PRODUCTION AND OPERATIONS MANAGEMENT

Vol. 0, No. 0, xxxx–xxxx 2012, pp. 1–17
ISSN 1059-1478 | EISSN 1937-5956 | 12 | 0 | 0001

DOI 10.1111/j.1937-5956.2012.01356.x

Increasing Revenue by Decreasing Information in Procurement Auctions

Ernan Haruvy, Elena Katok
Jindal School of Management, The University of Texas at Dallas, Richardson, TX 75080, USA, eharuvy@utdallas.edu, ekatok@psu.edu

We report on results of several laboratory experiments that investigate on-line procurement auctions in which suppliers bid on price, but exogenous bidder quality affects winner determination. In procurement auctions, bidder quality may or may not be publicly known to all bidders, and the effect of this quality transparency on the auction outcome is one aspect of auction design that we examine. The second aspect of auction design that we examine is the effect of price visibility on the auction outcome, and the interaction between price visibility and quality transparency. In terms of price visibility, we consider two extreme cases: the sealed bid request for proposals (RFPs), and the open-bid dynamic auction event. In terms of bidder quality transparency, we also consider two extreme cases: a setting in which bidder qualities are publicly known and the case in which they are private. We find that in our laboratory experiments, the RFP format is consistent in generating higher buyer surplus levels than does the open-bid dynamic format. This advantage is independent of the quality transparency. In contrast, the open-bid format is highly sensitive to quality transparency, generating significantly lower buyer surplus levels when the information about bidder quality is public.

Key words: bidding; procurement auctions; reverse auctions; multi-attribute auctions; behavioral game theory; experimental economics
History: Received: February 2010; Accepted: January 2012 by Kalyan Singhal, after two revisions.

1. Introduction

Internet-enabled competitive sourcing mechanisms are fast becoming an essential part of the procurement toolkit. A recent large-scale study that surveyed close to 200 companies in a wide cross-section of industries (CAPS 2006) reported that nearly 65% use electronic procurement mechanisms and over 60% regularly use on-line procurement auctions. The average total amount spent through on-line procurement auctions was reported to be almost \$9 billion (about 1% of gross sales), and growing at about 20% per year. Beall et al. (2003) report that although on-line procurement auctions account for less than 10% of the actual total amount spent, for some firms this figure can potentially increase to as much as 50%, indicating high growth potential. In addition to substantial cost savings, on-line procurement auctions also deliver a number of other benefits, including an increase in the supplier base (CAPS 2006) and accelerated transaction time (Shugan 2005).

One of the most important factors that make most procurement auctions fundamentally different from most consumer auctions is that price is typically not the main attribute used to award contracts. Exogenous non-price attributes (e.g., distance from buyer, incumbency status, and reputation) have a major effect on the expected buyer's surplus. Some of these non-price attributes, such as the effect of the distance from the buyer on transportation costs, may be easily quantifiable, whereas others, such as reputation, may not be. Although the bidding in most on-line procurement auctions is in terms of price (this is different from pure score auctions, in which bidders can submit bids in terms of price and quality, e.g., Che 1993, Branco 1997), most awards are not made based only on price. Jap (2002) was the first to point out that "…the vast majority of [on-line reverse] auctions … do not determine a winner" (p. 510). The buyer reserves the right to award a contract on any basis, and sometimes the lag between the end of the auction and the announcement of the winner may be as long as 4–6 weeks (Jap 2002). The long lag between the end of the auction and the award decision may indicate that sometimes buyers themselves are unable to quantify the value of quality until after the auction. This may be because some bidders have not been fully vetted before the auction (Wan and Beil 2009, Wan et al. 2012). For the purpose of this article, we collectively group non-price attributes and label them *quality* (see

Correction added on 24 July 2012 after initial online publication on 1 May 2012: authors' affiliation have been updated.

also Tunca et al. 2008, and Engelbrecht-Wiggans et al. 2007, who use a similar approach). It is a reasonable approximation of reality to model bidders as having some private information about their own quality that is not fully known to other bidders, although it will eventually be revealed to the buyer.

The degree to which information about bidder quality is private or full varies and depends both on the industry in which the buyer operates and on auction design decisions by the buyer. For example, according to a senior category managed at Ariba, a major provider of e-sourcing solutions, the amount of bidder-specific information that buyers are able and willing to provide to the suppliers depends on the industry. For direct materials, buyers usually identify supplier-specific transformation factors and provide feedback to bidders during the auction, as well as make final awards, based on transformed bids (Elmaghraby et al. 2012). In industries in which a relatively small number of suppliers compete repeatedly, such as the highway construction industry (Bajari and Lewis 2011), bidders may know the identity of their competitors, and consequently bidder-specific attributes may be well known. In those settings, even though bidders will undoubtedly have some private information about their quality advantage or disadvantage, they will also have sufficient information to make reasonable adjustments to their price bids, to allow for quality advantage or disadvantage relative to competitors (thus it is reasonable to model bidder information as public).

In contrast, in Ariba auctions for indirect materials and services, buyers tend to not reveal bidder-specific information before or during the auction, and incorporate non-price attributes into their final selection decision after the auction (Elmaghraby et al. 2012). Similarly, Tunca et al. (2008) report that in auctions for legal services, General Electric (GE) identifies a bidder-specific quality measure that serves as an input to the final decision, but is not revealed to the bidders. More generally, in industries that are not very concentrated, the same bidders do not interact repeatedly and bidders are less likely to have knowledge about the non-price attributes of their competitors. In these kinds of settings, price adjustments to account for quality differences are more difficult to make. The buyer can contribute to this uncertainty by revealing or concealing key quality attributes of each participating bidder (such as incumbency status, size of joint investments, etc.) or by revealing or concealing the identity of the bidders.

The importance of price visibility has to do with the strong relationship between time-varying factors related to observed bids (such as bid concentration and bidding rate) and subsequent bids (Bapna et al. 2008a, Bradlow and Park 2007, Park and Bradlow 2005, Wang et al. 2008). This relationship implies that bidders may infer something about their competition by observing price dynamics in the auction. This dramatically changes the theoretical properties of the equilibrium with implications for auction design. It is the purpose of the present work to address some of these issues. Price visibility may also affect supplier relationships (see Zhong and Wu [2009] for an empirical study of the effect of procurement auctions on relationships with preferred suppliers), although the relationship aspect is outside the scope of the present work.

Price visibility can range from one extreme—full visibility, where all bids are displayed to all competitors—all the way to the opposite extreme of private visibility, where each bidder submits a single bid (a sealed-bid auction, or a RFP). Much of the literature on sealed-bid procurement auctions deals with private sector procurement. For example Bajari and Lewis (2011) report on the California Department of Transportation procurement for highway repair contracts that incorporates timing consideration in winner selection. Katzman and McGeary (2008) report on Medicare auctions for durable medical equipment, prosthetics, orthotics, and supplies. Open-bid auctions are common in procurement. For example, most of Ariba's auctions are open bid, and about 30% of them have full price visibility (Elmaghraby et al. 2012). Jap and Haruvy (2008) report on open-bid full visibility online procurement auction for automobile components, specifically metal parts and plastics, and Tunca et al. (2008) investigate the GE auctions for legal services.

Jap (2003) writes that the full price visibility format is more damaging to the buyer-supplier relationship than the sealed-bid format because under full price visibility suppliers are more suspicious about buyer opportunism and are less willing to make customized investments. In two large-scale studies that consider partial price visibility formats and the effect of the number of bidders, Jap (2007) and Jap and Haruvy (2008) find that price visibility choices are closely related to both bidding intensity and the resulting relationship between buyer and supplier.

The investigation into price visibility is related to studies that look into revenue equivalence or lack thereof between sealed-bid and open-bid formats. The existing format comparison literature largely pertains to single dimension price auctions with symmetric information. With asymmetric information (Maskin and Riley [2000] provide the seminal theoretical treatment), the closest empirical comparison has been by Athey et al. (2004). They study entry and bidding patterns in sealed bid and open auctions with heterogeneous bidders. Using data from U.S. Forest Service timber auctions, they find that sealed bid auctions attract more small bidders, shift the allocation toward these bidders, and can also generate higher revenues. The reason for this is that weaker bidders have a greater chance of winning in the sealed bid auctions since strong bidders shade their bids. In our setting, we keep participation exogenous and focus on the interaction between price visibility (open vs. sealed) and quality transparency.

In the next section, we present a model of procurement auctions that incorporates both quality transparency and price visibility, followed by analytical results that guide our laboratory investigation. In section 3 we describe the design of our laboratory experiments and the protocol we followed in conducting this study. In section 4 we present data analysis, and in section 5 we summarize and discuss our results as well as directions for future research and managerial implications.

2. Analytical Background

2.1. Buyer-Determined Auctions

In a buyer-determined auction, the buyer takes into account non-price attributes according to an explicit or implicit scoring rule, and awards the contract to the supplier whose bid provides the highest value. We assume, following Kostamis et al. (2009) and Engelbrecht-Wiggans et al. (2007), that at some point before the contract is awarded, the buyer will articulate non-price attribute trade-offs clearly enough to be able to compare the final bids and award the contract to the supplier whose bid delivers the highest surplus; suppliers know this allocation rule. We also assume that some of the non-price attributes are sufficiently objective for suppliers to know something about the value of their own non-price attributes to the buyer. The transparency of this quality information, one of the factors in our study, depends on whether or not suppliers know the buyer's values for the non-price attributes of their competitors.

The second factor in our study is price visibility. Following Jap (2003) we look at two extreme cases: the RFP format is the *first-price sealed bid* buyer-determined auction that has zero price visibility, and the *open-bid* format that has full price visibility. In the open-bid format, suppliers participate in an auction event, in which they can submit bids, observe other suppliers' bids, and submit new bids in real time (Rangan 1998).

Figure 1 shows an example of bidding activity in an open-bid buyer-determined auction from freemarkets. The key observation from Figure 1 is that bids are not monotonically decreasing over time. Even toward the end of the auction, we observe that some bids are significantly above the standing lowest bid. This is suggestive of the fact that some suppliers (those submitting high bids) do not compete on price alone—they count on their quality advantage to win.

For the remainder of this article we use the term *quality* to refer to bidder-specific exogenous non-monetary attributes. We refer to a bidder's quality minus cost as the *score* and to quality minus bid as the *score bid*. However, although we use the word "score" to describe the position of the firm, please note that this is not a score auction because bidders cannot adjust their quality attributes in the bidding. Score auctions (beginning with Che 1993, Branco 1997) are popular in the auction literature and are occasionally used in government auctions, but are not common in B2B procurement practices (Jap 2002). As we pointed out in the introduction, the vast majority of on-line procurement auctions are buyer determined. Buyer determined auctions generate more value to the buyer than price-based auctions do, as long as there are enough suppliers competing for the contract (Engelbrecht-Wiggans et al. 2007),

Figure 1 An Example of Bidding Activity in an Open-Bid Buyer-Determined Auction. Reproduced with Permission from Jap (2002)

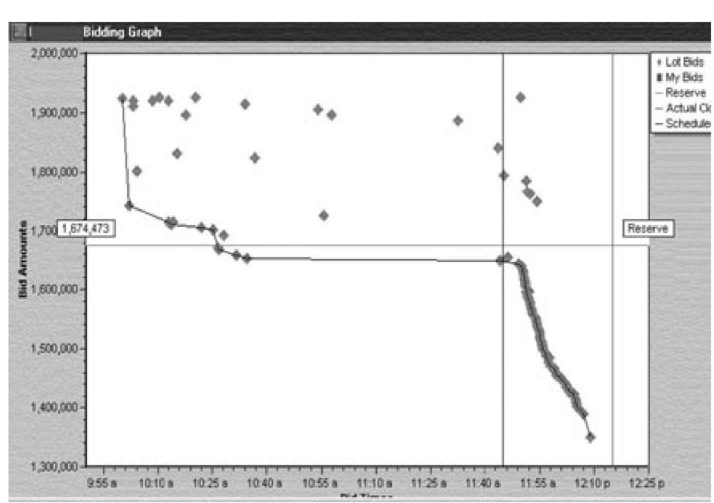

which explains their popularity and success in practice. Our goal in this study is to understand the effect of the quality transparency, and price visibility in these auctions, through the use of analytical modeling and laboratory experiments.

Specifically, we focus on two research questions. First, which format, open bid (full price visibility) or sealed bid (no price visibility), generates higher levels of buyer surplus? Although this question is related to the literature on testing auction revenue equivalence in the experimental economics (see Kagel [1995] for a review),[1] the test in our article involves buyer-determined procurement auctions (with the quality component), so even though we report results that are in line with the early literature (the open bid vs. sealed bid comparison is theoretically isomorphic to forward auction comparisons), the fact that these regularities continue to hold in the new setting of buyer-determined procurement auctions is an important result.

Second, we examine how quality transparency affects the performance of the two auction mechanisms. In contrast to Engelbrecht-Wiggans et al. (2007), we specifically examine the effect of quality transparency on sealed bid buyer determined auctions, rather than assuming, as in Engelbrecht-Wiggans et al. (2007), that this information is always private. Our work is different from Kostamis et al. (2009) because we do not assume that the buyer selects the auction format based on the realization of bidder types, but instead we study the effect of the auction format under different information structures. Although the first question of open bid vs. sealed bid has parallels in the forward auction literature, the second set of questions is unique to buyer-determined procurement settings.

2.2. Structural Properties of Buyer Determined Auctions

Consider a setting in which N bidders compete to provide a contract to a buyer. This contract can be thought of as a single unit of an indivisible good. Bidders are heterogeneous in quality and cost, and are pre-qualified and verified to meet the minimum quality requirements the buyer has specified. Bidder i has a privately known, non-negative cost C_i of providing the contract. The buyer's value for the contract depends on which bidder provides it. We refer to the value the buyer places on these non-monetary, bidder-specific, exogenously fixed attributes as the *quality* Q_i. We assume that both the buyer and bidder i know Q_i, and one of the issues we investigate is the effect of full information about Q_i on the auction outcome.

Each bidder's type (C, Q) is assumed to be a random vector identically distributed, independent and non-degenerate. In general, the C's are not independent of the Q's. In the sealed-bid first-price auction, each bidder i submits a single bid $B_i \equiv B(C_i, Q_i)$. The buyer selects the bidder i whose bid generates the highest buyer surplus $Q_i - B_i$, and pays the winning bidder i an amount B_i.

Let us start by examining a setting in which the buyer and bidder i both know the quality Q_i, and other bidders do not. This is the setting studied in Engelbrecht-Wiggans et al. (2007). In this setting, since buyer and bidder i both know Q_i, we can think of each bidder i as bidding in the score space (compete on the difference between quality and bid, $Q_i - B_i$) directly. Consequently, the buyer-determined (hereafter BD) mechanism is isomorphic to a regular (price-based forward) sealed bid auction with independent, privately known valuations. Engelbrecht-Wiggans et al. (2007) showed that the usual results for auctions with independent private values apply to risk-neutral bidders: (i) Truthful bidding in the score space is the dominant strategy in the sealed bid second-price auction. (ii) There is expected buyer surplus equivalence between first- and second-price sealed bid BD auctions. (iii) The first-price sealed bid BD auction has a unique pure-strategy symmetric risk neutral bidding equilibrium.

Now consider a setting in which all qualities, Q_i, are known to all bidders and bidding takes place via an open-bid auction. This is the open-bid auction setting of Kostamis et al. (2009). Since Q_i's are known, even though the bidding takes place in the price-space, each bid B_i can be translated into the score space by looking at the difference between a bidder quality and her bid: $Q_i - B_i$. Therefore, bidders in this auction have a weakly dominant strategy to bid truthfully (enter quality minus cost) in the score space.

Below we state two structural characterizations of the behavior in the settings we are studying. These results are independent of Q and C distributions. Let us denote $\Pi^{Bidding}_{Information}$ as the expected buyer's surplus in a buyer determined auction using a *bidding* format that is either sealed- or open-bid (SB or O) with *Information* about quality that is either private or full (P or F), so for example Π^{SB}_{P} is the expected buyer's surplus in the sealed bid BD auction with private information. Thus, we have four possible configurations for auction formats: Sealed-Bid Private, Open-Bid Private, Sealed-Bid Full, and Open-Bid Full.

CHARACTERIZATION 1 (EXPECTED BUYER-SURPLUS EQUIVALENCE) *If bidders are risk-neutral, the expected buyer surplus is the same in the Open-Bid Full format and in the Sealed-Bid Private format, or formally:* $\Pi^{O}_{F} = \Pi^{SB}_{P}$.

This expected surplus equivalence is no different from the standard revenue equivalence result for forward auctions under risk neutrality because both open-bid full information format and the sealed-bid private information format can be reduced to a single-dimension (score) space. That is, in either of these auction formats, bidders can be thought of as placing a score bid instead of a price bid and all variables can be simply transformed. Thus, the equivalence proof is isomorphic to the well-known single dimension (price-only) revenue equivalence. The same cannot be said for the other two formats in our study because they cannot be transformed to a single dimension due to bidder information asymmetries.

In addition to the expected surplus equivalence in the above two cases, Kostamis et al. (2009) showed that in the open-bid BD auction with full information, bidders have the dominant strategy to bid the minimum increment in the score space above the maximum of the others' bids in the score space, down to their cost. Engelbrecht-Wiggans et al. (2007) showed that in the sealed bid auction with private information, the optimal bid depends on the distribution of bidder scores and for some distributions (including the distribution in our laboratory setting) it can be derived analytically. In the rest of this section, we describe structural properties of the bidding behavior in the open-bid auction with private information and the sealed-bid auction with full information.

CHARACTERIZATION 2 (SEALED-BID AUCTION WITH FULL INFORMATION) *If score and quality are not independent, then in the Sealed Bid Full format, bids depend on the observed qualities of the other bidders.*

The proof is in Appendix 1. Intuitively, Characterization 2 follows because in the sealed-bid BD auction with full information, monotonicity in the score space may not generally hold. That is, higher ($Q_i - C_i$) may not mean higher bid in the score space ($Q_i - B_i$) if Q_j enters the bid by bidder i for $j \neq i$. For example, bidder i with the same score ($Q_i - C_i$) may bid lower when Q_i is relatively low, or higher when Q_i is relatively high. Given the lack of monotonicity in the score space, standard approaches to analytically deriving optimal bid functions are not feasible and expected profit-maximizing bids may be analytically intractable (Cantillon 2008).

The corresponding theory for open-bid auction with private information can similarly be constructed in a static framework. That is, we can show that conditional on observing others' bids and conditional on these bids, the unique score to bid mapping cannot be the solution. However, it is more constructive to think of bids in the open-bid auction in a dynamic framework. We can then see that bids cannot fully reveal one's quality or score. If they did, any forward-looking bidder (see e.g., Zeithammer 2006) would be able to profitably deviate from such equilibrium by bidding as a bidder with slightly lower quality. All bidders trying to outbid this clever bidder by the smallest increment will thus fail. This means that open-bid auctions with private information are close in spirit to sealed-bid auctions since incremental bidding can be ruled out. We will refer to this insight in the following section and when reporting our results.

2.3. Equilibrium Predictions

In this section we focus on the equilibrium analysis for the specific setting we implemented in the laboratory. We cross two factors that pertain to BD auctions, price visibility and quality transparency, and vary each at two levels, for a total of four treatments. The auction format is either sealed bid (*SB*) that has zero price visibility, or open bid (*O*), that has full price visibility. Quality transparency deals with the knowledge of the realization of the bidders' Q's. In the *Full* transparency conditions (F) all bidders know the realizations of all their competitors' Q's, and in the *Private* condition (P) bidders know only the realization of their own Q's.

As in Engelbrecht-Wiggans et al. (2007), the cost and the quality in our experiments follow a linear relationship $Q = C + 300X$ in which $C \sim$ Uniform (0,100) and $X \sim$ Uniform (0,1). This setting is realistic in that Q increases in C, but the two are not perfectly correlated. In all our treatments $N = 4$, to generate sufficient competition to make the BD auction more profitable to buyers than the price-based auction. We conducted each treatment with two different (Q, C) profiles. We pre-generated a (Q, C) combination from the distributions above for each participant and each auction period and used those identical pre-generated profiles in all treatments. In all treatments participants know the distributions of Q's and C's.

Even though we know, based on Characterization 1, that there is an expected buyer-surplus equivalence between the open-bid auction with full information and the sealed-bid auction with private information, this equivalence assumes risk neutrality, and the two mechanisms are not strategically equivalent in the sense that bidders in the two mechanisms do not have the identical optimal bidding strategies. We already mentioned that in the open-bid auction with full information bidders have the weakly dominant strategy (regardless of risk-aversion) to bid one bid increment above the current high score bid, and stop bidding when their score bid reaches their score (equivalently, when their price bid goes down to their cost), so in the open-bid auction with full information, the equilibrium bid function (for losing bidders) is simply

$$B_i = C_i. \tag{1}$$

Under the sealed-bid auction with private information, Engelbrecht-Wiggans et al. (2007) showed that for the (Q, C) distribution in our setting, the equilibrium profit maximizing bid function for risk neutral bidders is

$$B_i = C_i + \frac{1}{N}(Q_i - C_i). \tag{2}$$

Because bidders do not have identical bidding strategies under the two mechanisms, the buyer surplus levels for the two mechanisms are identical in expectation but are not exactly identical for a given (Q, C) realization.

Analytical expressions for equilibrium bid functions in the sealed-bid auction with full information and the open-bid auction with private information cannot be obtained using standard auction theory tools because, as we mentioned in the previous section, the two mechanisms violate the essential assumption of symmetry.

A bid function is a mapping between the known information and the bid. The problem is essentially one of asymmetric distributions, which is one of the most elusive problems in auction theory research. In recent years, there have been successful attempts to characterize the solution for procurement auctions with two bidders with the difference in qualities being common knowledge (e.g., Balestrieri 2008, Mares and Swinkels 2010). However, these solutions do not extend to more general cases. In general (see also Kaplan and Zamir 2011), the bidding solutions are shown to be non-linear, even for simple uniform distributions. Since we do not know the true mapping in equilibrium, and given that equilibrium bid functions in the open-bid auctions with private information and sealed-bid auctions with full information are surely *non-linear*, without knowing their functional forms we cannot provide a reasonable numerical approximation. Instead, we will analyze bidding behavior in the open-bid auction with private information and in the sealed-bid auction with full information, empirically.

An allocation is efficient when the bidder with the highest score wins the auction. Efficiency will be 100% in theory in the open-bid auction with full information because in theory losing bidders bid down to their cost. Likewise, efficiency will be 100% in the sealed-bid auction with private information because bids in the score space are monotonically increasing in the score. Efficiency is likely lowest in the sealed-bid auction with full information setting because bids in the score space are no longer monotonically increasing in the score due to quality information about others. Efficiency in the open-bid auction with private information case is expected to be nearly 100% because information about bids is far more limited in usefulness than information about qualities, and therefore should not create very much distortion relative to the sealed-bid auction with private information.

3. Research Hypotheses and Experimental Protocol

3.1. Research Hypotheses

We summarize the analytical results from section 2 by linking specific research hypotheses to these analytical results. The first hypothesis deals with the expected buyer surplus equivalence between the Open-Bid Full and the Sealed-Bid Private treatments (Characterization 1). The expected buyer surplus levels in those two treatments should not be statistically different (they should be close but not identical because the two formats are not strategically equivalent), so we will test this equivalence in our first hypothesis.

> Hypothesis 1 (Buyer Surplus Equivalence): Average buyer surplus in the Open-Bid Full treatment and in the Sealed-Bid Private treatment will not be statistically different.

As we mentioned in the introduction, our work contributes to the literature on revenue equivalence in auctions because we test it in buyer-determined procurement auctions—the presence of the quality component makes the setting more complex than standard forward auctions even though theoretically the mechanism is equivalent to bidding in the score space. It is valuable to check whether or not the regularities established in earlier forward auction studies continue to hold in the more complex setting.

The second hypothesis deals with the sealed-bid auction with full information. The revealed quality levels of the competitors affect bids under this format (Characterization 2). Intuitively, bidders should increase their score bid in response to the competitors' qualities.

> Hypothesis 2 (Sealed-Bid Full Information Format): Bidders in the sealed-bid full information setting increase their score bids in response to the competitors' quality.

The third hypothesis deals with the Open-Bid Private auction. As we discussed earlier, bids under this format can neither deterministically reveal quality nor score. A reasonable conjecture is that bidders under this format should decrease their score bid in response to the price bids by their competitors that they observe. If a bidder's own bid is the lowest, this bidder should decrease his score bid lesser than that when his bid is not the lowest.

Hypothesis 3 (Open-Bid Private Format): Bidders in the Open-Bid Private information setting decrease their score bids in response to bids they observe. They decrease their score bids less when their own bid is the lowest.

3.2. Experimental Protocol

We use a between-subjects design. Each treatment contained four independent cohorts. Each cohort was assigned to one of the two (Q, C) profiles, and each treatment contained two cohorts with each profile. Each cohort included eight participants for a total of 32 participants per treatment and 128 participants in the study. All human participants in our experiment were in the role of bidders. Participants were randomly assigned to one of the four treatments and to one of the four cohorts within a treatment. Each person participated in only one session. We conducted all experimental sessions at a major U.S. public university. We recruited subjects by posting flyers on campus inviting them to sign up for our study. Earning cash was the only incentive offered.

Upon arrival at the laboratory the subjects were seated at computer terminals. We handed out written instructions (see Appendix 3) to them and they read the instructions on their own. When all participants finished reading the instructions, we read the instructions to them aloud, to insure common knowledge about the rules of the game.

After we finished reading the instructions to the participants we started the actual game. In each session each participant competed in a sequence of 40 rounds of auctions. We used random matching that we kept the same for each profile. At the beginning of each round the eight participants in a cohort were divided into two groups of four bidders ($N = 4$) according to this pre-generated profile matching protocol. Each group of four bidders competed for the right to sell a single unit to a computerized buyer.

We programmed the experimental interface using the zTree system (Fischbacher 2007). The screen included information about the subject's cost and quality. In treatments with full information, subjects also saw information on all other bidders' qualities, but not costs. The screen also provided an input area to try different bids and a calculator to compute corresponding bids in the score space and the buyer surplus levels for different bids. In the sealed bid treatments, entering a bid concluded the auction. In the open bid treatments the bidding was allowed to start at any price (there was no reserve), and each auction lasted for at least 1 minute, and ended when no new bids were placed for 10 seconds.[2] In open bid treatments bidders could observe all price bids placed in real time. In the Open-Bid Full treatment, each bid entered was displayed accompanied by the associated quality and an anonymous but constant bidder ID and bids were binding in the sense that each bid placed had to increase the buyer surplus relative to the standing best bid. In the Open-Bid Private treatment, each bid had to be at least one bid decrement lower than the current low price bid. This feature is commonly assumed in auction research on dynamic bid formats (hence the terms "ascending" or "descending" bid auctions).

At the end of each round we revealed the same information in all conditions. This information included the bids and qualities of all bidders and the winner in that period's auction. The entire history of past winning prices and qualities in the session was also provided for explicit reference points (Dholakia and Simonson 2005).

For each auction in each round, the winning bidders earned the difference between their price bids and their costs, while the losing bidders earned zero. We computed cash earnings for each participant by multiplying the total earnings from all rounds by a pre-determined exchange rate and adding it to a $5 participation fee. Participants were paid their earnings from the auctions they won in private and in cash, at the end of the session.

4. Experimental Results

We begin by comparing buyer surplus levels across the four treatments.

In Table 1 we report the actual average buyer surplus levels and the percentage of efficient allocations in each treatment over all 40 rounds.[3] In Figure 2 we display the distributions of the log-transformed buyer

Table 1 Average Observed Buyer Surplus Levels, Log Surplus Levels, and Proportion of Efficient Allocations, and Theoretical Buyer Surplus for the Open-Bid Full and Sealed-Bid Private treatments. Those Two Treatments Should be 100% Efficient in Theory

	Open-Bid Full	Sealed-Bid Private	Open-Bid Private	Sealed-Bid Full
Actual buyer surplus (SE)	186.11 (4.17)	224.72 (2.73)	211.85 (6.12)	205.99 (7.97)
Log surplus (SE)	5.17 (0.025)	5.39 (0.013)	5.33 (0.033)	5.30 (0.039)
Theoretical buyer surplus (SE)	184.00 (0.66)	182.86 (0.22)	NA	NA
Theoretical log buyer surplus (SE)	5.17 (0.020)	5.18 (0.008)	NA	NA
Actual proportion of efficient allocations	86.88%	88.43%	85.94%	84.38%

Figure 2 Distributions of Log-Transformed Buyer Surplus Levels

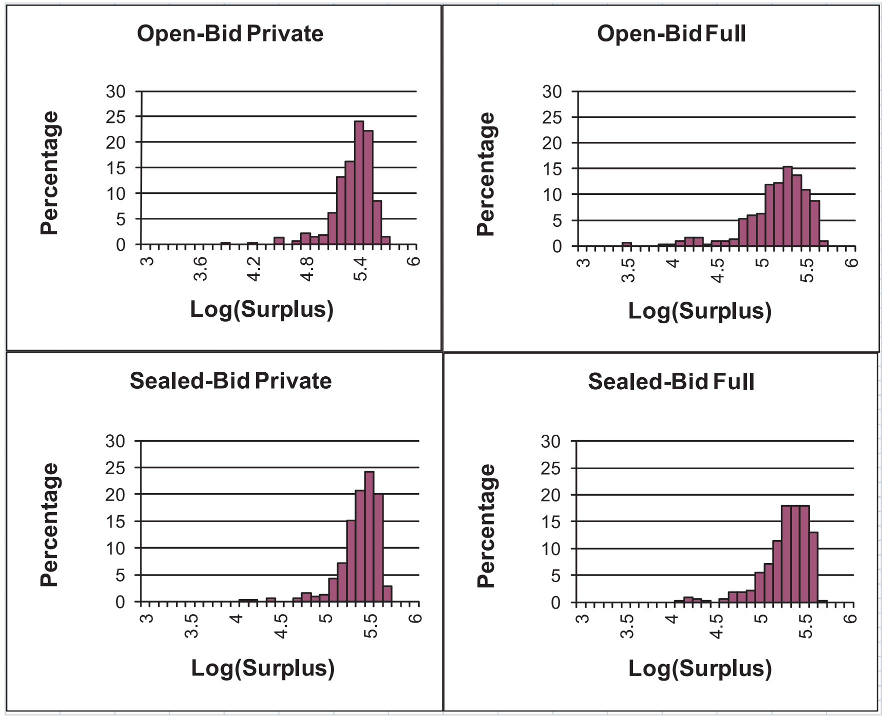

surplus in the four treatments (for discussion and details on the need for log transformation, see Bapna et al. 2008b). We see in Figure 2 that in line with the findings of Bapna et al. (2008b), the log surplus distributions have long left tails (corresponding to the right tails in Bapna et al.'s consumer surplus in forward auctions). In other words, surplus levels are clearly asymmetric and skewed.

Actual buyer surplus and log buyer surplus levels are not significantly different from theoretical predictions in the Open-Bid Full treatment. In the Sealed-Bid Private treatment, actual buyer surplus and log buyer surplus levels are significantly *higher* than the theoretical benchmarks. The proportion of efficient allocations is below 100% in the Open-Bid Full and Sealed-Bid Private treatments; actual efficiency does not significantly vary across our four treatments.

We continue by estimating individual bid functions in each of the four treatments. We do not restrict the form of the bid functions to be linear, but instead we allow them to be polynomials in all their terms. The bid, B_{it}, by bidder i in auction t is the dependent variable in all bid function regressions. Functions of the quality minus cost (Q_{it}, $-$ C_{it}) and cost, C_{it}, of bidder i at time t, are included as explanatory variables in all treatments except the open-bid auction with full information (where the theoretical bid function for losing bidders is $B_{it} = C_{it}$).[4] We avoid colinearity by estimating C and (Q-C) as regressors—as opposed to C and Q, which are correlated. The variables C and (Q-C) are independently distributed in our experimental setting.

For the sealed-bid auction with private information, no information other than own cost and own quality is given to the bidders and so we are left with the following general model:

$$B_{it} = \alpha(C_{it}) + \beta(Q_{it} - C_{it}) + \eta_i + \varepsilon_{it} \tag{3}$$

In theory, $\alpha()$ and $\beta()$ should be linear functions of the cost C and the score (Q-C), respectively, and the linear model will be one of the models we estimate. We also estimate models in which these functions are polynomials of various degrees.

In the Sealed-Bid Full treatment, bidders know their own quality disadvantage relative to other bidders. Let $Q^{(k)}$ be the k^{th} highest quality ($Q^{(1)}$ is the highest quality) and let $q_{it}^{(k)} = Q_{it}^{(k)} - Q_{it}$ represent bidder i's quality disadvantage relative to the k^{th} highest quality ($q_{it}^{(1)}$ is bidder i's quality disadvantage relative to the best competitor's quality). We estimate the following general model for the Sealed-Bid Full treatment:

$$B_{it} = \alpha(C_{it}) + \beta(Q_{it} - C_{it}) + \gamma(q_{it}^{(k)}) + \eta_i + \varepsilon_{it} \tag{4}$$

We estimate models in which functions $\alpha()$, $\beta()$, and $\gamma()$ are polynomials (see Appendix 2).

For the Open-Bid Private treatment, we no longer have the quality rank information but instead we now have competitors' bid information. Due to endogenous bid ranks, we use the average of the competitors' bids in the regressions we estimate. We denote the average of bidder i's competitors' bids at time t by $Avg(B_{-i,t})$.[5] Let the variable LB_{it} be 1 if bidder i's bid is the lowest bid in period t and 0 otherwise. The general regression model for the Open-Bid Private treatment is then:

$$B_{it} = \alpha(C_{it}) + \beta(Q_{it} - C_{it}) + \phi(Avg(B_{-i,t})) + \xi(Avg(B_{-i,t}) \times LB_{it}) + \eta_i + \varepsilon_{it} \quad (5)$$

Again, we estimate models in which functions $\alpha()$, $\beta()$, $\phi()$, and $\xi()$ are polynomials (see Appendix 2).

Note that in all the regressions we estimate, there are two error components: one that is independent across all observations, ε_{it}, and one that is participant-specific, η_i. Both error terms have means of zero and positive standard deviations. This treatment of the individual effects is known as the random effects model. Having estimated polynomial functions of various degrees, we find no significant improvement beyond cubic (third degree) polynomials in terms of R^2 fit. Appendix 2 reports the R^2 for each setting under different functional assumptions about the degree of the polynomials.

In Table 2, we report estimates for the following models: For Open-Bid Full and Sealed-Bid Private, we report the linear estimates because the linear bidding functions correspond to well-defined theoretical bidding functions. Although we find improvements in R^2 with added degrees in the polynomial, for exposition purposes, we find it useful to report the best fitting models in Appendix 2 and theoretically supported models in Table 2.[6] Appendix 2 reports the best fitting models for these settings, which are cubic for open-bid full and quadratic for sealed-bid private.

For Open-Bid Private and Sealed-Bid Full, we do not have clear theoretical predictions and so we report the best fitting cubic polynomial models here. We find (reported in Appendix 2) that treating cost as a non-linear term in these settings does not result in significant coefficients on the non-linear cost terms nor adds improvement to the fit in terms of R^2. This is to be expected since bidders can be thought of as bidding in the score space, which is the bid minus the cost. Indeed, in all cases, the estimated coefficient on cost is near 1.

In contrast to the linearity with respect to cost, captured by the estimated function $\alpha()$, the $\beta()$, $\gamma()$, $\phi()$, and $\xi()$ functions appear to be highly non-linear and are generally cubic with the exception of sealed-bid private, which appears to be quadratic.

Moving on to analyzing estimated bid functions in Table 2, we note that in the linear model of the Open-Bid Full auction (column 2 of Table 2), the intercept is positive and significant. This is contrary to the theoretical prediction of zero intercept. A likely reason for this positive intercept is that bidders deviate from the dominant strategy by placing jump bids, where a jump bid is defined as a bid in the score space that exceeds the previously highest score space bid by more than the minimum increment. The average jump amount (difference between current and previous bid) in the Open-Bid Full treatment is 19.6. This jump amount declines over time within an auction. A regression of jump amounts, taking into account the period, the number of bids submitted, and time elapsed, finds that every second elapsed in the auction decreases the jump amount by 0.03 and this is significant at the 5% level. Thus, jump bidding is not independent of the time the bid is submitted.

Table 2 Estimated Bid Functions Corresponding to a Best-Fitting Polynomial. Dependent Variable: Price Bid

Parameter	Open-Bid Full (losing bidders)	Open-Bid Full (all bidders)	Sealed-Bid Private	Sealed-Bid Full	Open-Bid Private
Intercept	24.11** (3.164)	−14.16** (6.760)	4.463* (2.458)	9.5105** (3.588)	−6.403** (2.6527)
C_i	1.051** (0.0476)	0.8311** (0.0388)	1.010** (0.0278)	0.9352** (0.0179)	0.7577* (0.0200)
$(Q_i - C_i)$		0.5489** (0.1487)	0.0696** (0.0093)	0.2138** (0.0602)	0.1553** (0.0576)
$(Q_i - C_i)^2$		−0.0040** (0.0011)		−0.0014** (0.0005)	−0.0011** (0.0004)
$(Q_i - C_i)^3$		1.02e-05** 2.25e-06		3.50e-06** (9.76e-07)	4.56e-06** (9.54e-07)
$Avg(B_{(-i)})$		0.4359** (0.0489)			0.4734** (0.0361)
$Avg(B_{(-i)})^2$		−0.0008** (0.0001)			−0.0004** (0.0001)
$Avg(B_{(-i)})^3$		7.00e-07** (1.09e-07)			3.57e-07** (9.58e-08)
$Avg(B_{(-i)}) \times LB$		−0.1948** (0.01518)			−0.2923** (0.0169)
$q^{(1)}$				−0.1431** (0.0416)	
$q^{(1)2}$				0.0010** (0.0003)	
$q^{(1)3}$				−2.05e-06** (8.07e-07	
$q^{(2)}$				−0.1138** (0.01599	
$q^{(2)2}$				0.0011** (0.0001)	
$q^{(2)3}$				−2.43e-06** (3.94e-07)	
R^2	0.3939	0.5654	0.5060	0.8016	0.7877

$*p < 0.1$; $**p < 0.05$.

Despite significant jump bidding activity, the resulting average buyer surplus is not statistically different from the theoretical buyer surplus (Table 1). However, jump bidding is possibly responsible for the lower efficiency indicated in Table 1 (see Kwasnica and Katok 2007, who report similar jump bidding behavior in forward auctions).

The second treatment for which we have theoretical benchmarks is the Sealed-Bid Private treatment. In theory, the intercept should be 0, and in the model we fitted (column 4) the intercept is only slightly above 0 and insignificant at the 5% level. Likewise, the value of the C_i coefficient should in theory be 1, and the estimated coefficient is not significantly different from 1. The $\beta()$ function should theoretically be linear, with the (Q-C) coefficient equal to 0.25. When we restrict the model to be linear, we observe the (Q-C) coefficient of 0.070, which is significantly below 0.25. The linear model in Table 2 has a slightly worse fit than the quadratic model (Appendix 2), which shows that $\beta()$ is non-linear in the (Q-C) term. The quadratic model is difficult to interpret, but we can tell from the linear model that the bidding behavior is more aggressive than it should be due to the small linear coefficient in the $\beta()$ function, resulting in a buyer surplus that is significantly higher than the theoretical prediction, and also higher than the average buyer surplus in the Open-Bid Full treatment. Thus, we reject Hypothesis 1. The overly aggressive bidding behavior is consistent with risk aversion (see Kagel [1995] and references therein), but several recent studies found that sensitivity to regret provides a richer explanation (see Engelbrecht-Wiggans and Katok [2008] and references therein).

We do not have theoretical predictions for the Sealed-Bid Full treatments, except a qualitative one that says that bidders should pay attention to quality levels of their competitors. Indeed, we observe (column 5) that bidders lower their bids based on their quality disadvantage toward the highest and the second-highest quality competitor—consistent with Hypothesis 2. The $\beta()$ function is non-linear, and is similar to the $\beta()$ function in the Sealed-Bid Private treatment, but unlike the Sealed-Bid Private treatment, the intercept is positive and significant.

Lacking theoretical guidance, we analyze bidding behavior in the Open-Bid Private treatment by fitting Equation (5) to that data (column 6), as well as to the Open-Bid Full data (column 3). The two sets of estimates are (surprisingly) similar in that: (i) both intercepts are negative and significant; (ii) both C coefficients are fairly high, but significantly below 1; and (iii) the $\phi()$ and $\xi()$ estimates are non-linear, and again, estimated coefficients are quite close. Where the two Open-Bid treatments differ is in the estimates of the $\beta()$ function—the linear term is much larger in the open-bid Full than in the Open-Bid Private treatment. This difference in $\beta()$ accounts for the significantly higher buyer-surplus level in the Open-Bid Private than that in the Open-Bid Full treatment. The form of the bidding function is consistent with Hypothesis 3: the linear $\phi()$ coefficient is positive and large, whereas the linear $\xi()$ coefficient is negative.

Another pattern in comparing the two open-bid treatments is that bidders place fewer bids with private information than with full information. In Figure 3 we plot the frequency of the number of bids placed by a bidder in an auction for the two open-bid treatments in our study. Unlike the Open-Bid Full treatment, the number of bids per auction in the Open-Bid Private treatment is extremely low (with the mode at 1). It appears that bidders in the Open-Bid Private auction place very few bids, so it may be that some aspects of their bidding behavior are similar to bidding behavior in sealed-bid auctions.

In Table 3 we present pair-wise comparisons of buyer surplus levels in the four treatments. Here the unit of observation is a cohort and in Table 3 we report the differences between the treatment in the row (Treatment X) and in the column (Treatment Y),

Figure 3 Frequency of Bidding per Auction per Person

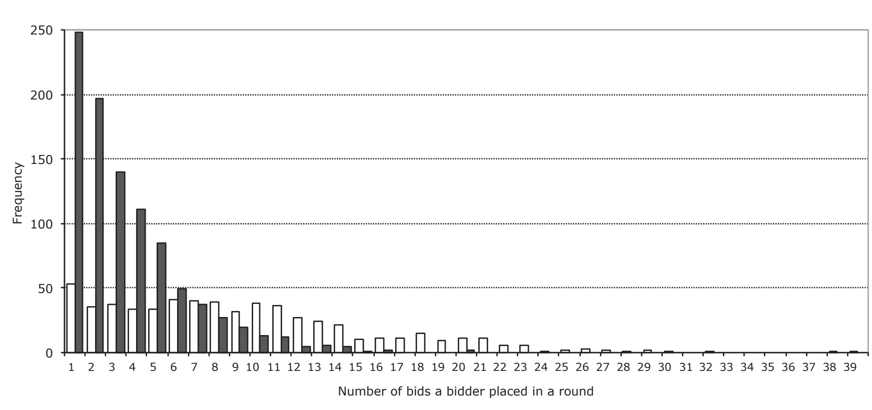

Table 3 Pair-Wise Difference in Buyer Surplus Levels

	Treatment Y		
	Sealed-Bid Private	Open-Bid Private	Sealed-Bid Full
Treatment X			
Open-Bid Full	−38.61**(4.98)	−25.74**(7.40)	−19.88**(9.00)
Sealed-Bid Full	−18.73**(8.43)	−5.86 (10.04)	
Open-Bid Private	−12.87**(6.70)		

*$p < 0.05$; **$p < 0.01$.

Table 4 The Effect of Sealed-Bid and Private Information Factors; Estimates of Equation (6)

Parameter	Variable	Empirical estimation from the data (Eq. (6))
β^{OBFull}	–	186.11** (2.78)
β^{SB}	SB_t	19.77* (3.93)
β^{PR}	PR_t	25.73** (3.93)
$\beta^{(SB \times PR)}$	$SB_t \times PR_t$	−7.02 (5.56)
R^2		0.865

*$p < 0.05$; **$p < 0.01$

and the standard error of the difference. We repeat the same analysis using log surpluses in Appendix 3 (comparisons are unaffected by the log transformation).

We conclude this section by summarizing the three main results as they relate to our three research hypotheses: (i) We can reject Hypothesis 1 because the average buyer surplus under the Sealed-Bid Private treatment is significantly higher than the average buyer surplus under the Open-Bid Full auction. (ii) The data support Hypothesis 2—bidders lower their bid in response to their own quality disadvantage relative to the highest and the second highest competitor quality. (iii) The data partially support Hypothesis 3, since bids in the Open-bid Private auction are affected by the bids of the competitors.

Returning to our original research question, of how price visibility and quality transparency affect buyer's surplus, we measure the effect of these two factors by fitting a linear regression model to our entire dataset, using a cohort as a unit of analysis:

$$\Pi_t = \beta^{OBFull} + \beta^{SB} SB_t + \beta^{PR} PR_t + \beta^{(SB\times PR)}(SB_t \times PR_t) + \varepsilon_t \tag{6}$$

The intercept term β^{OBFull}, and its standard error, measure the average buyer surplus in the Open-Bid Full information treatment. Variable SB_t is equal to 1 if auction t is in one of the sealed-bid treatments and zero otherwise. Variable PR_t is equal to 1 if auction t is one of the private information treatments and zero otherwise. We also included the interaction variable between the sealed-bid and private factors. We report estimates of the model in Equation (6) in Table 4.

Below is the summary of our results as they pertain to the effect that price visibility and quality transparency have on the buyer surplus given the cost and quality draws in our laboratory experiment. Note that these effects cannot be due to different random draws in the different treatments. As explained earlier, all random draws took place prior to the experiments, and were kept exactly the same across treatments:

1. The absence of price visibility increases buyer surplus by 19.77.
2. The lack of quality transparency increases buyer surplus by 25.73.
3. Interaction effect between price visibility and quality transparency is not significant.

Revealing less information seems to be better for buyers, whether this information is about the price visibility or about the quality transparency, which we can think of as a winner-determination criterion.

5. Discussion and Conclusions

Buyer-determined procurement auctions are widely used in industry in both open- and sealed-bid formats. In a laboratory experiment, we consider the effect of the price visibility (open- vs. sealed bid) as well as the effect of quality transparency on the performance of these buyer-determined auction formats. In theory, the expected buyer surplus levels should not differ between the formats of Open-Bid Full and Sealed-Bid Private. However, we find that in the laboratory average buyer surplus levels are generally not the same in the two corresponding experimental treatments. The actual buyer surplus is close to the theoretical prediction in the Open-Bid Full even though there is a fair amount of jump bidding. In the Open-Bid Private treatment, we find that average buyer surplus levels are higher than those that theory predicts.

What the Open-Bid Private, the Sealed-Bid Full, and the Sealed-Bid Private treatments have in common is that bidders do not know whether the bid they are about to place will win (they do not know their winning status). In the two sealed-bid treatments, not knowing the winning status is due to the lack of price visibility. In the Open-Bid Private treatment, it is due to the lack of quality transparency. Specifically, in the Open-Bid Private treatment, if bidders knew which bidder had the leading score in the auction at any given time and could observe all bids, then with a single switching point it would be possible to deduce the exact quality of the leading score bidder. However, bidders cannot observe who is the highest score bidder in the auction while the auction is taking place. Not knowing the winning status produces a sealed-bid "flavor" in the Open-Bid Private auction.

What the Open-Bid Private and Open-Bid Full formats have in common is that timing is possibly a part of the strategy space for bidders. For example, bidders could respond to every competing bid that comes in at above their cost, they could bid aggressively at the beginning, they could wait to bid near the end of the auction, or they could bid intermittently. In the Open-Bid Full treatment, these timing strategies do not affect the theoretical prediction for what the optimal bid would be. Given that there are neither monetary costs nor opportunity costs to bidding, the optimal bid conditional on bidding is always to bid the minimum increment below the current winning score. It was interesting to observe that bidders did not exactly follow this optimal strategy and occasionally placed jump bids, as we discussed in section 4, and these jumps were related to the timing of the bid. It was also interesting that these deviations did not have a major impact on surplus.

The Sealed-Bid Private treatment serves as a good benchmark, because we can test established theoretical predictions for bidding behavior in this treatment. We observe that the bidding in this treatment is somewhat aggressive relative to the risk-neutral prediction. This finding in itself is not new and is generally documented in the forward auction literature. For example, it may be consistent with risk aversion (Cox et al. 1988), regret (Engelbrecht-Wiggans and Katok 2007), impulse-balance equilibrium (a type of a learning explanation, Ockenfels and Selten 2005), or quantal response equilibrium (Goeree et al. 2002), as well as possibly other causes. Instead of focusing on a behavioral explanation, we offer a practical implication: Auctions that have a sealed-bid flavor (bidders do not know their winning status) are likely to result in bidding that is more aggressive than auctions in which bidders know their winning status and therefore have a weakly dominant strategy.

One managerial implication is that purchasing managers facing the choice between open-bid and sealed-bid auctions should prefer sealed-bid auctions as these tend to generate higher levels of buyer surplus in our experiments, due to overly aggressive bidding. This result nicely complements earlier findings in the literature that suggest that sealed-bid auctions are better for the buyer-supplier relationship and suppliers prefer them as well (Jap 2003, 2007). Consequently, it seems that sealed-bid auctions offer a superior format for both sides of the market. When open bidding is required (perhaps due to the dynamic nature of pricing or to government regulations), our laboratory findings indicate that it is best to leave bidder qualities unknown or ambiguous. That is, buyers should avoid revealing the identities of the bidders to the extent possible, and avoid posting differentiating features of the various bidders on the auction site or set individual reservation prices for different suppliers. Note that, as discussed in the introduction, the buyer does not always have a choice in the matter. If the industry is concentrated and the set of suppliers is fixed and well known, the buyer may not have a great deal of choice in what suppliers know about each other. However, it is our casual observation that the buyers do try to add relatively unknown suppliers to the mix of bidders and this could be a strategy intended to keep bidders guessing.

One direction for future research is to further investigate the effect various degrees of price visibility have on auction outcomes. In addition to the two extreme cases we looked at in this article, there are other partial price visibility formats currently used in industry. For example, sometimes bidders bid during a live auction event, but during the auction they know only the current low bid, but not the entire bid history (Jap 2007). Alternatively, bidders may know the rank of their price bid (whether it is the lowest, the second lowest, and so on) but are not told any of the bid amounts (see Elmaghraby et al. [2012] who report on a laboratory investigation of this setting).

Another direction for future research is to investigate the effect of ambiguity regarding the winner selection criteria. Sometimes buyers themselves are not certain about the exact trade-offs between non-price attributes (Elmaghraby 2007), which implies that bidders cannot be certain about their own quality. In an extreme case, this uncertainly may facilitate tacit collusion: if bidders believe the winner will be selected using a process that is, essentially, random, not bidding below the reservation price may be in equilibrium.

Finally, for dynamic bids in open-bid auctions, it is important to address the timing decisions of bidders (Bradlow and Park 2007, Park and Bradlow 2005). Our analysis, in accordance with common practice, examined final bids only. However, the timing of the bid is a critical component of the bidding strategy and one worth investigating in future research.

Acknowledgments

The authors gratefully acknowledge support from the Smeal College of Business at the Pennsylvania State University and the Naveen Jindal School of Management at the University of Texas at Dallas. Elena Katok gratefully acknowledges support from the National Science Foundation [Award 0849054]. We thank John Kagel, Tim Salmon, Mark Isaak, and John Duffy for helpful feedback.

Appendix 1

Characterization 2. In First Price Sealed Bid BD Auction with Full Information, Bids Depend on Observed Qualities

(a) The unique score-to-bid mapping is shown by the parallel of the score space to the price space. This mapping was computed by Engelbrecht-Wiggans et al. (2007) and we present it from bidder 1's perspective:

$$(Q_1 - B_1) = 3/4(Q_1 - C_1) \tag{A1}$$

(b) Engelbrecht-Wiggans et al. (2007) found that the mapping from score to bid is unique. Hence, in such equilibrium, if it exists, bidder 1 may solve the problem under the assumption that all other bidders are entering the bid prescribed in (a). We show that given this assumption, combined with bidder 1's explicit knowledge of the others' qualities, it is not the best response for bidder 1 to enter the bid prescribed by (a) as well.

Let P^U denote the unconditional probability of bidder 1 winning $P(Q_1 - B_1 > \max\{Q_2 - B_2, Q_3 - B_3, Q_4 - B_4\})$. Let P^C denote the parallel conditional probability $P(Q_1 - B_1 > \max\{Q_2 - B_2, Q_3 - B_3, Q_4 - B_4 \mid Q_2, Q_3, Q_4\})$. Computing the first order conditions of the maximization of bidder 1's profit, $(B_1 - C_1)$ P(bidder 1 wins), for the unconditional case (Sealed-Bid Private) and the conditional probability, and making them equal (as is required by uniqueness) gives us Equation (A2). The proof requires showing the equality in Equation (A2) cannot hold under the unique score-to-bid mapping.

$$\frac{P^{C'}}{dP^{C'}/dB_1} = \frac{P^U}{dP^U/dB_1} \tag{A2}$$

The proof rests on the simple fact that (A1) and (A2) are two equations with one unknown (B_1) and so there is no solution for general parameter values.

Appendix 2

Models of the Bidding Functions, with Bid as the Dependent Variable

Model fit indicated by R^2. The * indicates the best-fitting model relative to adding one more degree in the polynomial, based on an F-test. Sealed Private fits best with a quadratic bidding function. All others fit best with a cubic function.

	Open Full	Sealed Private	Open Private	Sealed Full
Linear	0.394	0.506	0.773	0.736
2nd degree polynomial with linear cost term	0.394	0.518*	0.783	0.791
2nd degree polynomial in all terms	0.465	0.518	0.785	0.792
3rd degree polynomial with linear cost term	0.394	0.519	0.788*	0.802*
3rd degree polynomial in all terms	0.502*	0.519	0.790	0.804
4th degree polynomial in all terms	0.513	0.520	0.790	0.807

For open-bid with full information and sealed-bid with private information we estimate the following polynomial models:

$$B_{it} = Intercept + \sum_{d=1}^{D} \alpha_d (C_{it})^d + \sum_{d=1}^{D} \beta_d (Q_{it} - C_{it})^d + \eta_i + \varepsilon_{it}$$

We report the best fitting models: *Cubic* (D = 3) for Open-Bid Full, *quadratic* (D = 2) with linear cost term for Sealed-Bid Private. The best-fitting models for Sealed-Bid Full and Open-Bid Private are reported in Table 2.

Parameter	Open-Bid Full (losing bidders) 3rd degree	Sealed-Bid Private 2nd degree with linear cost term
Intercept	-25.004^{**} (8.164)	15.2448^{**} (3.0612)
C_i	1.6764^{**} (0.4555)	1.0129^{**} (0.0274)
C_i^2	-0.0139 (0.0106)	–
C_i^3	8e-5 (7e-5)	–
$(Q_i - C_i)$	1.0755^{**} (0.1708)	-0.1441 (0.0375)
$(Q_i - C_i)^2$	-0.0087^{**} (0.0013)	0.0007^{**} (0.0001)
$(Q_i - C_i)^3$	$2e\text{-}5^{**}$ (3e-6)	–
R^2	0.502	0.518

$^{*}p < 0.05$; $^{**}p < 0.01$.

Comparison of the log surplus:

	Treatment *Y*		
	Sealed-Bid Private	Open-Bid Private	Sealed-Bid Full
Treatment *X*			
Open-Bid Full	-0.222^{**} (0.029)	-0.161^{**} (0.042)	-0.125^{**} (0.046)
Sealed-Bid Full	-0.097^{**} (0.041)	-0.036 (0.051)	
Open-Bid Private	-0.061^{*} (0.035)		

$^{*}p < 0.05$; $^{**}p < 0.01$.

Production and Operations Management 0(0), pp. 1–17, © 2012 Production and Operations Management Society

Appendix 3: Experimental Instructions

Compendium Instructions

Overview

You are about to participate in an experiment in the economics of decision making. If you follow these instructions carefully and make good decisions you will earn a considerable amount of money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and I will answer it. We ask that you not talk with one another for the duration of the experiment.

On your desk you should have a check-out form, a pen and two copies of the consent form.

How to Make Money

You are a supplier of some product. To make money you sell that product to a computerized buyer. You are competing against three other participants in the role of supplier. To sell the product, you offer the buyer a price; this offer is *your bid*.

Before you make *your bid*, you will find out your *cost* of supplying one unit, and the quality of your product. [Full information treatments: You will also find out the qualities of the other three supplier, but not their costs.]

The value of your product to the buyer depends on the difference between your *quality* and *your bid*. We call this value the *Buyer Net Value*. The formula for *Buyer Net Value* is

$$\text{Buyer Net Value} = \text{Quality} - \text{Bid}$$

The supplier with the highest *Buyer Net Value* wins and supplies one unit to the buyer.

You will bid in *40* rounds. Your *cost* in each round is an integer from 0 to 100. Your *quality* in each round is an integer from your *cost* to your *cost + 300*. For example, your cost in one round is 50, your quality in that round is an integer from 50 to 350, each integer in that range equally likely. Your cost and quality, as well as the costs and qualities of the other participants will be determined the same way and will change each round.

You make money by winning at a good price. If you do not win in a round, your profit for the round is zero. If you win, then your profit is

$$\text{Your Profit} = \text{Your Bid} - \text{Your Cost}$$

For example, if your cost is 50 and you win with a bid of 80, then your profit is

$$80 - 50 = 30$$

Caution: if your bid is below your cost and you win, you will lose money. Bid carefully.

The Mechanics Placing a Bid

Below is a sample snapshot of the screen you will be facing. As you can see under Your Information, you are given your Participant Number, your Cost and your Quality. [Full information treatments: On the top right hand corner of the screen, under the heading, "Bidder Qualities," you see the qualities of the four bidders, including yours. This allows you to determine your ranking on quality relative to the other bidders.]

To place a bid, you enter your bid amount in the Enter Bid box and then press the CALCULATE but-

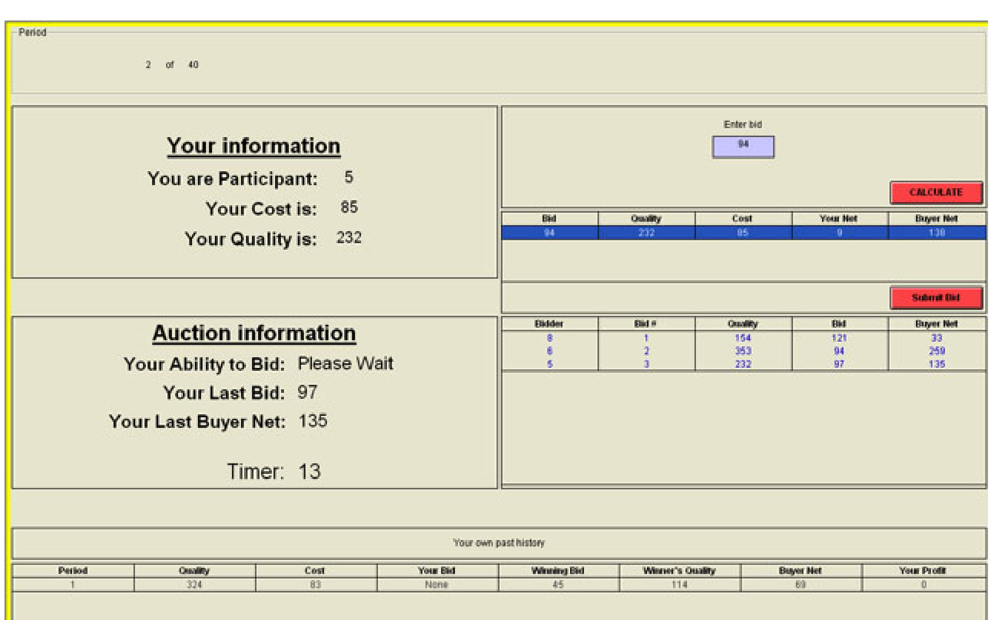

ton. The software will display calculations for your bid on the right of the bidding box. Let's go over these computations:

1 *Bid* = This is simply the bid you entered
2 *Quality* = This is your *Quality* as displayed in your participant information above.
3 *Cost* = This is your *Cost* as displayed in your participant information above.
4 *Your Net* = *Bid* − *Cost*. This is your profit if your bid wins.
5 *Buyer Net* = *Quality* − *Bid*. Your bid is ranked against the other three bids based on the Buyer Net. The bid with the highest Buyer Net wins the auction.

You can use the CALCULATE tool for any number of bids you wish. When you are ready to submit your bid, highlight it and click the SUBMIT BID button.

[Sealed-bid treatments: The auction ends when all four bidders have submitted their bids.]

[Open-bid treatments contained the following section:

The Dynamics of Bidding

1 Every time someone enters a bid, you will see that bid on your screen in the table under the Submit Bid button.
2 The entries in each row of that table are: (a) The bidder ID. (b) The bid number which indicates the number of bids entered so far. (c) The quality of that bidder. (d) The bid entered. (e) The buyer net this bid generates. Remember that the bidder with the highest buyer net wins.
3 You can enter as many bids as you wish, with the following exceptions:
   i Your bids must give a positive buyer net.
   ii Each bid you enter must be lower than the previous bid you entered.
   iii You can only revise your bid if someone else has entered a bid after your previous bid (in other words, the software will not let you outbid yourself).
4 The clock will count down from 60 seconds to 0 seconds. When someone enters a bid, if there are fewer than 10 seconds remaining to the end of the auction, the clock will reset to 10 seconds.
5 The auction ends when no bids have been placed for 10 seconds.]

Information You Will See at the End of Each Round

At the end of each round you will see a list of the four bids that have been placed and the corresponding buyer net values those bids generated.

Remember that the bid that generates the highest buyer net value wins.

You will also see the following information:

- The Winning Bid and the Quality of the Winner
- The Buyer's Net Value from the winning bid.
- The Buyer's Net Value from the second best bid
- Your Bid, Your Quality, and Your Cost.
- MLOT—money left on the table. This is the difference between the highest bid that would have allowed you to win the auction and your actual bid. That is, it is money you could have made but did not.
- Your Profit for the round.

On the bottom of each screen you will also see this information for all previous rounds. Below is a sample screen.

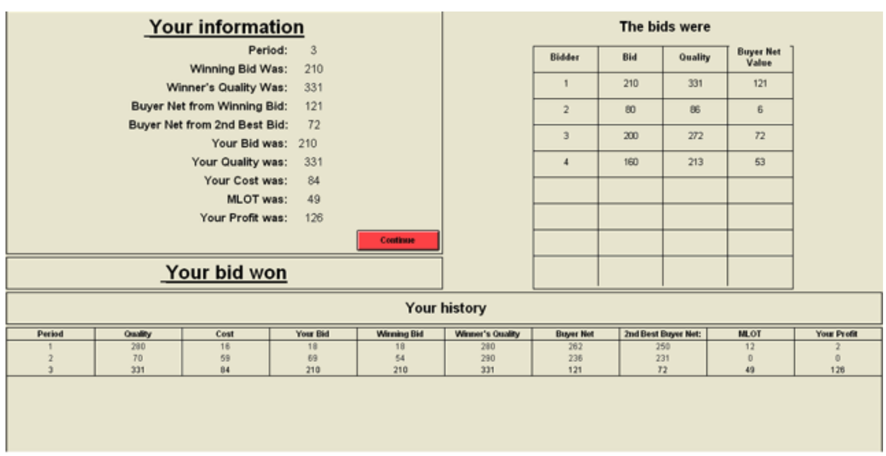

How You Will be Paid

At the end of the session, the computer will calculate the total profit you earned in all rounds and will convert it to US dollars at the rate of 5 cents per token. Your dollar earnings will be added to your $**5** participation fee and displayed on your computer screen. Please use this information to fill out the check-out form on your desk. All earnings will be paid in cash at the end of the session.

Notes

[1]Other studies challenge revenue equivalence between open bid and sealed bid forward auctions. Lucking-Reiley (1999) found that the Dutch format generates higher revenue than the sealed-bid format in Internet auctions. The open bid (English) auction is known to yield lower revenues than the second-price auction, as bidders tend to bid their valuations in the English (Coppinger et al. 1980, Kagel et al. 1987), but higher than their valuations in the second-price sealed bid format (Kagel and Levin [1993], Cooper and Fang [2008], but see exception in Garratt et al. [2012]. Thus, laboratory results show the dynamic open bid auction raising less revenue than the equivalent sealed-bid auction. First-price auctions consistently earn higher revenues than do second-price auctions (Coppinger et al. [1980], Cox et al. [1982], and Kagel and Levin [1993]). The result tends to diminish when the number of bidders becomes large, or when bidders' private values are affiliated (Kagel and Levin 1993).

[2]This is known as the "soft close" and is used to prevent sniping (see Roth and Ockenfels 2002).

[3]We repeated the same analysis for the last 10 periods, to check whether the actual buyer surplus in the Sealed-Bid Private treatment is closer to the theoretical buyer surplus toward the end of the session. It continues to be significantly lower ($p = 0.0004$). Average buyer surplus in the Open-Bid Private treatment continues to be in line with theoretical buyer surplus at the end of the session ($p = 0.984$).

[4]For the Open-Bid Full Information treatment we estimate bid functions for losing bidders only because losing bidders have the weakly dominant strategy of bidding down to their cost, one bid increment at a time.

[5]Strictly speaking, bids are sequential, so each bid should depend only on the bids that preceded it (so this is not endogeneity in the pure technical sense—the kind that would have clear econometric fixes), but nevertheless bids are endogenous choices to the bidders and this makes them problematic as independent variables. The main goal for our average bid specification is to improve robustness without sacrificing fit. When we enter the three competitors' bids, ordered from highest to lowest, in the regression with separate coefficients for the three ordered bids, as opposed to restricting them to a single coefficient, the R^2 is nearly identical to when we use the average competitors' bid and the sign of the single coefficient is highly robust to specification. That means that by the F-test we cannot reject a restricted model with the same coefficient on each competitor's bid. This is the same as multiplying that single coefficient by the sum of bids, or in our case—the average (which is the sum of bids rescaled by the number of bidders).

[6]We also found evidence that looking at log-price as the dependent variable can potentially improve the fit as well. However, this transformation has no mapping to any of the theoretical benchmarks and so we do not report it herein. Lastly, with sufficient data, one would be able to provide non-parametric kernel regression estimates which would provide the exact shape of the bidding function. This is outside the current scope of the article and the data we currently have is insufficient for such an attempt.

References

Athey, S., J. Levin, E. Seira. 2004. *Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions*. Stanford University, Palo Alto, CA.

Bajari, P., G. Lewis. 2011. Procurement contracting with time incentives: Theory and evidence. *Quart. J. Econ.*, **126**(3): 1173–1211.

Balestrieri, F. 2008. Essays on mechanism design, chapter 3. PhD Dissertation Thesis, MIT, Cambridge, MA.

Bapna, R., W. Jank, G. Shmueli. 2008a. Price formation and its dynamics in online auctions. *Decis. Support Syst.* **44**: 641–656.

Bapna, R., W. Jank, G. Shmueli. 2008b. Consumer surplus in online auctions. *Inf. Syst. Res.* **19**(4): 400–419.

Beall, S., C. Carter, P. L. Carter, T. Germer, T. Hendrick, S. D. Jap, L. Kaufmann, D. Maciejewski, R. Monczka, K. Petersen. 2003. *The role of reverse auctions in strategic sourcing*. CAPS Research, Tempe AZ.

Bradlow, E. T., Y.-H. Park. 2007. Bayesian estimation of bid sequences in internet auctions using a generalized record breaking model. *Mark. Sci.* **26**(2): 218–229.

Branco, F. 1997. The design of multi-dimensional auctions. *Rand J. Econ.* **28**(1): 63–81.

Cantillon, E. 2008. The effect of bidders' asymmetries on expected revenue in auctions. *Games and Econ. Behav.* **62**: 1–25.

CAPS. 2006. CAPS Research: Focus on eProcurement. CAPS: Center for Strategic Supply Research, Tempe, AZ.

Che, Y.-K. 1993. Design competition through multi-dimensional auctions. *Rand* **24**(4): 668–680.

Cooper, D., H. Fang. 2008. Understanding overbidding in second price auctions: An experimental study. *Econ. J.* **118**: 1572–1595.

Coppinger, V. M., V. L. Smith, J. A. Titus. 1980. Incentives and behavior in English, Dutch and sealed-bid auctions. *Econ. Inq.* **18**(1): 1–22.

Cox, J. C., B. Roberson, V. L. Smith. 1982. Theory and behavior of single object auctions. V. L. Smith, ed. *Research in Experimental Economics*. JAI Press, Greenwich, CT, 1–43.

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

Dholakia, U. M., I. Simonson. 2005. The effect of explicit reference points on consumer choice and online bidding behavior. *Marketing Sci.* **24**(2): 206–217.

Elmaghraby, W. 2007. Auctions within e-sourcing events. *Prod. Oper. Manag.* **15**(4): 409–422.

Elmaghraby, W., E. Katok, N. Santamaria. 2012. A laboratory investigation of rank feedback in procurement auctions. *Manuf. Serv. Oper. Manage.*, in press.

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

Engelbrecht-Wiggans, R., E. Katok. 2008. Regret and feedback information in first-price sealed-bid auctions. *Manage. Sci.* **54** (3): 808–819.

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

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

Garratt, R., M. Walker, J. Wooders. 2012. Behavior in second price auctions by highly experienced eBay buyers and sellers. *Experimental Econ.* **15**(1): 44–57.

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

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

Jap, S. D. 2003. An exploratory study of the introduction of online reverse auctions. *J. Mark.* **67**: 96–107.

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

Jap, S., E. Haruvy. 2008. Interorganizational relationships and bidding behavior in industrial online reverse auctions. *J. Mark. Res.* **45**(5): 550–561.

Kagel, J. H. 1995. Auctions: A survey of experimental research. J. H. Kagel, A. E. Roth, eds. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 501–585.

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

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.

Kaplan, T., S. Zamir. 2011. Asymmetric first-price auctions with uniform distributions: Analytic solutions to the general case. *Econ. Theor.*, forthcoming

Katzman, B., K. A. McGeary. 2008. Will competitive bidding decrease Medicare prices? *South. Econ. J.* **74**(3): 839–856.

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

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

Lucking-Reiley, D. 1999. Using field experiments to test equivalence between auction formats: Magic on the internet. *Am. Econ. Rev.* **89**(5): 1063–1079.

Mares, V., J. Swinkels. 2010. On the analysis of asymmetric first price auctions. Working paper, Kellogg School of Management, Georgia Tech, Atlanta, GA.

Maskin, E., J. Riley. 2000. Asymmetric auctions. *Rev. Econ. Stud.* **67**(3), 413–438.

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

Park, Y-H., E. T. Bradlow. 2005. An integrated model who bids and whether, when, and how much to bid in internet auctions. *J. Mark. Res.* **42**(4): 470–482.

Rangan, V. K. 1998. *Freemarkets Online*. Harvard Business School Publishing, Boston, MA, Case #598109, 1–20.

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. *Am. Econ. Rev.* **92**(4): 1093–1103.

Shugan, S. M. 2005. Marketing and designing transaction games. *Mark. Sci.* **24**(4): 525–530.

Tunca, T. I., D. J. Wu, F. V. Zhong. 2008. An empirical analysis of price, quality, and incumbency in service procurement auctions. Working paper, Georgia Tech, Atlanta, GA.

Wan, Z., D. Beil. 2009. RFQ auctions with supplier qualification screening. *Oper. Res.* **57**(4): 934–949.

Wan, Z., D. Beil, E. Katok. 2012. When does it pay to delay supplier qualification? *Theory and experiments. Manage. Sci.*, in press.

Wang, S., W. Jank, G. Shmueli. 2008. Explaining and forecasting online auction prices and their dynamics using functional data analysis. *J. Bus. Econ. Stat.* **26**(2): 144–160.

Zeithammer, R. 2006. Forward-looking bidding in online auctions. *J. Mark. Res.* **43**(3): 462–476.

Zhong, F. V., D. J. Wu. 2009. E-Sourcing: Impact of non-price attributes and bidding behavior. Working paper, Georgia Tech, Atlanta, GA.