

Auctions of Homogeneous Goods with Increasing Returns: Experimental Comparison of Alternative “Dutch” Auctions

Elena Katok

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

Alvin E. Roth

Department of Economics, Harvard University, and Harvard Business School, Boston, Massachusetts 02163, al_roth@harvard.edu

Most business-to-business (B2B) auctions are used to transact large quantities of homogeneous goods, and therefore use multiunit mechanisms. In the B2B context, bidders often have increasing returns to scale, or synergies. We compare two commonly used auction formats for selling multiple homogeneous objects, both sometimes called “Dutch” auctions, in a set of value environments that include synergies and potentially subject bidders to the “exposure” and “free-riding” problems. We find that the descending-price auction, best known for its use in the Dutch flower auctions, is robust and performs well in a variety of environments, although there are some situations in which the ascending uniform-price auction similar to the one used by eBay better avoids the free-riding problem. We discuss the factors that influence each mechanism’s performance in terms of the overall efficiency, the informational requirements, the seller’s revenue, and the buyer’s profit.

Key words: multiunit auctions; procurement auctions; supply chain management; experimental economics

History: Accepted by William S. Lovejoy, operations and supply chain management; received May 2004. This paper was with the authors 3 months for 1 revision.

1. Introduction

In recent years the use of auctions and dynamic pricing for procurement and sales has been steadily increasing.¹ Much of the emphasis of e-sourcing is on creating ways for improving access to suppliers, facilitating competition, and driving down prices through the use of procurement Internet auctions. Because Internet auctions have lower participation and transaction costs and provide ready access to large markets, companies often use them to buy and sell excess inventory, as well as first-run goods and commodities (Lucking-Reiley 2000). Business-to-business (B2B) auctions are predominantly multiunit and often involve synergies. For example, a manufacturing firm that purchases raw materials may wish to buy everything from the same supplier to ensure uniform quality, save on transportation costs, reduce fixed setup costs, or simplify order tracking (Keskinocak and Tayur 2001). Another supply chain setting in which multiunit auctions with synergies play a critical role

is in procuring transportation services: Shippers typically wish to win a group of continuous lanes, or specific lanes that complement their existing network, and their costs dramatically depend on their ability to do so (Ledyard et al. 2002, Cantillon and Pesendorfer 2005). Perhaps the best-known example of multiunit auctions where synergies play a critical role are the FCC auctions for spectrum licenses, where synergies often arise from owning licenses in adjacent geographical areas or for adjacent frequencies (Rothkopf et al. 1998, Milgrom 2000, Krishna and Rosenthal 1996).

Our study focuses on two commonly used mechanisms for selling multiple homogeneous objects. Although both are sometimes called “Dutch,” the mechanisms are quite different. Perhaps the largest single institution specifically geared to selling multiple units of homogeneous goods is the flower auction in Aalsmeer, Holland, where approximately 4 billion flowers and 400 million plants worth well over a billion Euros are sold by about 8,000 different sellers annually. Close to 2,000 buyers trade in the auction every day, completing about 50,000 transactions between 6:30 a.m. and 10:15 a.m. (van den Berg et al. 2001). The mechanism used is sometimes called a *descending* or *reverse clock auction* because the price

¹ According to a recent report by AMR Research, the revenues from strategic e-sourcing tools have increased from \$191 million in 2000 to \$952 million in 2003 and are projected to increase to \$2.8 billion by 2005 (Kraus and Mitchell 2002).

descends until a bidder decides to accept the current price by stopping the auction and buying all or part of the lot at that price per unit for as many units as desired.² If a part of the lot remains, the price goes back up, and the process is repeated until either the entire lot is sold or the price reaches its minimum level. If the price falls to the minimum the remainder of the lot is destroyed.

The Internet auction site eBay sells multiple units using a different mechanism it calls “Dutch.” The eBay version is an ascending uniform-price auction for multiple units with a fixed end time, in which buyers bid by specifying a *price* and a *quantity*, and the units are sold to the bidders with the highest prices whose quantities add up to the amount available. All bidders pay the price of the bidder who bids the lowest winning price, and that bidder is not guaranteed the entire quantity demanded. (In case of tie bids, priority is given to the bidder who demanded the greater quantity, and when bids are tied in both price and quantity, the earlier bid wins.³)

The eBay mechanism guarantees that all winners pay the same price, but it does not guarantee a winning bidder the entire quantity on which he bids. In contrast, the descending Dutch auction does not yield a uniform price, but it does guarantee that the winning bidder receives the entire quantity demanded. That is, in the descending-price auction, a bidder who stops the clock obtains the full quantity he desires at the price at which he stopped the auction. Thus, the descending Dutch auction is a simple version of a package-bidding (or “combinatorial”) auction in that, for homogeneous goods, it allows bidders to submit bids on any relevant package of items, i.e., on packages of any size, from a single unit to the entire quantity available at the time of the bid (see, for example, Rothkopf et al. 1998 and Milgrom 2004 for a discussion of combinatorial auctions and package bidding). While the descending Dutch auction has received a good deal of attention in the economics literature, this

has focused almost exclusively on auctions of a single item,⁴ and as far as we know its package-bidding feature has not previously been studied.⁵

When bidders’ preferences include synergies, most simple mechanisms that do not allow package bidding subject bidders to the *exposure problem*. The exposure problem, in the case of homogeneous goods, occurs when a bidder bids on a specific (large) quantity at a price per unit that exceeds his unit value for a smaller number of units. In this case a mechanism, like eBay’s, that does not guarantee a winning bidder the desired quantity, can leave a bidder “exposed” to potential losses.⁶ Milgrom (2000) discusses the exposure problem in the general context of auctions for heterogeneous goods and points out that it can lead to substantially less aggressive bidding and, hence, lower revenue.

As we will discuss below, the package-bidding feature of the descending-price auction (i.e., the fact that bidders can bid on any quantity) implies that bidders never suffer from the exposure problem in auctions of a homogeneous good. Part of the descending Dutch auction mechanism’s success in the Aalsmeer Flower Auction may have to do with the fact that it protects buyers of large quantities from the exposure problem. The total number of buyers at Aalsmeer is close to 2,000 per year, but only (approximately) 50 buyers are responsible for about 50% of the total volume in guilders, while about 700 smaller buyers are responsible for less than 1% of the total volume (see van den Berg et al. 2001).

Other advantages of the descending auction mechanism include speed (over 50,000 transactions are completed in Aalsmeer daily in fewer than four hours) and the fact that it allows bidders to get “immediate gratification,” meaning that the winning bidder knows his winner status immediately, and in the case of multiunit auctions, he knows that he won exactly the number of units he wanted. Additionally, a feature of the descending Dutch auction that may be especially attractive in the procurement auction setting is

² Price is per stem, units are boxes of a certain number of stems, and there may be a minimum quantity of boxes required of any bid.

³ The third “Dutch” mechanism, but one we do not discuss in this paper, is often used for common stock tender offers, in which a firm buys back its own shares by announcing a price range and the number of shares it wishes to repurchase (Best et al. 1998). Shareholders interested in selling back their shares specify the lowest price within the range at which they are willing to sell, as well as the quantity. The firm selects the price at which it can repurchase the required number of shares, and all transactions are conducted at this price. The eBay “Dutch” auction considered here operates under different rules than the eBay auctions for single objects studied in Roth and Ockenfels (2002) and Ockenfels and Roth (2002, 2003), and data from multiple-unit auctions were not included in those studies.

⁴ For auctions of a single item, Vickrey (1961) observes that the descending-price Dutch auction is strategically equivalent to the sealed-bid auction, but significant differences are observed both in the laboratory and the field (see Footnote 10). Carare and Rothkopf (2004) describe a decision-theoretic model of a slow Dutch auction that can explain some of the differences.

⁵ Many studies of the descending-price auction focus on the case of a single unit for sale or multiple units in an environment in which no bidder desires more than one (see, for example, McCabe et al. 1990).

⁶ eBay has recently added a provision allowing the marginal bidder to withdraw a bid when the entire quantity desired is unavailable. This is presumably in recognition of the exposure problem and of eBay’s limited ability to enforce the completion of contracts.

that it does not reveal information about bidders' true valuations.⁷

In spite of its many attractive properties, the descending auction is not very common in practice.⁸ A reverse clock mechanism requires the knowledge of a good upper bound on the eventual purchase price to set the starting price on the clock. If this price is too low, the seller runs the danger of selling the item immediately and foregoing potential revenues, but if this price is set too high, the auction may take an unnecessarily long time. This requirement makes descending auctions more appropriate for selling commodity-like objects (flowers) than unique objects (art) and is less of an issue in the procurement version of these auctions (in which the auction is run by the buyer, the bidders are the sellers, and the low price wins) because a natural lower bound on the price (zero in the extreme) is more readily available. A potential disadvantage of fast-descending auctions is that they require bidders to assemble at a specific time and place. Internet B2B auction platforms, however, have recently started to experiment with slow-descending auctions (for example, Keskinocak and Tayur 2001 describe Fairmarket.com's "AutoMark-down" mechanism).

The potential downside of the descending-price Dutch auctions on which we specifically focus in this paper is that it can leave the bidders open to a different problem that has been discussed in the context of package bidding—the *free-rider* (or "threshold") problem (Milgrom 2000). This problem occurs when the efficient allocation requires a group of small bidders

(or *local* bidders, to use the FCC terminology, reflecting the geographic aspect of spectrum licenses) to coordinate to beat a big (or *global*) bidder who has a disproportionately high value for a large number of units. If enough local bidders purchase units early at high prices so that the remaining quantity is less than the global bidder requires, the global bidder stops being competitive, and any remaining local bidders could potentially transact at a very low price. This situation may create an incentive among the local bidders to wait for other local bidders to buy at high prices and, if too many of them wait too long, the global bidder may purchase the units even when this results in an inefficient allocation.

In this paper, we investigate the performance of two mechanisms of interest in supply chain management, the descending-price Dutch auction and eBay's ascending uniform-price "Dutch" auction⁹—two mechanisms specifically designed for, and used in practice for selling multiple homogeneous goods. We focus on situations in which bidder preferences include synergies because, as we mentioned earlier, synergies are often observed in the procurement contexts. We test the mechanisms in the laboratory on a set of environments that include the potential for the exposure and the free-rider problems, thus creating a substantial challenge for both mechanisms. Although the set of environments in our laboratory experiments is necessarily a limited one, we are nevertheless able to draw some conclusions and provide insight into the relative advantages of the different auction mechanisms. We find that the descending-price Dutch mechanism is quite robust and performs well even in situations in which the potential for free riding is present. The ascending mechanism, predictably, does not perform as well in environments in which the exposure

⁷ The only information it reveals is the winning price, and because bidders pay what they bid, this winning price is only a lower bound on the bidder's true valuation. Rothkopf et al. (1990) write, "People of our acquaintance with experience in conducting business are reluctant to reveal their true costs and valuations. They are strongly conditioned to keep this information confidential. ... [and] ... such conditioning will not normally lead bidders to err ... [because] ... a truth revealing strategy may give away valuable information. It could reveal to potential competitors the extent to which a firm's technology was superior. Most important, it could reveal to others with whom the firm must subsequently negotiate precisely how much it can yield" (pp. 102–103). So even though in the affiliated value-setting mechanisms that reveal a large amount of information have been shown to generate higher revenues in equilibrium (Milgrom and Weber 1982), whether these higher revenues are likely to be observed in practice is a matter of some debate.

⁸ Although a variation of the slow reverse clock descending mechanism may be an appropriate model for retailers and clearance sales (Lazear 1986). A well-known retail use of a slow Dutch auction explicitly is Filene's, which runs a slow reverse clock auction in the basement (called Filene's Basement). Items in the basement are discounted by 25% of the original price every two weeks. Other retailers use a form of a slow reverse clock auction also, and although they do not discount according to a prespecified schedule, they do discount merchandise if it has not been sold after a period of time.

⁹ Some of the other mechanisms used to sell multiple units include simultaneous and sequential ascending mechanisms, the share mechanism, and possibly others. The simultaneous ascending auction, for example, is used by the FCC to sell spectrum licenses (Milgrom 2000). Hausch (1986) compares simultaneous and sequential multiunit auctions from the seller's perspective, and Engelbrecht-Wiggans and Weber (1979) demonstrate that sequential ascending auctions can be quite inefficient when bidders have non-linear utility functions. The share auction (Wilson 1979) involves each bidder submitting a sealed bid that includes his entire demand schedule (or alternatively the quantity demanded at each possible price), and the price that clears the market is then chosen. We do not study simultaneous or sequential ascending auctions here. Like the eBay "Dutch" mechanism that we do study, they are vulnerable to the exposure problem. The share auction is not vulnerable to the exposure problem, but submitting the entire demand schedule may not be very practical in the procurement context. We concentrate on the eBay mechanism (before the modification, which does not require partially filled bids to be transacted) because, of the mechanisms that exhibit the exposure problem and are practical for procurement, it appears to be the one most often used to sell homogeneous goods.

problem exists. More surprisingly, it sometimes allows the exposure problem to appear in situations in which it would not be an issue if all players played equilibrium strategies of the one-time auction. We also identify environments, with a very pronounced free-rider problem, in which the ascending auction produces efficient outcomes more often than the Dutch auction.

Related Literature

The descending auction mechanism has not been previously studied in the context of selling multiple homogeneous items in environments with complementarities. Lucking-Reiley (1999) conducted a field experiment comparing revenue in four auction mechanisms (English, first price sealed bid, second price sealed bid, and descending Dutch). The objects for sale in the field experiment were Magic Cards (which can be viewed as semihomogeneous objects) and a large number of cards were auctioned off in simultaneous auctions. Lucking-Reiley (1999) found that, contrary to the prediction of revenue equivalence (Vickrey 1961), revenues in the descending Dutch auction were higher than in the first-price sealed-bid auction. Revenue comparisons between the English and the second-price sealed-bid auctions were inconclusive, and efficiency comparisons are not possible in the field experiment setting.¹⁰

Various ascending and sealed-bid auction mechanisms for selling homogeneous objects have been studied theoretically. Ausubel and Crampton (2002) showed that the sealed-bid uniform-price auction mechanism is vulnerable to the “demand reduction” problem—potentially reducing competition and having a negative impact on the seller’s revenue. They also showed that often the uniform-price auction does not perform as well as a mechanism that permits price discrimination. Engelbrecht-Wiggans and Kahn (1998a) show that in pay-your-bid multiunit auctions there is a tendency for bids on different units to be close even when valuations for these units are not the same. Of course, there may be practical reasons to prefer a uniform-price mechanism, because price discrimination is often viewed as being unfair. Engelbrecht-Wiggans and Kahn (1998b) characterize equilibria that involve demand reduction for the uniform-price multiunit auctions when bidders have either increasing or decreasing returns. Krishna and

Rosenthal (1996) examine the simultaneous ascending auctions with synergies and heterogeneous bidders and find that the presence of global bidders leads to less-aggressive bidding. Rothkopf (1977) and Palfrey (1980) study multiobject auctions where bidders have a constraint on exposure.

A major focus of the empirical literature on auctions for multiple homogeneous units has been the so-called *declining-price anomaly*. When identical units are sold sequentially in practice, units sold earlier tend to be purchased at higher prices. This phenomenon appears robust to the auction format or the type of object being sold. Ashenfelter (1989) documents the declining-price anomaly in ascending auctions for wine. van den Berg et al. (2001) document the same phenomenon in descending “Dutch” auctions for roses at Aalsmeer. They also report that “at any round, the decline is stronger if the number of remaining units is smaller.” This finding is suggestive of increasing returns.

List and Lucking-Reiley (2000) use a field experiment to study demand reduction in uniform-price and Vickrey auctions for multiple units. Both mechanisms allow bidders to submit multiple sealed bids, but under the uniform-price rule, when m units are for sale, the top m bids win, and all units are sold at the price equal to the highest rejected bid. Under the Vickrey rules, the top m bids win, but the price for the k th unit equals the k th highest rejected bid (so under Vickrey rules the units are sold at different prices). List and Lucking-Reiley sell pairs of sports cards (football, basketball, and baseball) in two-person markets where each bidder would like to purchase both cards and find that, consistent with the theory, demand reduction is higher in the uniform-price auctions relative to the Vickrey auctions. They also find that, contrary to the theory, the bids on the first unit are higher in the uniform treatment than in the Vickrey treatment, and there are no significant differences in revenue between the two mechanisms. The field experiment setting does not provide a way for comparing efficiency. Kagel and Levin (2001) also find significant demand reduction in the laboratory in both sealed-bid and ascending auctions.

The experimental study most closely related to ours is by Kagel and Levin (2004). They compare the performance of the sealed-bid uniform-price and the ascending uniform-price multiunit auctions when bidders have synergies between the units. They conduct the auctions in four different environments, three with potential for the exposure problem. Two units are sold, and each market involves either four or six bidders: one human bidder with synergies (the *global bidder*) faces three of the four environments in each auction, and either three or five computerized bidders who demand a single unit (*local bidders*) only have

¹⁰ Cox et al. (1982) report that in laboratory experiments, contrary to the Lucking-Reiley (1999) result, the descending-price “Dutch” auction generates lower revenues than the sealed-bid auction. Katok and Kwasnica (2003) demonstrate that the clock speed in a descending Dutch auction has an effect on revenue in laboratory experiments—fast clock mechanisms generate lower revenues than slow clock mechanisms do—and offer a theoretical explanation based on a model of impatient bidders.

their valuations drawn from a uniform distribution and always follow the strategy of bidding up to their values. Kagel and Levin find that the ascending mechanism generally works better than the sealed bid in terms of the frequency of the optimal bidding behavior and participants' profit, and the revenues are consistently higher in the sealed-bid auctions. A major finding is that subjects in ascending auctions tend to bid too timidly in response to the exposure problem, but this does not happen in sealed-bid auctions. Bidders do learn to bid correctly in the environment without the exposure problem, and the learning is faster in ascending auctions than in the sealed bid.

Another related study by Fevrier et al. (2004) compares sequential auctions of two items with and without a "buyer's option" for the winner of the first auction to buy both units. They examine multiple auction formats and value environments, including increasing returns, in auctions with exactly two bidders, in which the free-rider problem can therefore never arise.

2. Design of the Experiment

Our experimental comparison of descending and ascending auctions is similar to Kagel and Levin's (2004) comparison of uniform-price sealed-bid and ascending auctions, except that to investigate behavior of local bidders (we will refer to them as *small bidders*) we use all human bidders. We also manipulate the small bidders' values to create environments with the free-riding problem.

Our design manipulates two factors: the auction mechanism and the distribution and correlation of bidder values. In all auctions three bidders compete for two units of an artificial asset. The global bidder (we will call him the *big bidder*) always has a high value for both units, but only if he wins both, and a low value for one unit. The two *small bidders* each want one unit of the asset and their values can be either high or low.

- In the *exposure* environment, $\frac{3}{4}$ of the time the big bidder's value for one unit is 20, and $\frac{1}{4}$ of the time it is 40. His value for two units is 100 or 120 (unit value of 50 or 60) with equal probability. The small bidders' values are correlated: $\frac{1}{3}$ of the time they are both low (either 35 or 45 with equal probability, independently drawn), and $\frac{2}{3}$ of the time one is high (65 or 75 with equal probability) and the other is low (35 or 45 with equal probability).

- The only difference between the exposure and the *free-riding* environments is that in the free-riding environment the two small bidders' values are either both low or both high—they are both low $\frac{1}{3}$ of the time and are both high $\frac{2}{3}$ of the time—conditional on being low or high, each bidder has an equal probability of having a value of 35 or 45 (if low) or 65

or 75 (if high). This implies that at equilibrium there is no danger of the exposure problem in the free-riding environment, because a big bidder who is outbid on one unit can expect to be outbid on both.

- The only difference between the free-riding and the *super-free-riding* environments is that the big bidder has a value of zero for one unit in the super-free-riding environment. This makes the free-riding problem worse for small bidders because the potential benefit from free riding successfully is higher.

All information about value distributions is public, and the roles, big or small bidder, do not change for the duration of the session. Big bidder wins in case of ties. We used the same randomly drawn values for all treatments in the same environment. Because value distributions were different in the exposure and the (super) free-riding environments, we used different draws for those environments. Values in the super-free-riding environment were the same as in the free-riding environment, but the big bidder's value for one unit was set to zero. The actual random draws were picked in a way that matched the announced distributions. Figure 1 summarizes the three environments in our experiment.

During each session participants are matched in groups of six (we call each group a *cohort*). Each cohort participates in a sequence of 20 auctions, with two separate groups of three bidders bidding in each round.¹¹ After each round the participants are randomly rematched within the cohort, in a way that no participant is matched with the same two participants for two consecutive auctions. The participants are told this.¹² The two auction mechanisms we study are the descending auction and an adaptation of eBay's ascending auction.

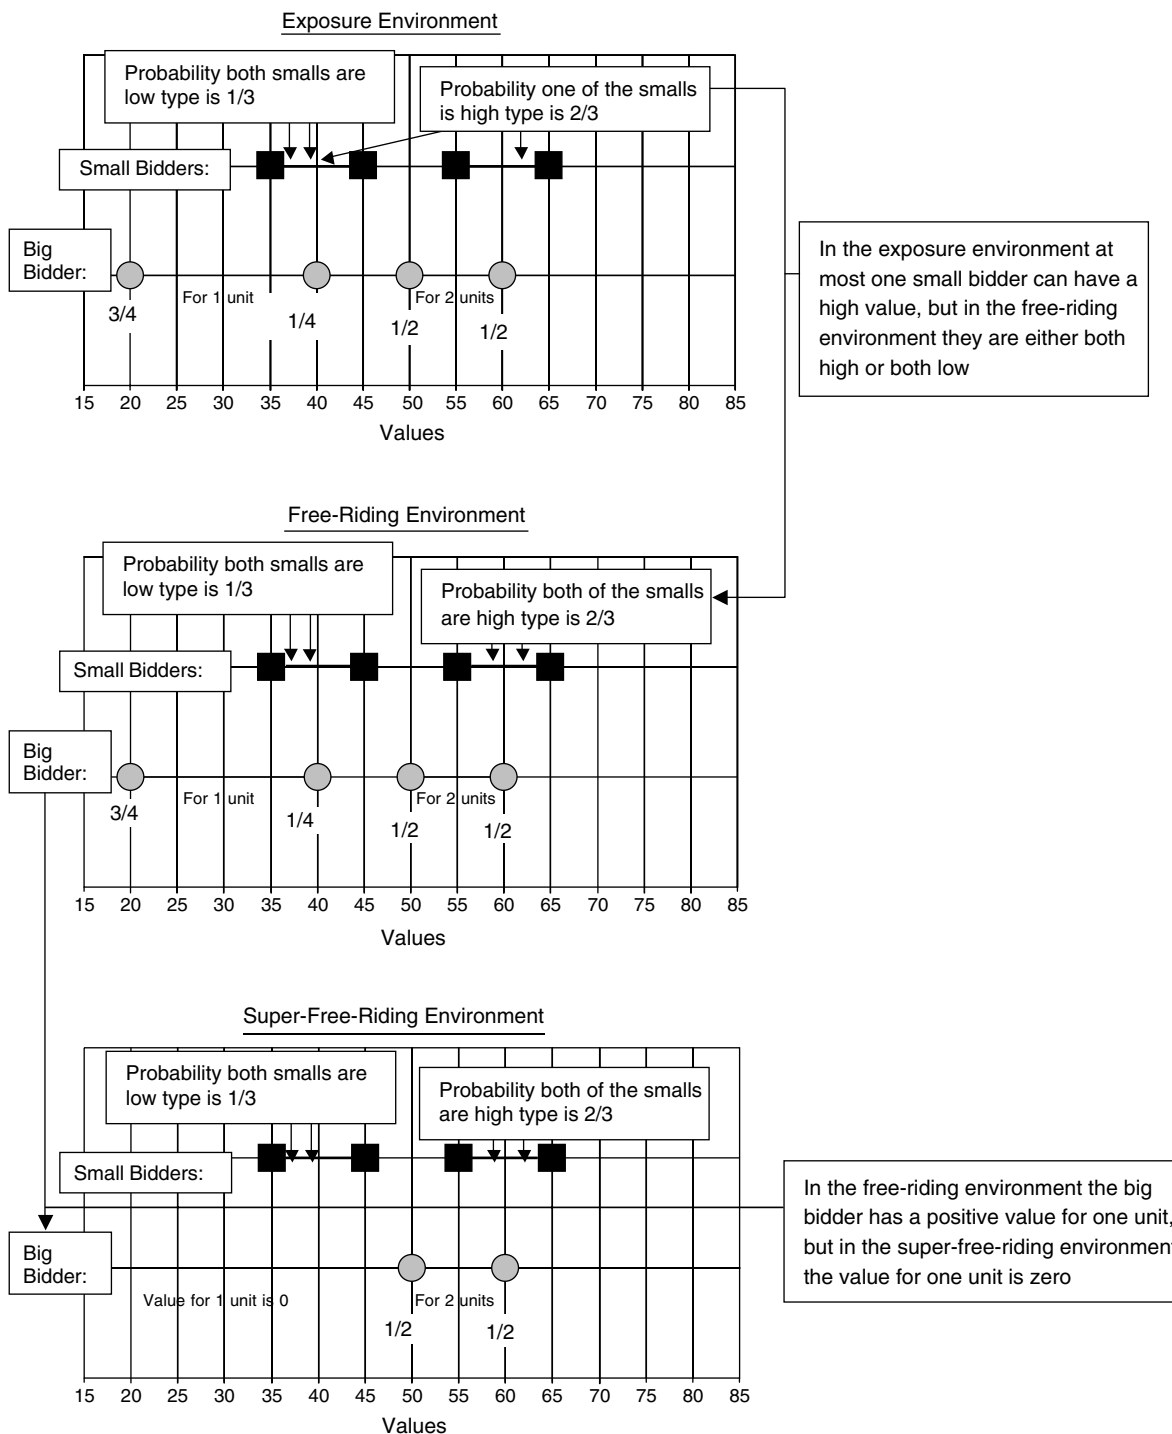
- In the *descending* mechanism the price starts at 85 and goes down by 0.5 every second.¹³ To purchase one or two units a participant clicks on a button. Big bidders have two buttons (for one and for two units) and small bidders have only one button. If a big bidder buys two units the auction ends. If a small bidder buys one unit (or if a big bidder buys one unit) the price of the remaining unit goes back up to 85 and the

¹¹ Each cohort of six subjects constitutes one independent observation (and we will be using this property in subsequent statistical analysis). There is a trade-off between having smaller cohorts and more independent observations versus larger cohorts where the same subjects are matched fewer times. In this case we chose to have smaller cohorts, although an argument can be made for either choice. In the context of procurement the same small set of bidders often compete repeatedly, and therefore small cohorts are consistent with the actual environment we have in mind.

¹² See complete instructions, including screen shots, at <http://mansci.pubs.informs.org/ecompanion.html>.

¹³ We introduced fractional tokens to decrease the probability of ties.

Figure 1 Distribution of Bidder Values in the Three Environments



process repeats.¹⁴ Once a bidder stops the clock (wins a unit) he cannot stop the clock again, and if only

one unit is available because one unit was already purchased, the big bidder receives an error message if he tries to purchase two units.

¹⁴ In the super-free-riding environment the price did not go back up to 85 but instead it continued to go down. This was done, following a pilot experiment, to save time because the remaining unit was always sold for a very low price, usually for 0.5 or 1 token.

- In the *ascending* mechanism the price starts at 10 and increases by 0.5 every second. To bid in the auction the bidders do nothing, but to stop bidding a bidder clicks a button. Once a bidder clicks a button

Figure 2 Experimental Design

Environment	Mechanism	
	Ascending	Descending
Exposure	5 cohorts = 30 subjects	5 cohorts = 30 subjects
Free-riding	4 cohorts = 24 subjects	5 cohorts = 30 subjects
Super-free-riding	4 cohorts = 24 subjects	4 cohorts = 24 subjects

Notes. The study contained 162 subjects. Each cohort (six subjects) participated in 20 rounds. Every round contained two separate auctions with three subjects in each auction. Each subject had the same role (big or small bidder) in every auction. No subject was matched with the same two subjects for two consecutive periods.

he becomes ineligible to bid again for the duration of the auction. The auction ends when the total number of units demanded by the eligible bidders falls to (or below) two. When this happens the remaining eligible bidder(s) win(s) units at the highest price that was reached before the auction ended, i.e., the price at which the last ineligible bidder dropped out. If the auction ended with demand by eligible bidders less than two, the remaining unit(s) goes to the bidder(s) who became ineligible last, at a price 0.5 below the selling price for the winning bidder (i.e., at the last price before the ineligible bidder dropped out). As in eBay, ties are broken based on quantity.

The exposure problem is realized when one of the small bidders becomes ineligible early, and the big bidder becomes ineligible next. In this case, the small bidder who did nothing wins one unit, and the big bidder wins the other unit and possibly loses money.

Figure 2 summarizes our experimental design, which varies both the auction type and the distribution of values (the value environment).

All sessions were conducted at Harvard Business School's Computer Laboratory for Experimental Research (CLER) between August 2001 and October 2001. Participants were recruited through flyers posted on billboards. Cash was the only incentive offered. Participants were paid their total individual earnings from the 20 auctions plus a \$10 show-up fee at the end of the session.¹⁵ The software was built using the zTree system (Fischbacher 1999). Sessions lasted about 90 minutes and average earnings were approximately \$25.

¹⁵ An additional \$5 payment was added to all big bidders' earnings in the ascending/exposure treatments and to all participants' earnings in the descending/exposure treatment. Participants in all treatments were informed that this fixed additional payment would be made in some treatments (see instructions for exact wording).

3. Theoretical Analysis

3.1. Structural Properties of Ascending and Descending Auctions for Homogeneous Goods

In this section, we present some simple theory of ascending and descending auctions for multiple units of a homogeneous good with increasing returns. The problems arising from increasing returns can be illustrated when there are two units for sale and at least three bidders. (The generalization to more goods and bidders is straightforward because the good is homogeneous, so only the number of units is important in determining the value.)

We start with notation. Assume there are three bidders, B , s_1 , and s_2 . Let $v_B(2)$ be the value of two units to B , $v_B(1)$ the value of one unit to B , and $v_{s_i}(1)$ and $v_{s_i}(2)$ the value of one and two units, respectively, to bidder s_i , $i = 1, 2$.

Suppose that B has increasing returns to scale, i.e., $v_B(2) > 2v_B(1)$, where $v_B(k)$ is the total value to B of k units. And suppose that s_1 and s_2 want only a single unit: $v_{s_i}(2) = v_{s_i}(1)$. Bidder B faces the exposure problem when, in attempting to win two units at a unit price greater than $v_B(1)$, he runs the risk of acquiring only one unit at that price and consequently making a loss. If B is deterred by this risk from bidding up to a unit price of $v_B(2)/2$ for two units, then the auction may result in an inefficient outcome whenever $v_{s_1}(1) + v_{s_2}(1) < v_B(2)$, but the small bidders nevertheless win the auction due to B 's reluctance to bid up to $v_B(2)/2$.

Bidders s_1 and s_2 suffer from the free-rider or threshold problem whenever there is an opportunity to make a profit by bidding on one unit and winning it, but there is a possibility of making a larger profit by bidding less, hoping that the other small bidder wins one unit at a high price, thus reducing the unit value of the remaining unit to B , and allowing the second small bidder to win a unit at a lower price. If either small bidder is deterred from bidding near $v_{s_i}(1)$ for this reason, then the descending-price auction may result in an inefficient outcome whenever $v_{s_1}(1) + v_{s_2}(1) > v_B(2)$, but B nevertheless wins both units, because each small bidder hoped that the other would win one of the units first.

In the *ascending uniform-price auction*, B faces the exposure problem when the current bid price per unit exceeds $v_B(1)$, but is less than $v_B(2)/2$, and two units are available. If B makes a bid of e per unit for two units such that $v_B(1) < e < v_B(2)/2$, then if another bidder makes a bid for one unit at a price $f > [v_B(2) + e - v_B(1)]/2$, B 's best response will be not to bid any further. In this case, if no more bids are made, the auction will end with B winning one unit at a price of e and making a loss of $v_B(1) - e$. That is, in making the

attempt to win two units at the potentially profitable price of e such that $v_B(1) < e < v_B(2)/2$, B is exposed to the risk of winning only one unit and making a loss.

PROPOSITION 1 (EXPOSURE PROBLEM). *In a descending-price Dutch auction for multiple units of a homogeneous good, a bidder B who wishes to buy more than one unit never faces the exposure problem.*

PROOF. At any moment in the descending-price auction when more than one unit is still available, B can purchase up to the number available at the current price. So B never wins fewer units than he bid for. \square

PROPOSITION 2 (NO STRATEGIC EQUIVALENCE). *The descending-price Dutch auction for multiple units is not strategically equivalent to the uniform-price sealed-bid auction in which the lowest-winning bidder may win only part of the quantity he bid on; neither is it equivalent to the first-price sealed-bid auction (despite their strategic equivalence when only one unit is to be purchased).*

PROOF. A bidder who bids on more than one unit can suffer losses from the exposure problem in the uniform-price sealed-bid auction, but not in the descending Dutch auction. Once, e.g., bidder s_1 buys a unit in the descending-price auction, s_2 can wait to bid until the clock nears the big bidder's reservation price for one unit, whereas in the first-price sealed-bid auction he cannot condition his bid on that of the other small bidder. \square

However, while the descending price Dutch auction eliminates one problem faced by bidders in uniform-price auctions, it raises another.

In the descending-price Dutch auction, bidders s_1 and s_2 face the free-rider problem whenever $v_{s_i}(1) > v_B(2)/2$ for both $i = 1, 2$. For prices p such that $v_B(2)/2 < p < v_{s_i}(1)$, either bidder s_i could profitably stop the clock just before the price reaches $v_B(2)/2$, and win one unit at a profit of $v_{s_i}(1) - v_B(2)/2$. If one of the small bidders does so, then the unit value to the large bidder is immediately reduced to $v_B(1)$, allowing the second small bidder j to guarantee himself a profit of $v_{s_j}(1) - v_B(1) > v_{s_j}(1) - v_B(2)/2$. Consequently, it is more profitable to let the other small bidder stop the clock. If both bidders delay until the price has dropped below $v_B(2)/2$, then bidder B may profitably bid on both units. The outcome of the auction will be inefficient when $v_B(2) < v_{s_1}(1) + v_{s_2}(1)$ if the free-rider problem causes the small bidders to defer bidding until the large bidder jumps in.

PROPOSITION 3 (FREE-RIDER PROBLEM). *The free-rider problem cannot occur in the ascending uniform-price auction (or in any uniform-price auction).*

PROOF. Because the auction is uniform price, no small bidder can obtain a unit at a lower price than any other small bidder. \square

So far we have discussed only the *structural* properties of the auctions, i.e., the properties concerning which outcomes and strategies are feasible. We now consider equilibrium predictions for our specific laboratory environments.

3.2. Equilibrium Analysis

We introduce some additional notation. Let the price increment in an auction be δ , and let $\Pr(H)$ be the probability that at least one small bidder is a high type, so $1 - \Pr(H)$ is the probability that both are low types. Also, let P_t be the price at time t , $\Pr_t(B)$ be the probability that B stops the clock at time t , and $\Pr_t(s)$ be the probability that one of the small bidders stops the clock at time t . The small bidders are either of low type, with the value of $v_{iL}(1)$ for one unit, or high type, with the value of $v_{iH}(1)$ for one unit. Consider a setting with three bidders, B with the value of $v_B(2)$ for two units and $v_B(1)$ for one unit, and two small bidders. Suppose high-type small bidders always have per-unit valuations above B , $v_{iH}(1) \geq v_B(2)/2 + \delta$, and low-type small bidders always have per-unit valuations below B , $v_{iL}(1) < v_B(2)/2$. Suppose B 's value for one unit is $v_B(1) < v_{iL}(1)$. (For most of our analysis it will be sufficient to speak about the players' realized values; but in Proposition 6 we will have to pay closer attention to the fact that our experimental environment is a game of incomplete information with two potential high and low values for each player.)

Descending Mechanism. We provide some general results about equilibrium bidding behavior in descending auctions with heterogeneous bidders. Our laboratory setting described in §2 has no pure strategy equilibrium (see Proposition 6) and multiple mixed-strategy equilibria.¹⁶

PROPOSITION 4 (DOMINATED STRATEGIES). *There is no equilibrium in undominated strategies for the descending auction at which a risk-neutral bidder B bids above*

$$\bar{P} = \frac{v_B(2)}{2} - \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right) (1 - \Pr(H)).$$

At an equilibrium in undominated strategies, a high-type small bidder never bids above $\bar{P} + \delta$. If bidders are risk averse, then

$$\bar{P} > \frac{v_B(2)}{2} - \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right) (1 - \Pr(H)).$$

¹⁶ The computations of some specific examples of mixed-strategy equilibria is straightforward but tedious, and we make them available for interested readers at: <http://mansci.pubs.informs.org/ecompanion.html>.

PROOF. If B bids at $\max_i v_{iL}(1)$, then his expected profit is at least

$$(1 - \Pr(H)) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right),$$

so if B is risk neutral, he must earn at least as much as the expected value of bidding $\max_i v_{iL}(1)$ whenever he bids. Therefore, any price P at which B might bid must satisfy

$$\frac{v_B(2)}{2} - P \geq (1 - \Pr(H)) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right).$$

So,

$$\bar{P} = \frac{v_B(2)}{2} - \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right) (1 - \Pr(H))$$

if B is risk neutral and

$$\bar{P} > \frac{v_B(2)}{2} - \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right) (1 - \Pr(H))$$

if B is risk averse. \square

Note that to find $\bar{P} + \delta$, a small bidder has to determine big bidder's \bar{P} and then apply one more iteration of the dominance argument. For this reason fewer small bidders may be able to do these calculations successfully.

PROPOSITION 5 (LOWER BOUND). *There exists a price $\underline{P} \geq \min_i v_{iL}(1) - \delta$ such that at any equilibrium of the descending-price auction, the clock is stopped for the first time at some $P_t > \underline{P}$.*

PROOF. We prove this by induction. At $P_t = \delta$ big and small bidders stop the clock with certainty because if they wait, the auction ends and they earn zero, so $\Pr_t(B) = \Pr_t(s) = 1$. Because B wins in case of ties, B always wins at $P_t = \delta$. Now, suppose there is some price $P_t < \min_i v_{iL}(1)$ at which a small bidder stops the clock with certainty. In this case, B stops the clock with certainty at P_t also, because he makes a positive profit at the price P_t , because $v_B(2)/2 > \min_i v_{iL}(1) > P_t$. But if B stops the clock with certainty at P_t , then the small bidder will stop the clock with certainty at $P_{t-1} = P_t + \delta$ if he makes a positive profit at P_{t-1} , $v_{is}(1) - P_t - \delta > 0$, which is always true if $P_t < \min_i v_{iL}(1) - \delta$. \square

Proposition 5 implies that in equilibrium we should not see any prices below the minimum value of the low-type small bidder (in our experiment, this is 35).

PROPOSITION 6 (PURE STRATEGY EQUILIBRIUM). *Pure strategy equilibrium in the descending auctions can only exist if the proportion of high-type small bidders is sufficiently small so as to make the big bidder's strategy of never bidding above $\max_i v_{iL}(1)$ profitable. If*

$$\frac{v_B(2)}{2} - \max_i v_{iL}(1) - 2\delta > (1 - \Pr(H)) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right),$$

no pure strategy equilibrium exists.

PROOF. For simplicity, we first prove the proposition for the slightly simpler game in which the big bidder's realized value $v_B(2)$ is common knowledge, and then extend the argument to our experimental environment in which it is not. Suppose there exists a price P^* at which B bids with certainty. Then, a high-type small bidder with a value $v_{iH}(1)$ bids with certainty at $P^* + \delta$ if this is profitable (which it always would be at equilibrium because a high type's valuation is always above B 's valuation). Therefore, at P^* , if the auction has not ended, B knows with certainty that the small bidder has a low value, and therefore B has no reason to bid above $\max_i v_{iL}(1)$, the maximum value a low type may have. So, it cannot be an equilibrium for B to stop the clock for the first time with certainty at any $P^* > \max_i v_{iL}(1)$. Hence, one candidate for P^* is $\max_i v_{iL}(1)$. If $P^* = \max_i v_{iL}(1)$, then the high-type small bidder bids with certainty at $\max_i v_{iL}(1) + \delta$ (which is always profitable for the high type). But, if

$$\begin{aligned} & \frac{v_B(2)}{2} - \max_i v_{iL}(1) - 2\delta \\ & > (1 - \Pr(H)) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right), \end{aligned}$$

B prefers to bid at $\max_i v_{iL}(1) + 2\delta$ and win for certain instead of waiting until $\max_i v_{iL}(1)$, and therefore bidding at $P^* = \max_i v_{iL}(1)$ is not a pure strategy equilibrium strategy. If, on the other hand,

$$\begin{aligned} & \frac{v_B(2)}{2} - \max_i v_{iL}(1) - 2\delta \\ & \leq (1 - \Pr(H)) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right), \end{aligned}$$

then any

$$P_t \in \left[\min_i v_{iL}(1), \max_i v_{iL}(1) \right]$$

is a candidate for a pure strategy P^* at which the big bidder stops the clock for the first time. Whether a pure strategy equilibrium exists and, if so, the value of P^* depends on the distribution of the low-type small bidders.

Now consider the case of our experimental environment, in which $v_B(2)/2$ may equal either 50 or 60 with equal probability. The proof above has to be modified to consider a big bidder strategy of bids at $P^*(50)$ and $P^*(60)$. A high-value small bidder's best response will be a bid with certainty at δ over one of these prices. If the small bidder's strategy were to bid above the higher of these two prices, then if no such bid were forthcoming, the big bidder would know that the small bidder was a low type, and so neither of the two prices can be above $\max_i v_{iL}$, based on the argument above. If the small bidder's strategy were to bid just above the lower of the two prices, then a best response by the big bidder would be to set the higher

price to coincide with the small bidder's bid, which again yields a contradiction. So, the only pure strategy equilibrium strategies for the big bidder to stop the clock for the first time must be in the interval

$$P_i \in \left[\min_i v_{iL}(1), \max_i v_{iL}(1) \right],$$

and cannot exist except as described above. \square

In particular, no pure strategy equilibria exist in our experimental environment.

Ascending Mechanism. Obtaining equilibrium benchmarks in ascending auctions is much easier than in descending auctions because pure strategy equilibria always exist. Although there are multiple equilibria, small bidders have a dominant strategy of bidding up to their value and no further. So, at any equilibrium in undominated strategies, small bidders play their dominant strategy and the big bidder plays a best response. (There remain multiple equilibria because the big bidder is indifferent about when to stop bidding when he knows that he will ultimately lose the auction, but for most purposes, as far as allocations are concerned, we can regard the collection of equilibria in undominated strategies as a single equilibrium.) See the appendix for a detailed description of the computation of equilibrium in undominated strategies (i.e., for a computation of the big bidder's best reply strategy to the small bidders' dominant strategy).

PROPOSITION 7 (EQUILIBRIUM IN UNDOMINATED STRATEGIES IN ASCENDING AUCTIONS). (A) *In the free-riding and the super-free-riding environments, it is a dominant strategy for small bidders to bid up to their values. The best reply for big bidders in these experimental conditions is to bid up to $\max_i v_{iL}(1) + \delta$. This weakly dominates the strategy of bidding up to $v_{iB}(2)/2$, in which case B also earns zero if high types are present, but high types earn less. If high types are not present, bidding in undominated strategies never continues beyond $\max_i v_{iL}(1) + \delta$.*

(B) *In the exposure environment, it is still the dominant strategy for small bidders to bid up to their value, and the best response strategy of big bidders depends on their values and is described in detail in the appendix. We summarize the strategies here:*

(i) *Big bidders with the value of 20 for one unit and 50 per unit for two units should stop bidding between 20 and 35. Following this strategy ensures that the bidder never suffers losses from the exposure problem, but it also predicts that 33% of the auctions end in inefficient allocations.*

(ii) *Big bidders with the value of 20 for one unit and 60 per unit for two units should bid up to 35, and if neither of the two small bidders drops out at 35, drop out at 35.5. However, if one of the small bidders does drop out at 35, the big bidder should fully go for the two units, and not drop*

out at all (and suffer losses from "exposure" 33% of the time (when one of the small bidders is of a high type and the low-type small bidder has the value of 35). Auctions end in inefficient allocations 25% of the time in this case.

(iii) *Big bidders with the value of 40 for one unit and 50 per unit for two units should bid up to 40 (they are actually indifferent between stopping anywhere between 35.5 and 40). There are no losses associated with this strategy, and 25% of the auctions end in inefficient allocations.*

(iv) *Big bidders with the value of 40 for one unit and 60 per unit for two units should bid up to 40 and if one small bidder drops out at 35, then continue bidding up to 45.5, which implies that big bidders suffer losses from the exposure problem when one of the small bidders is a high type and the other has the value of 35, which happens 33% of the time. The resulting allocation is inefficient 58.33% of the time.*

Overall in equilibrium, big bidders suffer losses from the exposure problem about 16.5% of the time in our environment, and the allocation is inefficient 32.17% of the time.

PROOF. The proof involves straightforward calculations presented in the appendix. \square

3.3. Efficiency

Equilibrium analysis of the individual bidding behavior in the preceding section allows us to draw conclusions about aggregate measures of efficiency for the two mechanisms.

PROPOSITION 8 (INEFFICIENCY AT EQUILIBRIUM IN DESCENDING AUCTIONS). *In a descending auction and free-riding environment (i.e., in which small bidders are always of the same type) there are efficiency losses due to free riding. In a descending auction and exposure environment (i.e., in which small bidders are not always of the same type) there are smaller efficiency losses due to the absence of a pure strategy equilibrium.*

PROOF. In settings where there is no pure strategy equilibrium (Proposition 6), B wins with some positive probability when two high-type small bidders are present. \square

PROPOSITION 9 (INEFFICIENCY AT EQUILIBRIUM IN ASCENDING AUCTIONS). *There are efficiency losses in ascending auctions and exposure environments because big bidders sometimes suffer losses from the exposure problem when they follow the equilibrium strategy in ascending auctions. They are also sometimes deterred from bidding up to their value for two units.*

PROOF. In our setting, big bidders with the value of 60 per unit for two units make a higher expected profit if they bid above their value for one unit, $v_{iB}(1)$, than if they bid up to $v_{iB}(1)$. Big bidders have a positive probability of being forced to purchase one unit at a loss when they bid above $v_{iB}(1)$ in our exposure environment. \square

3.4. Hypotheses

To summarize the hypotheses these theoretical results suggest for our experimental environments, recall that $v_B(2)/2 =$ either 50 or 60; $v_{iH} =$ either 65 or 75; $v_{iL} =$ either 35 or 45; $\Pr(H) = \frac{2}{3}$; and $\delta = \frac{1}{2}$, so $\min_i v_{iL}(1) = 35$ and $\max_i v_{iL}(1) = 45$. By Proposition 6 we know that there is no pure strategy equilibrium for the descending auction, because for big bidders with the value of 60,

$$\begin{aligned} \frac{v_B(2)}{2} - \max_i v_{iL}(1) - 2\delta &= 60 - 45 - 1 = 14 \\ &> (1 - \Pr(H)) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right) &= \frac{1}{3}(60 - 45) = 5. \end{aligned}$$

For a big bidder with a value of 50,

$$\begin{aligned} \frac{v_B(2)}{2} - \max_i v_{iL}(1) - 2\delta &= 50 - 45 - 1 = 4 \\ &> 1 - \Pr(H) \left(\frac{v_B(2)}{2} - \max_i v_{iL}(1) \right) &= \frac{1}{3}(50 - 45) = \frac{5}{3}. \end{aligned}$$

HYPOTHESIS 1 (BIDDING IN DESCENDING AUCTIONS). *There is no bidding below 35 or above 55 by a risk-neutral big bidder (because to do so would be to play a dominated strategy) and above 55.5 by a small bidder. By Proposition 4,*

$$\bar{p} = v_B(2) - \frac{(v_B(2) - \max_i v_{iL}(1))}{(1 - \Pr(H))} = 60 - \frac{60 - 45}{3} = 55.$$

HYPOTHESIS 2 (BIDDING IN ASCENDING AUCTIONS). (A) *In ascending auctions small bidders should follow the undominated strategy of bidding up their values. For the purposes of comparing our data with existing laboratory data for single-unit ascending-bid auctions, however, we note that low-type small bidders are sometimes indifferent between bidding up to their values or stopping below their values.¹⁷*

(B) *Big bidders should bid up to 45.5 in the free-riding and super-free-riding environments (Proposition 7(A)) and follow the strategies outlined in Proposition 7(B) and the appendix in the exposure environment.*

¹⁷ Due to discreteness, the dominant strategy of bidding truthfully is not unique for low types. In the free-riding and the super-free-riding environments, low-type small bidders never win in equilibrium, and therefore any bid below the value is also a best reply strategy (by Proposition 7(A)). In the exposure environment, B should drop out at 20 only if $v_B(1) = 20$ and $v_B(2)/2 = 50$ (Proposition 7(Bi)), and otherwise he should bid up to at least 35.5. Consequently, a low-type small bidder with the value of 35 is indifferent between dropping out anywhere between 20 and 35. Similarly, in equilibrium, if a small bidder with the value of 45 observes the price above 35.5, he can conclude that B will not stop before 45.5 (Proposition 7(Bii/iv)), and therefore is indifferent between dropping out between 35.5 and 45.

HYPOTHESIS 3 (EXPOSURE AND EFFICIENCY IN ASCENDING AUCTIONS). *We say that B suffers losses from the exposure problem when B loses money by either purchasing one or two units at a loss, having been outbid after bidding for two units at a price above the unit value for one unit.*

(A) *In the exposure environment B suffers losses from the exposure problem about 16.5% of the time at equilibrium in undominated strategies, at which small bidders play their dominant strategy of bidding up to their values and big bidders follow a best response strategy (see Proposition 7(B) and the appendix).*

(B) *In the free-riding and super-free-riding environments B never suffers losses from the exposure problem.*

(C) *In the exposure environment, about 32% of ascending auctions end in inefficient allocations at equilibrium in undominated strategies (see Proposition 7(B) and the appendix).*

(D) *Ascending auctions are 100% efficient in the free-riding and super-free-riding environments.*

HYPOTHESIS 4 (EFFICIENCY COMPARISONS). (A) *In the exposure environment the efficiency of descending auctions is higher than the efficiency of ascending auctions.*

(B) *In the free-riding and super-free-riding environments the efficiency of ascending auctions is higher than the efficiency of descending auctions, and the efficiency of the ascending auction is the same (100%) in the free-riding and the super-free-riding environments.*

(C) *The efficiency of the descending auction is higher in the free-riding environment than in the super-free-riding environment.*

HYPOTHESIS 5 (REVENUE COMPARISONS). (A) *In the exposure environment the revenue in descending auctions is higher than the revenue of ascending auctions.*

(B) *In the free-riding and super-free-riding environments the revenue in ascending auctions is higher than the revenue of descending auctions, and the revenue of the ascending auction is the same in the free-riding and the super-free-riding environments.*

(C) *The revenue in the descending auction is higher in the free-riding environment than in the super-free-riding environment.*

4. Data Analysis

4.1. Individual Bidding Behavior and the Role of Information

We start by examining individual bidding behavior relative to equilibrium benchmarks summarized in Hypotheses 1 and 2. We analyze aggregate performance measures in §4.2.

The descending-auction bidding data are by and large consistent with Hypothesis 1. In Table 1 we summarize the ranges of prices that resulted from the first time the clock was stopped.

Table 1 The Ranges of Prices That Resulted from the First Time the Clock Was Stopped in the Descending Auction

Environment	Bidder who stopped the clock first	Number of occurrences	Below 35	Above 55	Above 55.5	Prices consistent with equilibrium prediction (%)
Exposure	Big	107	0	12		88.7
	Small	93	0		25	73.1
	Overall	200	0		37	81.5
Free-riding	Big	66	3	1		93.9
	Small	94	1		30	67.0
	Overall	160	4		31	78.1
Super-free-riding	Big	88	0	0		100.0
	Small	72	1		13	80.5
	Overall	160	1		13	91.9

Table 1 makes it apparent that Hypothesis 1 has somewhat more bite for big bidders than for small bidders, and this is not surprising in the laboratory setting because small bidders have to do one extra step of dominance reasoning.

We now examine the ascending-auction bidding data and compare it to the benchmarks in Hypothesis 2. At equilibrium big bidders in the free-riding environment under the ascending mechanism should bid up to 45.5. Under this strategy they always win two units when the two small bidders are low type and nothing otherwise. But the success of this strategy rests heavily on the assumption that high-type small bidders bid up to their value. Figure 3 summarizes the percentage of time small bidders in ascending auctions behave consistently with Hypothesis 2 (i.e., who either win the auction below their value, bid up to within one token of their values, or stop below their values when they have low values that cannot win the auction). Behavior of low-type small bidders is consistent with theoretical predictions in Hypothesis 2(A) 97% of the time,¹⁸ but behavior of high-type small bidders is less consistent. There was not a single instance of bidding above the value by high-type small bidders, but in the free-riding environment only 88% either win the auction or bid up to their value, which is substantially below the 97% in the super-free-riding environment.

We further investigate the reasons for this underbidding in Figure 4, which focuses on small bidders with high values who lose the auction and displays the distribution of their highest bids. In general, small bidders with high values should never lose the auction under the ascending mechanism in the free-riding environment, but in some cases this happens because big bidders continue to bid above their value

and eventually win two units at a loss. We see this happen seven times in both ascending/free-riding and ascending/super-free-riding treatments.

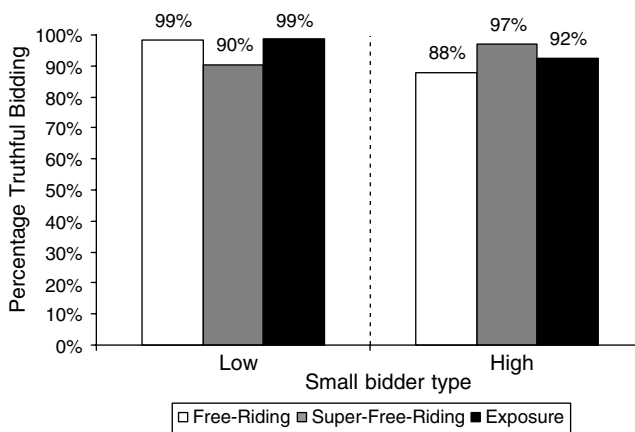
Of special interest is the number of times small bidders stop bidding between the values of 21 and 45.5. One way to interpret this behavior is that it is a deliberate attempt to “trap” the big bidder—make him think that both small bidders are low type—and force him to win one unit at a loss. There is a cost to doing this for the small bidder because he foregoes positive profit, but he may be deliberately incurring this cost to make the two big bidders in the cohort more conservative in future auctions¹⁹—or small bidders may simply be making mistakes or being lazy. Overall, small bidders with high values stop bidding between 21 and 45.5 18 times in the ascending/free-riding treatment, but only one time in the ascending/super-free-riding treatment. The difference is weakly significant (*t*-test *p*-value is 0.0910). Whatever the small bidders’ motive, this behavior does make big bidders more conservative in the free-riding environment: Big bidders bid up to 45.5 or above, consistent with Hypothesis 2(B), in 71% of the auctions in the free-riding environment and in 92% of the auctions in the super-free-riding environment. In the exposure environment big bidders’ optimal strategy is more complex and only roughly 50% of the big bidders follow it.²⁰

¹⁹ That is, although we have focused our analysis on the one-time auction, some bidders in our experiment may be responding to the fact that they will play future auctions with bidders drawn from the same pool, and hence have a positive probability of meeting each big bidder again.

²⁰ Percentage of auctions where big bidders follow the optimal strategy in the exposure environment: 50% (35 out of 70) of the big bidders of type 20/50 stopped on or before 20; of the 40/50 type big bidders, 50% (15 out of 30) bid up to 35 and stopped there; of the 40/60 type big bidders only 58% (11 out of 19) bid up to 45 and stopped there; of the 20/60 type big bidders 50% (40 out of 80) behaved consistent with the optimal strategy. Overall, in about 38% of the auctions, big bidders stopped bidding earlier than the optimal strategy required—thus they were deterred by the exposure problem.

¹⁸ There were three instances of bidding above the value by low types in the exposure environment (one won), four instances in the free-riding environment (one won), and 11 instances in the super-free-riding environment (zero won).

Figure 3 Ascending Auction



Note. Percentage of small bidders who either win the auction for a bid below their value, bid up to within one token of their values, or stop below their values when they cannot win the auction.

4.2. Aggregate Mechanism Comparison: Efficiency, Revenue, and Profit

Examining Hypothesis 3, in the exposure environment the ascending-auction equilibrium in undominated strategies predicts that 16.5% of the time big bidders should suffer losses from the exposure problem (Hypothesis 3(A)), and the actual percentage of time big bidders suffer losses from the exposure problem is almost twice as high—33%. In the free-riding and the super-free-riding environments big bidders should never suffer losses from the exposure problem (Hypothesis 3(B)). In fact, 17% suffer losses from the exposure problem in the free-riding environment and 6% in the super-free-riding environment. Overall, about 32% of the ascending auctions in the exposure environment should end in inefficient allocations if everyone were to follow equilibrium strategies (Hypothesis 3(C)), but the actual number is somewhat higher, at 42.5%. Auctions in free-riding

and super-free-riding environments should never end in inefficient allocation (Hypothesis 3(D)), but in fact 27% do in the free-riding environment and 7% in the super-free-riding environment. So, the data are not consistent with the point predictions corresponding to equilibrium (Hypothesis 3(A)–(D)). One noteworthy observation is that the efficiency of the ascending auction in the super-free-riding environment is much closer to the theoretical prediction than the efficiency of the ascending auction in the free-riding environment. Given individual bidding behavior we observed and discussed in the previous section, this difference is not surprising: Equilibrium benchmarks assume equilibrium behavior by small bidders (especially high-type small bidders) and we have noted that a sizable minority of high-type small bidders in the free-riding environments do not follow the equilibrium strategy of the isolated auctions.

Figure 5 classifies the causes of inefficient outcomes. Figure 5(a) confirms that the free-riding problem exists under the descending mechanism and is significantly higher in the environment in which the rewards from the free riding are high (the super-free-riding environment), but is also present, but to a smaller extent, in the free-riding environment. Figure 5(b) shows the extent of the exposure problem under the ascending mechanism. Big bidders are excessively conservative in the exposure environment, and also, interestingly, in the free-riding environment. That is, the results for the free-rider condition in Figure 5(b) suggest that big bidders may have sometimes feared losses from the exposure problem even though this would not happen if high-type small bidders always bid up to their values.

We now examine the differences in efficiency of the two auction mechanisms, as summarized in Hypothesis 4. Figure 6(a) compares the percentage of efficient

Figure 4 Distribution of Highest Bids By Small Bidders with High Values Who Lose the Auction in the Ascending/Free-Riding and the Ascending/Super-Free-Riding Treatments

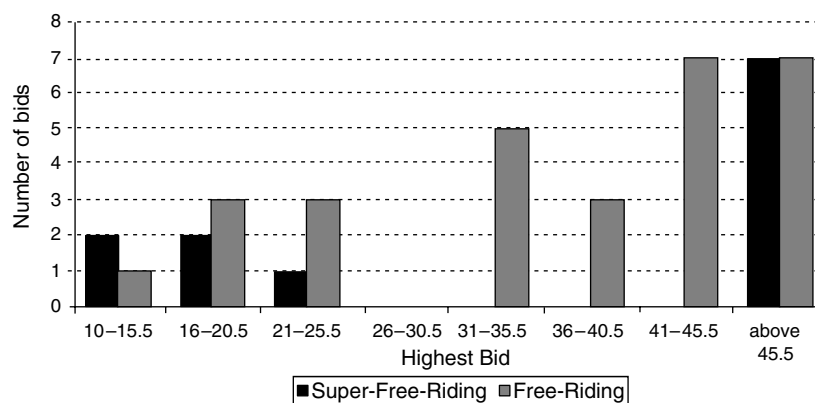
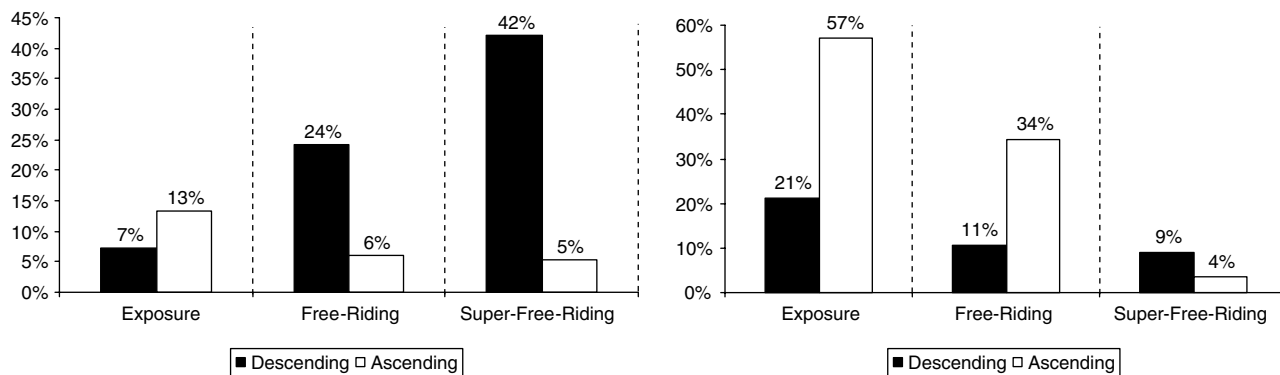


Figure 5 The Frequency of Inefficient Outcomes



(a) The big bidder wins when he should have lost

(b) The big bidder loses when he should have won

allocations over all 20 rounds for the two mechanisms, benchmarked using the efficiency of a random allocation. The descending mechanism is more efficient in the exposure environment (consistent with Hypothesis 4(A)), and the ascending mechanism is more efficient in the super-free-riding environment (partially confirming Hypothesis 4(B)), but contrary to Hypothesis 4(B), the two mechanisms are equally efficient in the free-riding environment. The differences in the exposure and the super-free-riding environments are highly significant (t -test p -values are 0.0011 and 0.0039) and the difference in the free-riding environment is not significant (t -test p -value is 0.1515).²¹ The efficiency is fairly constant over time in the exposure and the free-riding environments. In the super-free-riding environment the ascending mechanism becomes nearly 100% efficient after several initial periods. Consistent with Hypothesis 4(C), the descending auction is more efficient in the free-riding than in the super-free-riding environment (t -test p -value = 0.0096).

Figure 6(b) compares seller's revenue in periods 1–20. Consistent with Hypothesis 5(A), the descending auction yields higher revenue in the exposure environment (t -test p -value is 0.0222) and consistent with Hypothesis 5(B), the ascending revenues are higher in the super-free-riding environment (p -value is 0.0022). Contrary to Hypothesis 5(B), there is no significant difference between the mechanisms in the free-riding environment (p -value is 0.1567). Consistent with Hypothesis 5(C), the descending auction generates more revenue in the free-riding than in the super-free-riding environment (t -test p -value = 0.0389).

Table 2 summarizes average unit selling prices (top), their standard errors (second number in parentheses), the number of occurrences (third number), and the percentage of occurrences (bottom number in

parentheses) in the six treatments. In all the ascending treatments when the two small bidders win they pay the same price. In the exposure environment, the average purchase price of the small bidders when they both win is 23.68 tokens. This low price means that B often decides to leave the auction very close to the value for one unit (which is 20 75% of the time and 40 25% of the time). In other words, B cannot successfully compete in the exposure environment due to the exposure problem (the big bidders in the Kagel and Levin 2004 study behave similarly). When B wins only one unit in this treatment (this happened 63 times, 31.5% of the time), the average selling price is 44.6 tokens and B loses money. Because B is unable to compete successfully, the small bidders often win at very low prices, hence the low revenues in the ascending/exposure treatment.

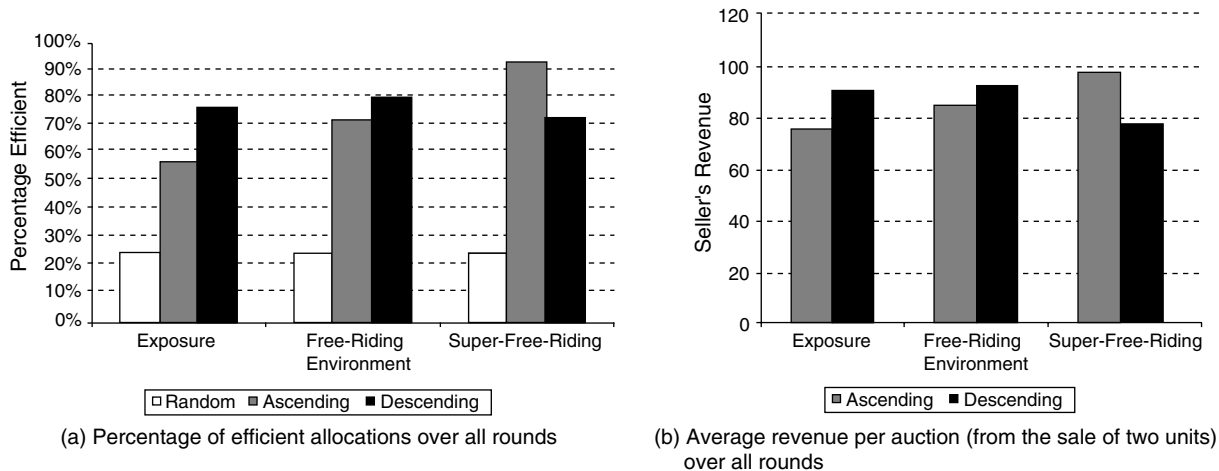
The descending mechanism eliminates the exposure problem for B 's and allows them to compete successfully. The average price B pays for two units in the exposure environment is very similar in the ascending and descending treatments (48.64 tokens versus 46.68 tokens). However, B wins two units more often in the descending/exposure treatment (107 times, 53.5% of the time versus 57 times, 28.5% of the time under the ascending mechanism), hence higher efficiency. When B only wins one unit in the descending/exposure treatment, the average purchase price is 30.54, and B 's never lose money. They win one unit when one of the small bidders has purchased a unit at a high price and B is able to successfully compete with the second small bidder.²²

As we mentioned in the introduction, an important difference between the ascending and the descending mechanisms is the ability of the descending mechanism to price discriminate. When the small bidders win in the descending/exposure treatments they on

²¹ Comparisons are done using a t -test for samples with unequal variance and the unit of observation we use throughout is a cohort.

²² The second small bidder in the exposure environment has a low value (35 or 45) and the big bidder's value for one unit is either 20 (75% of the time) or 40 (25% of the time).

Figure 6 Overall Efficiency Comparisons



average pay radically different prices (52.30 versus 24.43). They also have different values (one always has a high value of 65 or 75, while the other always has a low value of 35 or 45). So, in this case the seller is able to capture a great amount of the consumer surplus from the small bidder with a high value. This is in contrast to the ascending/exposure treatment when both small bidders end up paying 23.68 tokens on average, meaning that the small bidder with the high value gets to keep much of the consumer surplus.

Now let us examine the free-riding and the super-free-riding environments. First, note that the efficient

allocation is always the same in these two environments and never includes *B* winning one unit. In fact, *B*'s value for one unit should be irrelevant. Nevertheless, we do see *B* winning one unit under the ascending mechanism. This happens more often in the free-riding (40 times, 20% of the time) than in the super-free-riding (7 times, 4.4% of the time) environment, consistent with the results in §4.1. However, note that when *B* wins one unit in these environments, the average price is 43.62 in the free-riding and 48.43 in the super-free-riding environment. This highlights again that *B* does occasionally suffer from the

Table 2 Average Selling Prices, Their Standard Errors (Second Row in Parentheses), the Number of Occurrences (Third Row), and the Percentage of Occurrences This Number Represents (Third Row in Parentheses) in the Six Treatments

Environment	Auction type	Average price of a winning bidder			
		Small bidder 1*	Small bidder 2*	Big bidder (one unit)	Big bidder (two units)
Ascending					
Exposure	Avg. selling price	23.68	23.68	44.60	48.64
	(Std. error of price)	(4.98)	(4.98)	(3.86)	(7.52)
	# occur. (% occur.)	80 (40)	80 (40)	63 (31.5)	57 (28.5)
Free-riding	Avg. selling price	42.74	42.74	43.62	39.15
	(Std. error of price)	(7.94)	(7.94)	(6.56)	(4.07)
	# occur. (% occur.)	111 (55.5)	111 (55.5)	40 (20)	49 (24.5)
Super-free-riding	Avg. selling price	50.49	50.49	48.43	44.07
	(Std. error of price)	(4.52)	(4.52)	(4.97)	(5.25)
	# occur. (% occur.)	96 (50)	96 (60)	7 (4.4)	57 (35.6)
Descending					
Exposure	Avg. selling price	52.30	24.43	30.54	46.68
	(Std. error of price)	(2.86)	(3.68)	(5.92)	(2.93)
	# occur. (% occur.)	63 (31.5)	63 (31.5)	30 (15)	107 (53.5)
Free-riding	Avg. selling price	56.21	42.66	38.13	47.51
	(Std. error of price)	(5.06)	(7.11)	(8.99)	(3.87)
	# occur. (% occur.)	83 (51.9)	83 (51.9)	11 (6.8)	66 (41.3)
Super-free-riding	Avg. selling price	50.32	8.06	N/A	46.81
	(Std. error of price)	(4.63)	(8.42)		(2.54)
	# occur. (% occur.)	72 (45)	72 (45)		88 (55)

* Small bidder 1 is the first small bidder to stop the clock in a descending auction and small bidder 2 is the other small bidder. In the ascending treatments the two small bidders are labeled as 1 or 2 arbitrarily.

exposure problem even in the environments in which this problem would not arise at equilibrium.

The seller's ability to price discriminate actually has a negative impact on revenues in the super-free-riding environment (consistent with Hypothesis 5(C)). This is because after one small bidder wins one unit, B is completely out of the auction, so the second small bidder needs only to wait for the price to reach its minimum. We see this happening in the descending/super-free-riding treatment—one of the small bidders pays 50.32 tokens on average, while the second pays 8.06. We also see price discrimination in the descending/free-riding treatment, where one small bidder pays 56.21 tokens on average, while the other pays 42.66. However, in this case B is still somewhat competitive for one unit (with the value of 20 or 40), and this ensures that the second small bidder pays a higher price.

5. Discussion and Conclusions

Most B2B auctions are used to buy or sell large quantities of homogeneous goods, such as excess inventory, and therefore they use some form of a multiunit mechanism. In contrast to auctions that sell heterogeneous goods, such as the FCC spectrum auctions, there exists a simple mechanism that provides package-bidding capabilities in auctions for homogeneous goods. The descending-price Dutch auction is such a mechanism, and our experiments have shown its performance to be robust. Even though the laboratory environments in our study are highly artificial, we designed them with the specific purpose of challenging both mechanisms under investigation, and therefore we are able to draw conclusions that can be potentially useful to managers for deciding on the format of procurement auctions. When synergies exist, the descending auction is robust, even in situations with a potential for the free-riding problem. However, to conduct a descending auction the decision maker needs to know the upper bound on the price, because if the initial price is too low, the units will be sold immediately and the seller will lose revenue, but if the initial price is too high, the auction may last an unnecessarily long time. For this reason, descending auctions are most appropriate for handling commodity-like items for which markets are well established (for example, flowers, fish, or semiconductors). Descending auctions are likely to be less appropriate for market testing of new products, where the purpose of market testing is to better understand the demand structure for the product.

Another potential downside of the descending Dutch mechanism is that it price discriminates—all winners are not guaranteed the same price. In some contexts this feature may be viewed as unfair.

For example, when bidders in the same auction are also competitors in the larger marketplace, paying a higher price for a production input may put a firm at a competitive disadvantage. If this is the case, the free-riding problem may become exacerbated, making the descending Dutch mechanism less attractive.²³

Another important and surprising finding of our study has to do with the poor performance of the eBay ascending "Dutch" auction in environments with synergies. This mechanism performs better than the descending auction only in our super-free-riding environment. But the super-free-riding environment is, intentionally, extreme. The laboratory results demonstrate what it takes for the ascending auction to perform well when buyers have synergies: a large amount of credible information about competitors' willingness to pay. Without such information, there is the danger of the exposure problem appearing in practice, even when in theory it should not. Therefore, we conclude that eBay's "Dutch" auction will be susceptible to the exposure problem in many environments with synergies. Because eBay is primarily a business-to-consumer and consumer-to-consumer platform, this finding may not be highly relevant for eBay because consumers seem less likely to have large synergies. However, our finding is relevant to managers whose task is to design auctions for procurement and sales, when synergies exist, as they often do in B2B environments.

An electronic companion to this paper is available at <http://mansci.pubs.informs.org/ecompanion.html>.

Acknowledgments

The authors gratefully acknowledge the support from the eBusiness Research Center (eBRC), Smeal College of Business, Pennsylvania State University, the Computer Laboratory for Experimental Research at the Harvard Business School, and National Science Foundation grants to Pennsylvania State (Katok) and to the NBER (Roth). The authors thank John Kagel and Dan Levin for their useful comments on an earlier draft of this paper, and especially for encouraging them to provide equilibrium benchmarks. They also thank the associate editor and two anonymous referees.

Appendix

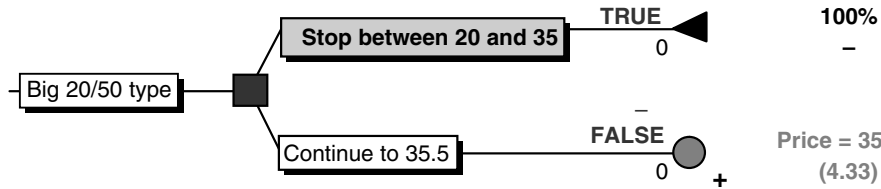
This appendix summarizes the equilibrium in undominated strategies in the ascending auction and the exposure environment. Table A1 summarizes equilibrium predictions for the four big bidder types, including a description of the big bidder's best reply strategies, the resulting allocation, the

²³ Ausubel (2002) proposes an ascending mechanism for selling multiple homogeneous items that implements the Vickrey outcome. In this mechanism, bidding truthfully is an equilibrium, so theoretically it is not susceptible to either the exposure or the free-riding problems. Like the descending Dutch auction, winners in the Ausubel auction pay different prices.

Table A1 The Summary of the Predictions for the Equilibrium in Undominated Strategies

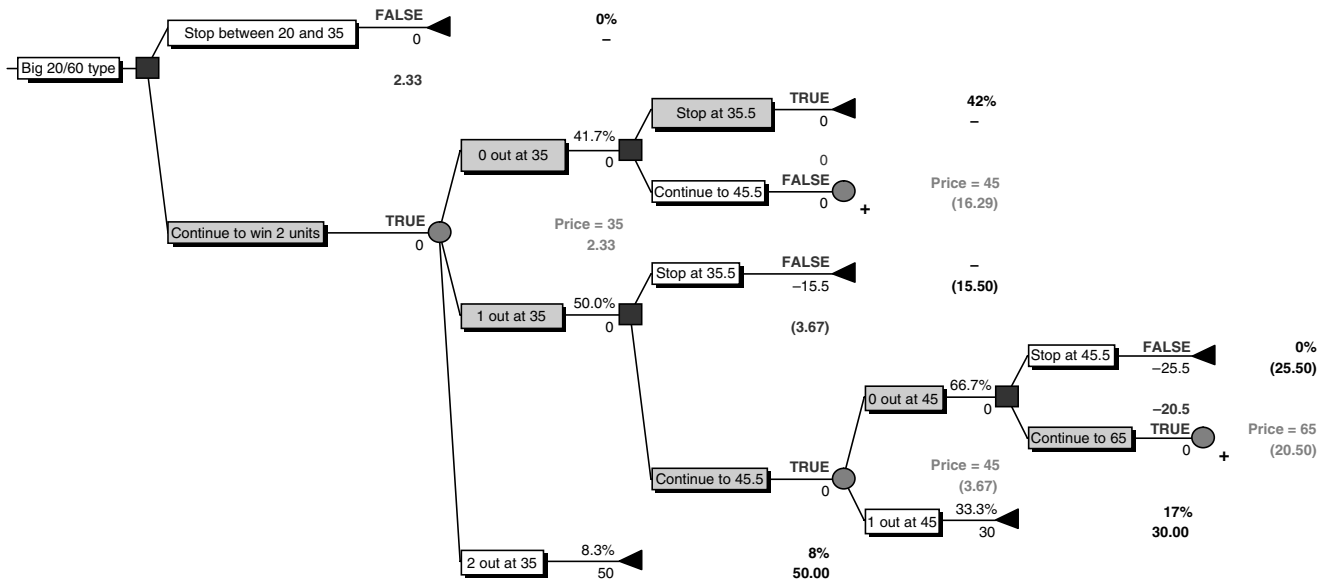
Strategy		Big bidder type																
		20/50				20/60				40/50				40/60				
Small bidders	Probability (%)	Stop between 20 and 35	Winner (units)	Exposure	Efficient?	If zero bidders drop out at 35, stop; otherwise bid to win two units	Winner (units)	Exposure	Efficient?	Bid up to 40 (indifferent between stopping between 35.5 and 40)	Winner (units)	Exposure	Efficient?	If zero bidders drop out at 35, stop between 35 and 40; otherwise bid up to 45.5	Winner (units)	Exposure	Efficient?	Overall
35	8.33	No	Small (2)	No	No	Yes	Big (2)	No	Yes	Big (2)	No	No	Yes	Big (2)	No	No	Yes	
35	8.33	No	Small (2)	No	No	Yes	Big (2)	No	Yes	Big (1) and Small (1)	No	No	No	Big (2)	No	No	Yes	
45	8.33	No	Small (2)	No	No	Yes	Big (2)	No	Yes	Big (1) and Small (1)	No	No	No	Big (2)	No	No	Yes	
45	8.33	No	Small (2)	No	No	No	Small (2)	No	No	Small (2)	No	No	No	Small (2)	No	No	No	
35	16.67	No	Small (2)	No	Yes	Yes	Big (2)	Yes	Yes	Big (1) and Small (1)	No	No	Yes	Big (1) and Small (1)	Yes	Yes	No	
35	16.67	No	Small (2)	No	Yes	Yes	Big (2)	Yes	Yes	Big (1) and Small (1)	No	No	Yes	Big (1) and Small (1)	Yes	Yes	No	
45	16.67	No	Small (2)	No	Yes	No	Small (2)	No	No	Small (2)	No	No	Yes	Small (2)	No	No	No	
45	16.67	No	Small (2)	No	Yes	Yes	Small (2)	No	Yes	Small (2)	No	No	Yes	Small (2)	No	No	Yes	
Probability of this type big bidder (%)		37.50		37.50							12.50	12.50			12.50			
Probability of suffering losses from the exposure problem (%)		0.00		16.67 + 16.67 = 33				0.00			0.00			16.67 + 16.67 = 33				16.50
Probability of an inefficient outcome (%)		8.33 + 8.33 + 8.33 + 8.33 = 33		16.67 + 8.33 = 25				8.33 + 8.33 + 8.33 = 25			8.33 + 8.33 + 8.33 = 25			16.67 + 16.67 + 16.67 + 8.33 = 58.33				32.17

Figure A1 Big Bidder Type with the Value of 20 for One Unit and 50 per Unit for Two Units



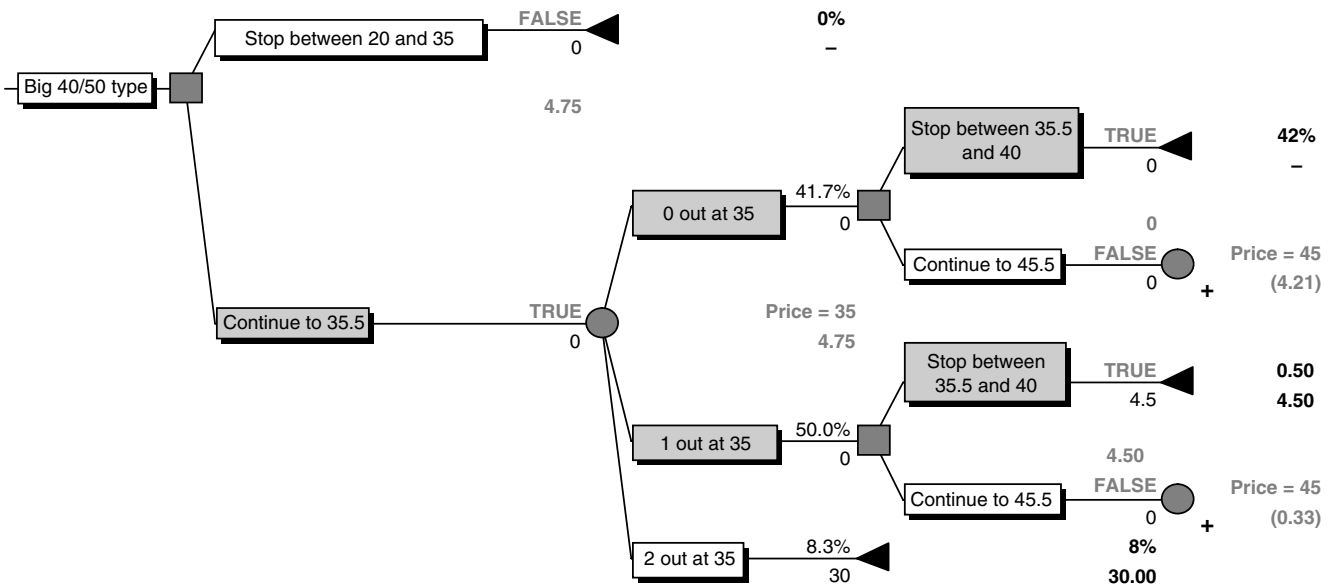
Notes. The profit-maximizing strategy is to bid up to 20. This type of big bidder is indifferent between stopping anywhere between 20 and 35 (or not bidding at all, for that matter) because he never wins in equilibrium, and therefore the best he can do is the expected profit of zero. The best he can do if he continues to bid beyond 35.5 is a loss of 4.33 tokens.

Figure A2 Big Bidder Type with the Value of 20 for One Unit and 60 per Unit for Two Units



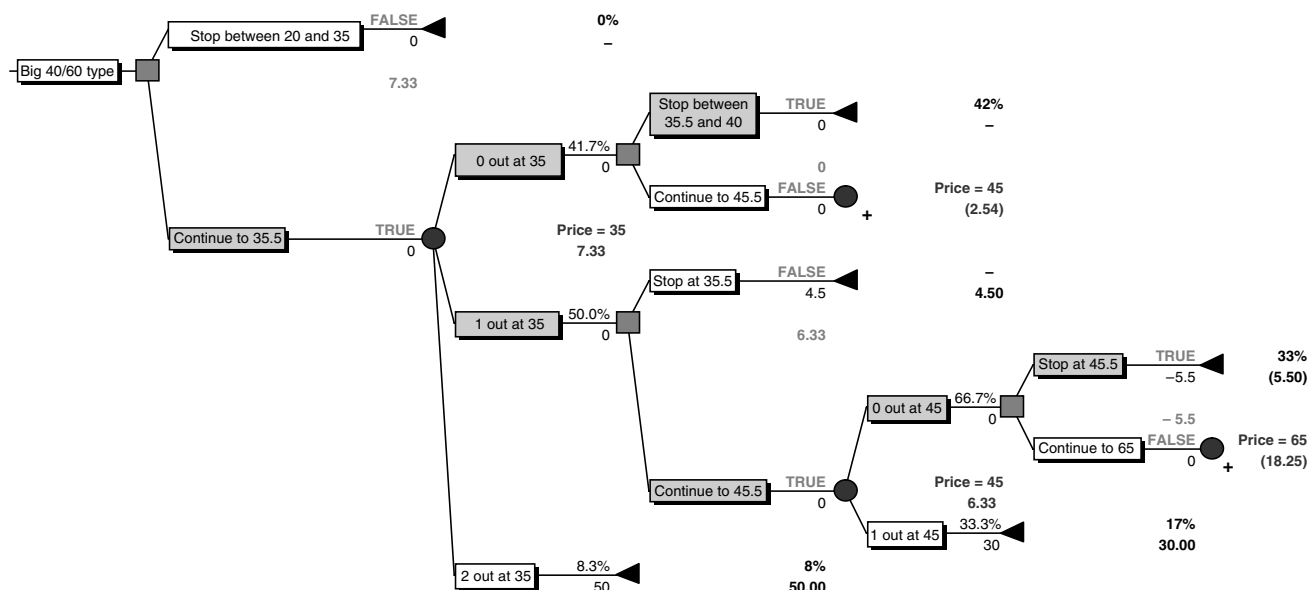
Notes. The profit-maximizing strategy is to drop out at 35.5 if no bidder drops out at 35, and otherwise bid to win two units. This strategy yields the expected profit of 2.33, which is higher than the expected profit from stopping at or below 35 (zero), stopping at 25.5 (loss of 15.5), or stopping at 45.5 (loss of 25.5).

Figure A3 Big Bidder Type with the Value of 40 for One Unit and 50 per Unit for Two Units



Notes. The profit-maximizing strategy for this big bidder type is to essentially forego the opportunity of winning two units, and simply bid to win one unit by bidding up to 40 (this bidder is actually indifferent between stopping anywhere between 35.5 and 40). The expected profit for this strategy is 4.75, which is higher than the expected loss of 4.21 from bidding beyond 35.5 if no bidders drop out, or a loss of 0.33 when one small bidder drops out at 35.

Figure A4 Big Bidder Type with the Value of 40 for One Unit and 60 per Unit for Two Units



Notes. The profit-maximizing strategy for this big bidder is to bid up to 35, if no small bidder exits then to stop bidding between 35.5 and 40 (his value for one unit), and if one small bidder exits to continue to bid up to 45.5. This strategy yields expected profit of 7.33, which is higher than the expected loss of 2.54 from continuing to bid above 35.5 when no small bidders exit at 35, or the expected loss of 18.25 bidding beyond 45.5.

potential for suffering losses from the exposure problem, and whether or not the equilibrium allocation is efficient. Overall, about 16.5% of the time big bidders suffer losses from the exposure problem in equilibrium, and about 32% of auctions end in inefficient allocation.

In Figures A1–A4 we present partial decision diagrams for the four big bidder types that illustrate how the optimal strategies are determined. We did the analysis and generated the diagrams using PrecisionTree software by Palisade Decision Tools (Palisade Corporation 2003).

References

- Ashenfelter, Orley. 1989. How auctions work for wine and art. *J. Econom. Perspectives* 3(3) 23–36.
- Ausubel, Lawrence M. 2002. An efficient ascending-bid auction for multiple objects. Working paper, University of Maryland, College Park, MD.
- Ausubel, Lawrence M., Peter Cramton. 2002. Demand reduction and inefficiency in multi-unit auctions. Working paper no. 96-07, University of Maryland, College Park, MD.
- Best, Ronald W., Roger J. Best, Charles W. Hodges. 1998. The effect of self-tender offers on earnings expectations. *J. Financial Res.* 21(2) 123–138.
- Cantillon, Estelle, Martin Pesendorfer. 2005. Auctioning bus routes. Peter Cramton, Yoav Shoham, Richard Steinberg, eds. *Combinatorial Auctions*. MIT Press, Cambridge, MA.
- Carare, Octavian, Michael Rothkopf. 2004. Slow Dutch auctions. *Management Sci.* Forthcoming.
- Cox, James, Bruce Roberson, Vernon L. Smith. 1982. Theory and behavior of single object auctions. *Res. Experiment. Econom.* 1 61–99.
- Engelbrecht-Wiggans, Richard, Charles M. Kahn. 1998a. Multi-unit pay-your-bid auctions with variable awards. *Games Econom. Behavior* 23 25–42.
- Engelbrecht-Wiggans, Richard, Charles M. Kahn. 1998b. Multi-unit auctions with uniform prices. *Econom. Theory* 12(2) 227–258.
- Engelbrecht-Wiggans, Richard, Robert J. Weber. 1979. An example of a multi-object auction game. *Management Sci.* 25(12) 1272–1276.
- Fevrier, Philippe, Laurent Linnemer, Michael Visser. 2004. Buy or wait: That is the question. Working paper.
- Fischbacher, U. 1999. z-Tree—Zurich toolbox for readymade economic experiments. Experimenter's manual. Working paper #21, Institute for Empirical Research in Economics, University of Zurich, Zurich, Switzerland.
- Hausch, Donald B. 1986. Multi-object auctions: Sequential vs. simultaneous sales. *Management Sci.* 32(12) 1599–1610.
- Kagel, John H., Dan Levin. 2001. Behavior in multi-unit demand auctions: Experiments with uniform price and dynamic Vickrey auctions. *Econometrica* 69(2) 413–454.
- Kagel, John H., Dan Levin. 2004. Multi-unit demand auctions with synergies: Behavior in sealed-bid versus ascending-bid uniform price auctions. *Games Econom. Behavior* Forthcoming.
- Katok, Elena, Anthony Kwasnica. 2004. Time is money: The impact of clock speed on seller's revenue in Dutch auctions. Working paper, Pennsylvania State University, University Park, PA.
- Keskinocak, Pinar, Sridhar Tayur. 2001. Quantitative analysis for Internet-enabled supply chains. *Interfaces* 31(2) 70–89.
- Kraus, Bob, Pierre Mitchell. 2002. The procurement and sourcing applications report, 2001–2006. *AMR Res.* (July 31).
- Krishna, Vijay, Robert W. Rosenthal. 1996. Simultaneous auctions with synergies. *Games Econom. Behavior* 17 1–31.
- Lazear, Edward P. 1986. Retail pricing and clearance sales. *Amer. Econom. Rev.* 76(1) 14–32.
- List, John, David H. Lucking-Reiley. 2000. Demand reduction in multi-unit auctions: Evidence from a sports card field experiment. *Amer. Econom. Rev.* 90(4) 961–972.
- Ledyard, John O., Mark Olson, David Porter, Joseph A. Swanson, David P. Torma. 2002. The first use of a combined-value auction for transportation services. *Interfaces* 32(5) 4–12.
- Lucking-Reiley, David H. 1999. Using field experiments to test equivalence between auction formats: Magic on the Internet. *Amer. Econom. Rev.* 89(5) 1063–1080.

- Lucking-Reiley, David H. 2000. Auctions on the Internet: What's being auctioned and how? *J. Indust. Econom.* **48**(3) 227–252.
- McAfee, R. Preston, Daniel Vincent. 1993. The declining price anomaly. *J. Econom. Theory* **60**(1) 191–212.
- McCabe, Kevin A., Stephen J. Rassenti, Vernon L. Smith. 1990. Auction institutional design: Theory and behavior of simultaneous multiple-unit generalizations of the Dutch and English auctions. *Amer. Econom. Rev.* **80**(5) 1276–1283.
- Milgrom, Paul. 2000. Putting auction theory to work: The simultaneous ascending auction. *J. Political Econom.* **108**(2) 245–272.
- Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Churchill Lectures in Economics, Cambridge University Press, Cambridge, U.K.
- Ockenfels, Axel, Alvin E. Roth. 2002. The timing of bids in Internet auctions: Market design, bidder behavior, and artificial agents. *AI Magazine* (Fall) 79–88.
- Ockenfels, Axel, Alvin E. Roth. 2003. Late and multiple bidding in second-price Internet auctions: Theory and evidence. *Games Econom. Behavior* Forthcoming.
- Palfrey, Thomas R. 1980. Multiple-object, discriminatory auctions with bidding constraints: A game-theoretic analysis. *Management Sci.* **26**(9) 935–946.
- Palisade Corporation. 2003. PrecisionTree: Decision analysis in Microsoft Excel. www.palisade.com/.
- Roth, Alvin E., Axel 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, Michael H. 1977. Bidding in simultaneous auctions with a constraint on exposure. *Oper. Res.* **25** 620–629.
- Rothkopf, Michael H., Alexander Pekeć, Ronald M. Harstad. 1998. Computationally manageable combinatorial auctions. *Management Sci.* **44**(8) 1131–1147.
- Rothkopf, Michael H., Thomas J. Teisberg, Edward P. Kahn. 1990. Why are Vickrey auctions rare? *J. Political Econom.* **98**(1) 94–109.
- van den Berg, Gerard J., Jan C. van Ours, Menno P. Pradhan. 2001. The declining price anomaly in Dutch rose auctions. *Amer. Econom. Rev.* **91**(4) 1055–1062.
- Vickrey, William. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16** 8–37.
- Wilson, Robert. 1979. Auctions of shares. *Quart. J. Econom.* **93**(4) 675–689.