **30**(4) 506–525.

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

Kagel, J. H. 1995. Auctions: A survey of experimental research. *Handbook Experiment. Econom.* **89**(5) 501–535.

Kagel, J. H, D. Levin. 1993. Independent private value auctions: Bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *Econom. J.* **103**(419) 868–879.

Kagel, J. H., A. E. Roth. 1995. *Handbook of Experimental Economics*, 1st ed. Princeton University Press, Princeton, NJ.

Kagel, J. H., R. M. Harstad, D. Levin. 1987. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica* **55**(6) 1275–1304.

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

Koppius, O. R., E. van Heck. 2003. Information architecture and electronic market performance in multidimensional auctions. Doctoral dissertation, Erasmus University Rotterdam, Rotterdam, The Netherlands.

Kostamis, D., D. R. Beil, I. Duenyas. 2009. Total-cost procurement auctions: Impact of suppliers' cost adjustments on auction format choice. *Management Sci.* **55**(12) 1985–1999.

Krishna, V. 2002. *Auction Theory*. Academic Press, San Diego.

Kwasnica, A. M., E. Katok. 2007. The effect of timing on jump bidding in ascending auctions. *Production Oper. Management* **16**(4) 483–494.

Lucking-Reiley, D. 1999. Using field experiments to test equivalence between auction formats: Magic on the Internet. *Amer. Econom. Rev.* **89**(5) 1063–1080.

Mithas, S., J. L. Jones. 2007. Do auction parameters affect buyer surplus in e-auctions for procurement? *Production Oper. Management* **16**(4) 455–470.

Neugebauer, T., R. Selten. 2006. Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets. *Games Econom. Behav.* **54**(1) 183–204.

Ockenfels, A., R. Selten. 2005. Impulse balance equilibrium and feedback in first price auctions. *Games Econom. Behav.* **51**(1) 155–170.

Parkes, D. C., J. Kalagnanam. 2005. Models for iterative multiattribute procurement auctions. *Management Sci.* **51**(3) 435–451.

Parkes, D. C., L. H. Ungar. 2000. Iterative combinatorial auctions: Theory and practice. *Proc. 17th Natl. Conf. Artificial Intelligence*, Association for the Advancement of Artificial Intelligence, Palo Alto, CA, 74–81.

Raviv, Y. 2008. The role of the bidding process in price determination: Jump bidding in sequential English auctions. *Econom. Inquiry* **46**(3) 325–341.

Roth, A. E., A. Ockenfels. 2002. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *Amer. Econom. Rev.* **92**(4) 1093–1103.

Rothkopf, M. H. 2007. Decision analysis: The right tool for auctions. *Decision Anal.* **4**(3) 167–172.

Rothkopf, M. H., A. B. Whinston. 2007. On e-auctions for procurement operations. *Production Oper. Management* **16**(4) 404–408.

Sterman, J. D. 1989. Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Sci.* **35**(3) 321–339.

Sterman, J. D. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. McGraw-Hill, Boston.

Teich, J. E., H. Wallenius, J. Wallenius, O. R. Koppius. 2004. Emerging multiple issue e-auctions. *Eur. J. Oper. Res.* **159**(1) 1–16.

Vickrey, W. 1961. Counter speculation, auctions and competitive sealed tenders. *J. Finance* **16**(1) 8–37.