Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Should Sellers Prefer Auctions? A Laboratory Comparison of Auctions and Sequential Mechanisms

Andrew M. Davis, Elena Katok, Anthony M. Kwasnica

MANAGEMENT SCIENCE
Vol. 60, No. 4, April 2014, pp. 990–1008
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.2013.1800

Should Sellers Prefer Auctions? A Laboratory Comparison of Auctions and Sequential Mechanisms

Andrew M. Davis
Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, New York 14853,
adavis@cornell.edu

Elena Katok
Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080,
ekatok@utdallas.edu

Anthony M. Kwasnica
Smeal College of Business, Pennsylvania State University, University Park, Pennsylvania 16802,
kwasnica@psu.edu

When bidders incur a cost to learn their valuations, bidder entry can impact auction performance. Two common selling mechanisms in this environment are an English auction and a sequential bidding process. Theoretically, sellers should prefer the auction, because it generates higher expected revenues, whereas bidders should prefer the sequential mechanism, because it generates higher expected bidder profits. We compare the two mechanisms in a controlled laboratory environment, varying the entry cost, and find that, contrary to the theoretical predictions, average seller revenues tend to be higher under the sequential mechanism, whereas average bidder profits are approximately the same. We identify three systematic behavioral deviations from the theoretical model: (1) in the auction, bidders do not enter 100% of the time; (2) in the sequential mechanism, bidders do not set preemptive bids according to the predicted threshold strategy; and (3) subsequent bidders tend to overenter in response to preemptive bids by first bidders. We develop a model of noisy bidder-entry costs that is consistent with these behaviors, and we show that our model organizes the experimental data well.

Data, as supplemental material, are available at http://dx.doi.org/10.1287/mnsc.2013.1800.

Keywords: behavioral operations management; auctions; experimental economics; behavioral mechanism design
History: Received September 6, 2011; accepted July 26, 2013, by Teck Ho, behavioral economics. Published online in *Articles in Advance* November 27, 2013.

1. Introduction

In this study we analyze a setting in which an asset or a contract is up for bid, and potential bidders must incur a cost prior to bidding to learn their valuations. This setting is used in a number of contexts. For instance, in procurement activities, suppliers must commit significant resources to estimate the value of a contract up for bid. In mergers and acquisitions, one firm must incur the due diligence cost to research the value of the other company. Similarly, in residential home sales, the buyer must visit the house, conduct a home inspection, and research the neighborhood to estimate the value of the property. Two of the more common mechanisms used by the bid taker in these (and many other) scenarios are an auction or a sequential mechanism (Bulow and Klemperer 2009).

Past empirical work suggests that bidders and sellers may differ in terms of how they view auctions. For example, Warren Buffet, when describing the Berkshire Hathaway acquisition criteria in his 2008 annual report writes, "We don't participate in auctions" (Berkshire Hathaway 2009). In a recent poll of private equity firms, 90% said that as bidders, they prefer to avoid auctions, but 80% of those same companies said that as sellers, they prefer auctions (Stephenson et al. 2006). Auctions appear to be one of the primary institutions preferred by sellers, but there is also evidence of sellers utilizing sequential mechanisms. Often referred to as "go-shop" institutions in corporate sales, these sequential mechanisms state that after a seller receives a bid from one firm, they are permitted to shop around for bids from other companies, to which the original firm may respond (Denton 2008). Subramanian (2008) finds that shareholder returns in mergers and acquisitions are 5% higher when companies are acquired through sequential mechanisms and argues that they are preferable for both buyers and sellers.

Auctions and sequential mechanisms have been well studied in the theoretical literature. Bulow and

Klemperer (2009) (see also Fishman 1988) develop a model that predicts that sellers should prefer auctions and bidders should prefer sequential mechanisms, because auctions are revenue maximizing for sellers, and sequential mechanisms result in higher bidder profits. Despite the attention these mechanisms have gained in the theoretical literature and their use in practice, no work has explored them from an experimental standpoint. In this study, we investigate how a standard English auction and sequential mechanism compare to theoretical predictions and each other in a controlled laboratory setting.

We adapt a simple version of the Bulow and Klemperer (2009) model, in which a single item is up for bid, and two potential bidders must incur entry costs prior to learning their valuations. The two mechanisms we compare are an auction, in which the two bidders must make entry decisions simultaneously, and a Bulow and Klemperer (2009) sequential mechanism, in which the bidders make entry decisions sequentially, and the first bidder has an opportunity to signal her valuation by placing a preemptive jump bid. We test the predictions of the model in the controlled laboratory setting under two different cost of entry treatments. In the Lowcost treatment, the cost of entry is low, so the expected differences in seller revenue and bidder profits are more modest than in the Highcost treatment, where the relatively high cost of entry makes it easy for a first bidder to deter entry by a second bidder, thus substantially lowering seller revenue in the sequential mechanism.

The objective of this study is to examine, using the controlled setting of an economics laboratory, how the two mechanisms compare to each other in terms of revenue and profits, and whether bidder behavior is similar to theoretical predictions of the Bulow and Klemperer (2009) model, which predicts that an English auction generates higher seller revenues and the sequential mechanism generates higher bidder profits. There are a number of reasons why this theoretical result may not translate into practice. Past experimental work has shown that when bidders in auctions make bidding decisions without knowing their bidding status, behavior frequently does not conform to standard game-theoretic predictions (for a survey of laboratory auction research, see Kagel 1995). Bidders in sealed-bid first-price auctions tend to bid more aggressively than they should, whereas bidders in English auctions quickly learn to follow the weakly dominant strategy. In our study, both bidders' entry decisions, as well as the first bidder's preemptive bid decision in the sequential mechanism have the "sealed-bid" flavor to them because bidders do not know their winning status, which would result from their decisions.

Similarly, the sequential mechanism model incorporates signaling behavior that assumes bidders are perfect optimizers who can make complex inferences and calculations related to entry decisions and preemptive bidding behavior. Issues of bounded rationality by bidders are likely to affect their behavior and, ultimately, the normative predictions of revenue and profits for the two mechanisms (for summaries of bounded rationality models, see Simon 1984 and Conlisk 1996). By comparing the performance of the auction and the sequential mechanism in a controlled setting of a laboratory, with well-defined rules that match the Bulow and Klemperer (2009) model, we can test whether this model is a good predictor of actual behavior. Previous experimental work, and the relative complexity of the environment we study, suggests that some deviations from theory is to be expected. The objectives of this study are to (1) ascertain whether these deviations are sufficient to reverse the normative predictions of the theory that a seller should prefer an auction whereas a bidder should prefer a sequential mechanism, and (2) develop a new model to explain the observed behavior.

Our main finding is that the preference of the two mechanisms for bidders and sellers is different from the theoretical prediction. In both cost treatments, we find that the sequential mechanism actually results in slightly higher seller revenues than does the auction, whereas average bidder profits are similar. Therefore, our laboratory results indicate that it may well be that sellers should prefer sequential mechanisms over auctions, whereas bidders should be indifferent between the two. We find that the differences between our data and theoretical predictions result primarily from three behavioral phenomena, which we incorporate into a new behavioral model. First, in the auction, bidders do not enter 100% of the time, as the standard theory predicts, thus driving its revenue below that of normative benchmarks. Second, in the sequential mechanism, the first bidders set positive preemptive bids different from the standard theory. Third, the second bidders enter the sequential mechanism auctions more often than they should. This second-bidder overentry, in particular, drives the revenues in the sequential mechanism to be significantly higher than they should be in theory.

Our main contribution is an alternative model of bidder behavior that better organizes our data. We show that if individual bidders derive some (random) benefit or cost from entering auctions,[1] we can

[1] The benefit could come from a variety of sources, such as the joy of competing, or to overestimating the probability of winning the auction, to name a few possible sources. We do not specifically model the source of the cost or benefit. Rather, our intent is to demonstrate that such factors might have a dramatic impact on both individual behavior and aggregate performance of the mechanism.

generate predictions that are largely consistent with what we observe in the laboratory. We use structural modeling to estimate this model and illustrate that it predicts behavior better than the standard theory.

Besides the models of Bulow and Klemperer (2009) and Fishman (1988), which our experimental environment is specifically designed to match, the work of Roberts and Sweeting (2013) is most closely related to ours. They develop a model in which bidders may have noisy estimates of their valuations prior to entering the auction. They show that this addition may also change the predictions of the standard theory, and then estimate the model from U.S. Forest Service timber auctions data. In many ways, our concurrently developed models and empirical studies are complementary in that they evaluate the performance of auctions and sequential mechanisms directly, and also demonstrate the fragility of the Bulow and Klemperer (2009) results to many small but realistic changes to the model.

Other related studies include Bernhardt and Scoones (1993), who present a specific application of a sequential mechanism to wage offers. Arnold and Lippman (1995) compare an auction to a sequential process with information asymmetries, discounting, and costly search by sellers. Hirshleifer and Png (1989) also theoretically study a sequential bargaining process with two bidders; however, they assume that bidding itself is costly. In their setting, the sequential mechanism can generate higher revenue compared to an auction. A concept related to the preemptive bidding in the sequential mechanism is the notion of jump bidding in an English auction (a bidder places a bid greater than the minimum required increment). Avery (1998) models how bidders can use jump bidding to signal in ascending auctions, with the goal of keeping other bidders out of the auction, and as a result earn higher profits, thus differentiating the revenue results of an English auction from a sealed-bid auction.

In addition to these theoretical papers on jump bidding, there is an empirical literature on jump bidding in ascending auctions documenting that jump bidding is commonly observed in practice. Isaac et al. (2007) examine 41 spectrum auctions conducted by the Federal Communications Commission and find that sometimes as many as 40% of the bids are jump bids. Easley and Tenorio (2004) use data from 236 Internet auctions and find that jump bidding is observed in over a third of their sample. Kwasnica and Katok (2007) observe that jump bidding in ascending auctions emerges as a way to decrease the auction duration in a treatment in which bidders have incentives to complete more auctions. However, Kwasnica and Katok (2007) do not find evidence that jump bidding is used for signaling.

In the next section, we describe our experimental design along with standard theoretical predictions for both the auction and the sequential mechanism. In §3, we present the results of all the treatments in our experiment. In §4, we present an alternative model that builds on the standard theory and show that it better describes our data using structural modeling techniques to estimate parameters. In §5, we conclude our investigation with a summary and comment on future research.

2. Experimental Design

In all treatments, two bidders compete to purchase a single indivisible object. We used two bidders to create the simplest possible environment in which the theory applies, thus giving the theory the best chance to be correct. Each subject in every treatment was randomly assigned the role of either a first or second bidder. In each round, a first and second bidder were randomly matched together.[2] In the auction treatments, each round began with both bidders making their entry decisions privately and simultaneously. If both bidders entered the auction, they were then shown their own private values, and proceeded to compete for the item via an ascending clock auction in which the initial price was 0. The bidder who dropped out of the auction first lost the auction, and this drop-out price established the winning bid for the other bidder.

In the sequential mechanism (Seqmech) treatments, each round began with the first bidder of each pair deciding whether or not to enter, and, if she chose to enter, setting an initial preemptive bid for the auction. After the first bidder made these decisions, the second bidder then made her entry decision after observing the first bidder's preemptive bid. If both bidders entered the auction, then they competed for the item in an ascending clock auction in which the initial price corresponded to the first bidder's preemptive bid. (Sample instructions are available upon request.)

The private values for all bidders were integer values, uniformly distributed from 1 to 100, independent and identically distributed, in each round of all treatments. Each subject participated in a single treatment only, and each treatment included 30 rounds. To eliminate the possibility of losses, we provided each subject with an initial endowment of 20 laboratory dollars per round in all four treatments in our study.

In both the auction and sequential mechanism, we ran one set of treatments with an entry cost of 3 ($c = 3$), which we refer to as Lowcost. In a second set of treatments, we set the entry cost to 10 ($c = 10$),

[2] We called the two bidders "bidder A" and "bidder B" in the experiment to avoid any framing effects.

Table 1 Experimental Design and Number of Participating Cohorts

Treatment	Auction	Seqmech	Total
Lowcost	10	10	20
Highcost	10	10	20
Total	20	20	40

which we refer to as Highcost. We varied the entry costs between treatments to help determine whether any potential results were influenced by entry costs rather than the selling mechanism. Table 1 summarizes our design of the experiment and sample sizes. Each treatment included 10 independent cohorts;[3] we use the cohort as the main unit of statistical analysis in our results section (§3).

After the completion of each round of each treatment, we provided the following information to the bidders: who entered the auction, the outcome of any auction (the winning bid was 0 if a single bidder entered), who won the object, and the resulting profits.

In roughly half of the sessions, we also administered a separate and independent second stage of the experiment. This second stage was comprised of the Holt and Laury (2002) risk-aversion elicitation exercise (see §3.1 for details). We performed this separate stage to determine whether any of our results could be attributed to risk aversion. In terms of our experimental procedure at the start of these sessions, subjects were informed that after they finished the 30 rounds of the first stage, there would be a second additional exercise. At this time, no details were provided for the second stage. Then, after all subjects completed all 30 rounds of the first stage, we distributed the instructions for the second stage, read them out loud, answered questions, and administered the exercise.[4]

We conducted all sessions at the Laboratory for Economics Management and Auctions at the Pennsylvania State University, Smeal College of Business, in the spring of 2010 and 2012. Subjects in all four treatments were students, mostly undergraduates, from a variety of majors. Before each session, subjects were allowed a few minutes to read the instructions themselves. Following this, we read the instructions aloud and answered any questions. We recruited participants through an online recruitment system where cash was the only incentive offered. Subjects were paid a $5 show-up fee plus an additional amount that was based on their personal performance for all 30 rounds. Average compensation for the participants, including the show-up fee, was $22. Each session lasted approximately 60 minutes, and we programmed the software using the z-Tree system (Fischbacher 2007).

2.1. Predictions

Given our experimental parameters we begin by expressing bidder behavior, seller revenue, and bidder profits for both the auction and sequential mechanism as predicted by the unique sequential perfect equilibrium identified in the more general model of Bulow and Klemperer (2009).[5]

Under both mechanisms, there are two potential bidders who must decide whether or not to pay a common cost c to learn their private valuations. Values are drawn independently from the continuous uniform distribution on 0 to 1.[6]

The timing of decisions under the two selling mechanisms are different. Under the sequential mechanism, the first bidder "arrives" first and has the opportunity to pay the cost c to learn her value (v_1) and then enter the auction. We denote the possibly mixed strategy between entry or not by the first bidder with the probability of entry of β_1 and not entry $1-\beta_1$. Contingent upon entry, the first bidder learns her valuation and has the opportunity to place a preemptive bid that might depend on her valuation and is denoted by $p(v_1)$. The second bidder arrives next and observes whether or not the first bidder entered, and the preemptive bid. She then decides whether or not to enter and learn her valuation (v_2). We denote the possibly mixed-entry strategy of the second bidder by $\beta_2(p)$ (enter) and $1-\beta_2(p)$ (not enter).[7] After the entry decision of the second bidder, the item is sold in an English auction with the starting price of either 0 (if the first bidder did not enter) or p if the first bidder entered. Assuming that both bidders play the weakly dominant strategy of bidding up to their value in the auction,

[3] In each of the two Highcost treatments, we had one cohort of 8 and one cohort of 10, but all other cohorts consisted of 6 subjects. In total, 252 participants were included in our study.

[4] We found no differences in the 30 rounds of data between these sessions and those in which the second stage was omitted.

[5] The interested reader is referred to Bulow and Klemperer (2009) for a more detailed development of the theoretical predictions in addition to the common equilibrium refinement (on out-of-equilibrium beliefs) that Bulow and Klemperer (2009) utilize to generate uniqueness of the equilibrium.

[6] In the actual experiment, valuations were drawn uniformly on the integer valuations from 1 to 100. When comparing theoretical results with our experimental predictions, we simply multiply the theoretical results by 100. As is standard in experimental auction studies, we assume that the application of the continuous theory to a discrete implementation is sufficiently precise.

[7] The entry strategy of the second bidder can depend upon the preemptive bid strategy and the entry decision of the first bidder in different equilibria (e.g., $p(v)$ and β_1). For notational clarity, we do not include the observed and equilibrium entry decisions of the first bidder. It is obvious that, contingent upon nonentry by the first bidder, the second bidder will have a dominant strategy to always enter for the parameter values of c in our experiment.

Figure 1 **Payoff Table from Auction Entry Stage**

		Bidder 2: Enter	Bidder 2: Not enter
Bidder 1	Enter	$\frac{1}{6}-c,\ \frac{1}{6}-c$	$\frac{1}{2}-c,\ 0$
	Not enter	$0,\ \frac{1}{2}-c$	$0,\ 0$

any auction with only one bidder will end at either 0 (in the event the first bidder did not enter but the second bidder did) or p (the first bidder enters but the second bidder does not). An auction with both bidders will proceed to the maximum of the second-highest valuation of the two bidders ($\min\{v_1, v_2\}$) and the preemptive bid p.

The auction mechanism is similar except that the preemptive bid opportunity is not available to the first bidder. Therefore, in effect, both bidders simultaneously decide whether or not to enter and, after learning their valuations, compete in an English auction. As before, the English auction will progress to a price of 0 (only one bidder entered) or to the second-highest valuation of the two entering bidders.

We first examine the equilibrium in the auction mechanism since it is easily derived from well-known auction results. The payoff table in Figure 1 depicts each player's (ex ante) expected profits from entry in the auction. Clearly, as long as $c < \frac{1}{6}$, it is a dominant strategy for both bidders to enter, resulting in expected bidder profits of $\frac{1}{6} - c$ and a seller expected revenue of $\frac{1}{3}$.[8]

Now consider the equilibrium in the sequential mechanism. The preemptive bidding strategy of the first bidder is the crucial element of the sequential mechanism since it allows for the first bidder to transmit information about her valuation to the second bidder, which might induce the second bidder to not enter. Note that the auction mechanism outcome can always be replicated by the first bidder entering and following a "pooling" preemptive bidding strategy of always bidding 0 (e.g., $p(v_1) = 0$ for all $v_1 \in [0, 1]$). On the other hand, a completely revealing preemptive bid strategy (e.g., $p(v_1)$ is an increasing continuous function of v_1) is not tenable since low-valuing first bidders would want to mimic high-valuing bidders who can discourage competition from the second bidder; the second bidder would *never* enter if she knew the first bidder's value was greater than $1 - \sqrt{2c}$. Therefore, the equilibrium preemptive bidding strategy is of a "partially pooling" nature where low-valuing bidders bid 0 and all others bid a common preemptive bid. Bulow and Klemperer (2009) show that in the unique perfect sequential equilibrium (under a standard refinement on out-of-equilibrium beliefs), the first bidder selects a preemptive bid of 0 if her value is below the cutoff value v_s (called the *deterring value*), and p^* otherwise. In equilibrium, the preemptive bid p^* is chosen in a way that makes the second bidder indifferent between not entering and paying c to compete against a bidder whose value is above v_s. At the same time, p^* is selected such that a first bidder with a value of v_s is indifferent between competing in the auction against the second bidder whose value if uniformly distributed on $[0, 1]$, or winning the auction outright with the bid of p^*.

Formally, the equilibrium preemptive bid has the following form:

$$p(v) = \begin{cases} 0 & v < v_s, \\ p^* & v \geq v_s, \end{cases} \tag{1}$$

where $p^* \leq v_s$ ensures individual rationality for the first bidder. Given this preemptive bid strategy, the second bidder can calculate her expected auction profits (denoted π_2^a) for each preemptive bid observed in equilibrium:[9]

$$\pi_2^a(p) = \begin{cases} \dfrac{(v_s)^2}{6} + \dfrac{1 - v_s}{2} & p = 0, \\[2ex] \dfrac{(1 - v_s)^2}{6} & p = p^*. \end{cases} \tag{2}$$

To make preemptive bidding worthwhile, the second bidder must be induced to not enter whenever p^* is observed. Assuming that the bidder will decide not to enter when she is indifferent between entry and not, we must have $\pi_2^a(p^*) - c = 0$, or

$$\frac{(1 - v_s)^2}{6} = c. \tag{3}$$

Note that since $p^* \leq v_s$, the second bidder's expected auction profits only depends on the cutoff value v_s, so solving for v_s yields

$$v_s = 1 - (6c)^{1/2}. \tag{4}$$

When the second bidder enters the auction, the first bidder's (interim) auction profits depends on the chosen level of the preemptive bid and is given by

$$\pi_1^a(p, v_1) = \frac{v_1^2 - p^2}{2}. \tag{5}$$

[8] If $\frac{1}{3} \leq c < \frac{1}{2}$, there are multiple equilibria where only one bidder enters and the other does not, resulting in a revenue of 0 for the seller. In this case, there is also a mixed-strategy equilibrium. Bulow and Klemperer (2009) and we do not consider this case explicitly.

[9] More generally, the second bidder's expected auction profits from competition against a bidder whose values lie uniformly in the subinterval of the original distribution given by $[\underline{v}, \bar{v}]$ is $\pi_2^a = (\bar{v} - \underline{v})^2/6 + ((1 - \underline{v})(1 - \bar{v}))/2$.

Table 2 Experimental Predictions

	Actual draws		Theory	
Treatment	Auction	Seqmech	Auction	Seqmech
Lowcost ($c = 3$)				
Seller revenue	33.05	30.54	33.33	30.79
First-bidder profit	14.23	16.15	13.67	16.21
Second-bidder profit	13.80	13.45	13.67	13.67
Deterring value, v_s	—	57.57	—	57.57
Preemptive bid, p^*	—	41.00	—	41.00
Efficiency	1.000	0.919	1.000	
Highcost ($c = 10$)				
Seller revenue	33.56	17.98	33.33	17.84
First-bidder profit	6.72	22.29	6.67	22.16
Second-bidder profit	7.00	6.80	6.67	6.67
Deterring value, v_s	—	22.54	—	22.54
Preemptive bid, p^*	—	20.00	—	20.00
Efficiency	1.000	0.678	1.000	

Table 3 Summary of the Data

Treatment	Auction	Seqmech
Lowcost ($c = 3$)		
Seller revenue	31.57 (0.85)	34.25 (1.30)
First-bidder profit	14.58 (1.09)	11.56 (1.50)
Second-bidder profit	14.00 (1.27)	12.59 (1.27)
Preemptive bid	—	7.97 (1.37)
First-bidder-entry proportion	0.959 (0.013)	0.983 (0.009)
Second-bidder-entry proportion	0.967 (0.015)	0.937 (0.022)
Efficiency	0.888 (0.018)	0.900 (0.017)
Highcost ($c = 10$)		
Seller revenue	26.61 (0.76)	30.30 (1.39)
First-bidder profit	8.15 (1.01)	7.36 (1.08)
Second-bidder profit	9.33 (1.18)	9.24 (1.46)
Preemptive bid	—	10.24 (1.12)
First-bidder-entry proportion	0.843 (0.032)	0.928 (0.025)
Second-bidder-entry proportion	0.907 (0.020)	0.862 (0.034)
Efficiency	0.823 (0.018)	0.841 (0.022)

Note. Standard errors in parentheses.

The equilibrium is therefore found by selecting a p^* that ensures that low-valuing first bidders (those with values below v_s) prefer a preemptive bid of 0 to p^*. Since Equation (5) is an increasing function of v_1, this is found by finding the p^* such that $\pi_1^a(0, v_s) = v_s - p^*$, where the right-hand side is the certain profit from bidding a preemptive bid of p^* and therefore deterring entry by the second bidder. Given the value of v_s from Equation (4), the preemptive bid p^* is given by

$$p^* = \tfrac{1}{2} - 3c \qquad (6)$$

whenever $c < \frac{1}{6}$. Given the enhanced profitability of this preemptive bidding strategy, the first bidder will always enter ($\beta_1 = 1$).

Expected profits of both the bidders and the sellers can be calculated given the equilibrium and our parameterizations. Table 2 summarizes predicted seller revenue, bidder profits (net endowments), deterring values, preemptive bids, and efficiency for our experimental parameters. We define efficiency as the proportion of time the bidder with the highest valuation wins the item.[10]

Note that the Bulow and Klemperer (2009) theory predicts that the seller revenue is higher in the auction, and bidder expected profit (particularly the first bidder's profit) is higher under the sequential mechanism. Furthermore, the difference in seller revenue between the two mechanisms should be increasing in c; we intentionally selected the cost parameters such that the expected differences between the two mechanisms was quite high in the Highcost treatment, whereas the difference was smaller in the Lowcost treatment. Finally, the predicted deterring value and optimal preemptive bid are both decreasing in c.

3. Results

Table 3 summarizes average seller revenue, bidder profits, preemptive bids, and entry rates in the experiment.[11]

We can see from Table 3 that in both cost conditions, the sequential mechanism generates higher revenue for the seller compared to the auction. Using the cohort average as the main statistical unit of analysis (we follow this approach for all statistical tests in this section), we find that a one-sided t-test comparing the Auction and Seqmech revenue results in $p = 0.0158$ in Highcost and $p = 0.0507$ in Lowcost.[12] This is counter to the predictions of the Bulow and Klemperer (2009) model, which predicts that the auction should generate higher seller revenues in both cost treatments. We also observe that the sequential mechanism has a slightly higher efficiency than the auction, although the differences are not significant for either cost condition.

Comparing Tables 2 and 3, we can also see that the auction, particularly in the Highcost treatment, generates seller revenue that is slightly below theoretical predictions (two-sided t-test Highcost $p < 0.001$ and Lowcost $p = 0.1153$). In our data, we see that bidders in the auction do play the dominant strategy

[10] In the data, we have access to all potential values, regardless of whether or not a bidder enters. Therefore, in a particular auction, if only one bidder enters, but this bidder happens to have the larger of the two potential values, this auction would be marked as efficient.

[11] We provide cohort-level data for revenue in Appendix C.

[12] A more conservative Mann–Whitney test results in $p = 0.0413$ in Highcost and $p = 0.1304$ in Lowcost.

Figure 2 **Proportion of Preemptive Bids Equalling Zero in Data (a) and in Theory (b)**

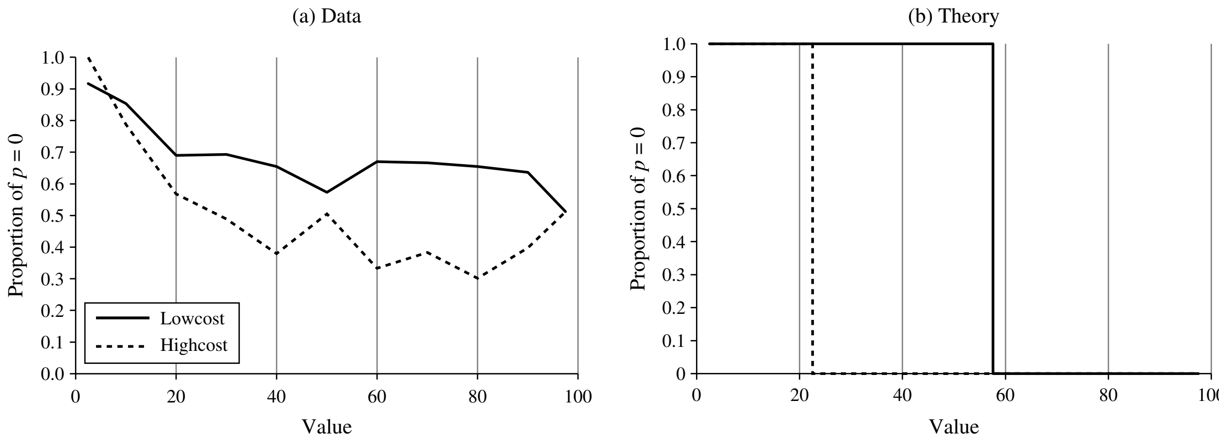

of bidding up to their values, so lower than predicted auction revenues are due to entry behavior.[13] Bidders enter the auction only 96.28% of the time in the Lowcost treatment, and 87.46% of the time in the Highcost treatment. Lower than 100% entry rates account for the auction's revenues being slightly below the predicted values, along with lower efficiencies, and suggest that bidders respond to the magnitude of the entry cost when making their entry decisions. We explore alternative models for this entry behavior in later sections.

Turning to the bidders' profits, the first bidders should fare better in the sequential mechanism than in the auction, because it provides them with the opportunity to set a preemptive jump bid, potentially deterring the second bidders from entering the auction. On the other hand, if the first bidders act optimally, second bidders earn the same average profits under the two mechanisms.

In the auction, both first- and second-bidder profits are largely in line with the predicted values. (There is a slight increase in bidder profits in our data due to the entry decisions mentioned previously; however, this does not cause any of these differences to be statistically significant.) In the sequential mechanism, second-bidder profits are also not statistically different from theoretical predictions. However, first-bidder profits in the sequential mechanism are far below theoretical predictions (two-sided t-test leads to $p < 0.0001$ in Highcost and $p < 0.0136$ in Lowcost). This last finding also results in the first bidders' average profits being roughly the same between the auction and the sequential mechanism.

Thus far, we have shown that the sequential mechanism results in higher seller revenues than an auction, which is contrary to the Bulow and Klemperer (2009) model. This higher revenue for the seller is achieved primarily at the expense of the first bidder. Next, we examine both bidders' decisions in comparison to the equilibrium predictions of Bulow and Klemperer (2009) to better understand the potential causes of these findings.

We begin by examining how first bidders set preemptive bids. In the Bulow and Klemperer (2009) model, first bidders follow a threshold strategy, as shown in Figure 2(b): first bidders should set the preemptive bid equal to 0 if their value is below v_s, and when their value is above v_s, they should set the preemptive bid equal to p^*. Figure 2(a) shows the proportion of preemptive bids set to 0 in our data, as a function of value.[14] It is clear from Figure 2 that bidders do not follow the threshold strategy, but instead, their probability of setting a preemptive bid of 0 decreases in value up to some point, and then levels off, never reaching a probability of 0.

Next we examine the magnitude of preemptive bids. The preemptive bids should follow a threshold strategy, as shown in Figure 3(b), specifically, preemptive bids should be constant when $v \geq v_s$, and positive preemptive bids should be different for the two cost conditions. However, in our data, summarized in Figure 3(a), we see that the magnitude of positive preemptive bids increases in v linearly, and moreover, there is no discernible difference in the two cost conditions. We confirmed this formally with a random effect regression with the preemptive bid as the dependent variable: the coefficient on v is positive and significant, the coefficient on *HIGHCOST* is not

[13] The second-lowest value less the winning bid was, on average, −0.11 across all observations in all four treatments.

[14] For all figures, we removed any observation where bidder 1 entered and set a preemptive bid equal to or above their value; this occurred 19 times out of all 1,890 decisions by first bidders.

Figure 3 **Magnitude of Positive Preemptive Bids in Data (a) and in Theory (b)**

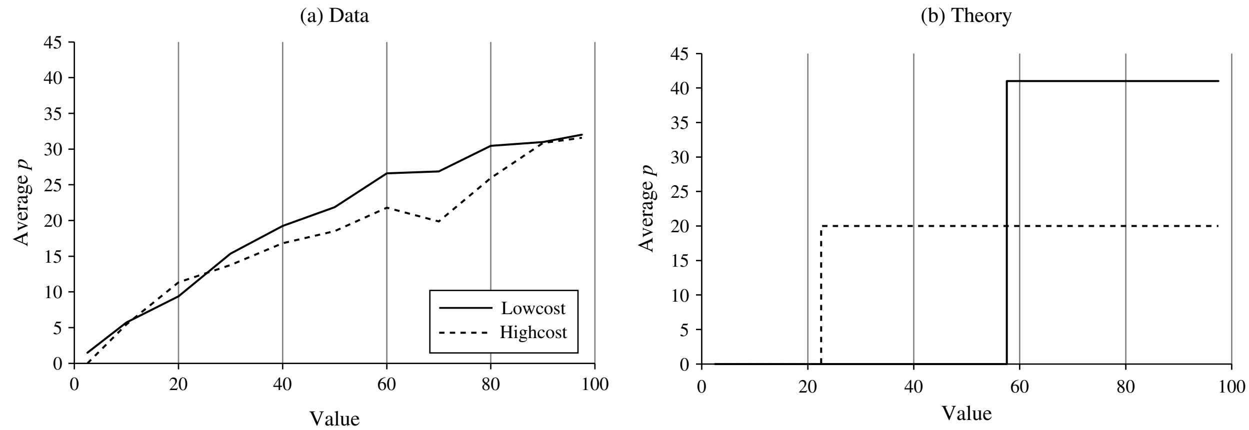

significant, and neither is the coefficient on the interaction variable $v \times HIGHCOST$.

Our conclusions about the behavior of first bidders in the sequential mechanism are summarized as follows: First bidders enter the auctions less than 100% of the time. When they do enter, they do not follow the threshold strategy in regards to either their decision to place a positive preemptive bid, or to the magnitude of the preemptive bid. For low v's, first bidders' probability of entering a positive preemptive bid increases in value, and for high v's, these probabilities reach a constant level, that is, significantly below 100%. Therefore, first bidders are more likely to place preemptive bids when their v's are low, and are not likely enough to place positive preemptive bids when their v's are high. The size of the preemptive bid itself increases in v and does not depend on the entry cost.

Moving on to the second bidder's behavior, as shown in Figure 4(b), according to the theory, second bidders should always enter as long as the preemptive bid is below p^*, and should never enter as long as the preemptive bid is above p^*. The critical difference between how our second bidders enter and how the Bulow and Klemperer (2009) model says they should enter is that they enter too often following high preemptive bids. Specifically, second bidders in the $c = 10$ condition should never enter when preemptive bids are above 20, and second bidders in the $c = 3$ condition should never enter when preemptive bids are above 41. However, as we can see from Figure 4(a), second bidders enter quite frequently when faced with preemptive bids exceeding 20 and 41. The expected profitability of entry depends on the beliefs of the second bidder about the first bidder's value given the observed preemptive bid. Since we know first bidders are not placing preemptive bids in accordance with the theory, it may not be irrational that second bidders are entering. Even under the most optimistic beliefs about the first bidder's value given a preemptive bid that $p = v_1 - c$, in the Highcost condition the second bidder would not want to enter after

Figure 4 **Entry Proportion of Second Bidders in Response to the Preemptive Bid in Data (a) and in Theory (b)**

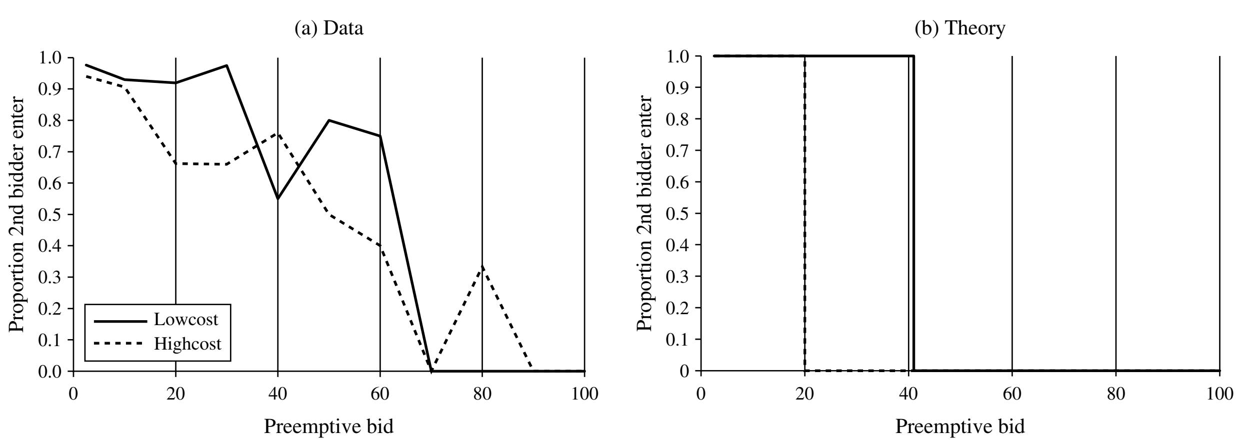

observing a preemptive bid of greater than 45.[15] As is evidenced by Figure 4(a), the entry rate for preemptive bids between 45 and 60 is quite high. We can also demonstrate overentry by second bidders empirically. In the Highcost (Lowcost) treatment, second bidders make losses, on average, whenever they enter following a preemptive bid of 19 (31), yet they continue to enter frequently after observing such preemptive bids (62.87% Highcost, 61.25% Lowcost).[16] Second bidders also sometimes fail to enter for low preemptive bids, but this effect is not very large.[17]

In sum, our data suggest that the sequential mechanism generates the same or higher revenue to sellers when compared to the auction, and roughly the same profits to bidders. These results stem from three systematic behavioral deviations from the theory: (1) in the auction, subjects do not enter quite enough, especially when entry costs are high; (2) in the sequential mechanisms, first bidders do not follow the threshold strategy in setting preemptive bids, and as a result they end up not setting preemptive bids frequently enough, and when they do set them, the size of the preemptive bid is positively correlated with the first bidder's value; and (3) in the sequential mechanism, second bidders enter even when first bidders set high preemptive bids.

3.1. Risk-Aversion Analysis

Risk aversion has often been cited as a potential cause of deviations in observed auction behavior from that of the standard theory where risk neutrality is typically assumed (Cox et al. 1982). In this section, we attempt to determine whether the observed behavior in our experiments is driven by risk aversion. We examine the issue both theoretically and experimentally.

First note that, since it is a weakly dominant strategy for a bidder to bid up to her value in an English auction, behavior in the auction-bidding stage should be unaffected by risk aversion. Now consider the potential impact of risk aversion on preauction behavior in the sequential mechanism. We consider constant relative risk-aversion (CRRA) preferences, which have been used extensively to study risk aversion in both laboratory settings and empirical studies (Binswanger 1980, Chen and Plott 1998, Campo et al. 2011). CRRA utility functions are given by $u(x) = x^{1-\alpha}$, where $0 \leq \alpha < 1$ is the measure of relative risk aversion. Given this specification, the calculation of the equilibrium described by Equations (1)–(6) can be replicated assuming risk aversion.[18] The impact of risk aversion on the equilibrium preemptive bid p^*, deterring value v_s, and expected revenue is depicted in Figure 5. Since $\alpha = 0$ is a risk-neutral decision maker, it is clear to see that both the deterring value and preemptive bid are decreasing as players become more risk averse, driving expected revenue down. For sufficiently high levels of risk aversion, the equilibrium-deterring values fall to 0. At those levels of risk aversion it is sufficient for the first bidder to enter (and place a preemptive bid of 0) to deter entry by the second bidder. For even higher levels of risk aversion, the first bidder may prefer nonentry, but for the level of entry costs considered, a very high level of risk aversion (e.g., $\alpha > 1$) is required. Previous studies have shown estimates for the level of risk aversion (α) to be between 0.45 and 0.67 and to vary depending on the setting (Cox and Oaxaca 1996, Goeree et al. 2002, Goeree and Holt 2004). Note that for those values of risk aversion under the Highcost condition, the equilibrium-deterring value is 0, meaning that simple entry by the first bidder (without a preemptive bid) results in zero revenue for the seller. The observed behavior under the sequential mechanism is obviously substantially different than that predicted by risk aversion.

Turning to the auction, as mentioned earlier, bidders failed to enter 100% of the time under the auction mechanism. Although under risk neutrality it is a dominant strategy for both bidders to enter, it is possible that such behavior may be the result of risk aversion. If players are sufficiently risk averse, there will exist two pure-strategy Nash equilibria where one bidder enters and the other does not (since sufficiently risk-averse bidders may prefer the certain payoff of 0 for nonentry versus the risky but positive expected payoff of participating in the auction versus another bidder). If this is the case, there will also be a mixed-strategy equilibrium where both bidders enter with some probability. Although this may rationalize behavior in the auction, these levels of risk aversion, however, contradict behavior in the sequential mechanism. If bidders are sufficiently risk averse to induce a mixed strategy in the auction, it also means that the

[15] The belief that $p = v_1 - c$ is the most optimistic second-bidder belief, under the assumption that first bidders do not place preemptive bids that guarantee losses. Under the Lowcost condition, the preemptive bid would have to be 72, which is rarely observed.

[16] Even in the second half of the 30 experimental periods, second-bidder entry following these preemptive bids is common at 61.26% (Highcost) and 71.43% (Lowcost).

[17] It is worth noting that the number of observations across values in Figures 2 and 3 is quite constant. However, in Figure 4, the number of observations is right skewed, so that the number of preemptive bids that were above 50, for example, only occurred roughly 1% and 2% of the time in Lowcost and Highcost, respectively.

[18] In Appendix A, we provide a specification of the equivalent equilibrium conditions under risk aversion. Analyses show that under most specifications of risk-averse preferences (such as constant abolute risk-aversion utility), expected revenue is decreasing with increases in risk aversion.

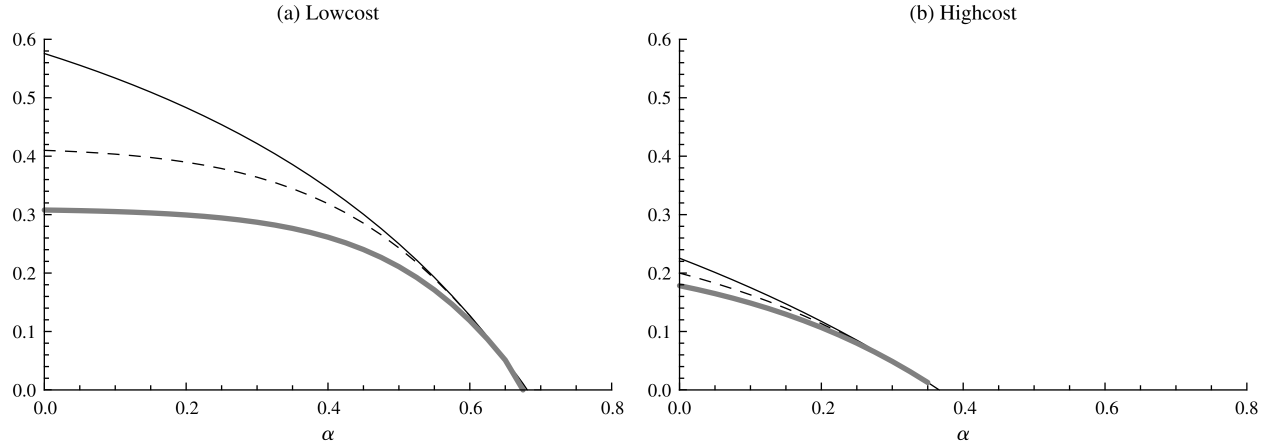

Figure 5 **Equilibrium Preemptive Bid (Dashed Line), Deterring Value (Solid Line), and Expected Revenue (Bold Line) as a Function of α in CRRA Utility**

equilibrium in the sequential mechanism will involve the first bidder entering and the second bidder being deterred simply by knowledge that another bidder entered. (See Appendix A for a formal proof.) In other words, under the assumption of CRRA preferences, the level of risk aversion must fall into the range of risk aversion where the deterring value is 0 (e.g., to the right of the curves in Figure 5). Considering the empirically observed mean frequency of entry of 0.963 for the Lowcost treatment and 0.875 for the Highcost treatment, the value of the risk-aversion parameter for CRRA preferences that would rationalize this level of entry as a mixed strategy would be $\alpha = 0.70$ (Lowcost) and $\alpha = 0.50$ (Highcost). Although these numbers are not too far from those observed in other experiments, it is clear that behavior in the sequential mechanism is inconsistent with that observed in the our experiments.[19]

As mentioned in §2, in some of our treatments, we had subjects complete a second stage of the experiment where we administered the Holt and Laury (2002) risk-aversion elicitation exercise. In this exercise, subjects were required to select their preference between 10 lottery pairs. In each pair, the "safe" option, A, resulted in a payoff of either \$2.00 or \$1.60, and the "risky" option, B, resulted in either \$3.85 or \$0.10. In the first pair listed, the chance of the higher payoff of both options (\$2.00 and \$3.85) was 10%. In the second pair, the chance of the higher payoff was 20%, in the third pair, it was 30%, and so on.

For each subject that completed the risk-aversion elicitation exercise, we calculated the number of times they selected option A. We use this as a proxy for risk aversion, where more selections of option A are linked to higher levels of risk aversion. We report logit regressions (with random effects) for the sequential mechanism with second-bidder entry as the dependent variable in Table 4.[20]

Table 4 **Logit Regressions Examining Whether Risk Aversion Is Related to Entry by Second Bidders in the Sequential Mechanism**

Variable	Description	Lowcost	Highcost
Constant	Intercept	2.169	4.039***
		[1.811]	[1.000]
SumA	Total number of "safe" options selected	0.393	−0.040
		[0.283]	[0.174]
Period	Decision period	−0.006	0.017
		[0.017]	[0.015]
Jump	Preemptive bid	−0.088***	−0.089***
		[0.012]	[0.010]

***$p < 0.01$.

In Table 4 we observe that the coefficient *SumA* is insignificant in both cost conditions.[21] However, note that the coefficient on *Jump* is negative and significant in both regressions. Combining this with our previous entry observations, it appears that subjects were somewhat deterred by higher preemptive bids, but not enough to coincide with the standard theoretical predictions. Therefore, considering that risk aversion is not a key driver in explaining second-bidder-entry decisions in the sequential mechanism, we now turn to a more formal model that may explain this behavior.

[19] It is possible that a model of heterogeneity of risk aversion may generate results that are qualitatively similar to the model that we develop in §4, but the parameter estimation suggests that at least some bidders would have to be assumed to be risk loving.

[20] For these regressions, we excluded those observations where the first bidder failed to enter.

[21] We ran a variety of different regressions on second-bidder entry. In no regression was the coefficient on *SumA* significant.

4. Modeling Bidding Behavior

The objective of this section is to develop a parsimonious and plausible model of bidder behavior that matches, at least qualitatively, the features of bidder behavior identified in §3. Primarily, we are asking the following question: Is there a model that deviates from the standard theory of Bulow and Klemperer (2009) in a realistic and minimal way that better organizes the experimental data?

As is typical with such an exercise, we could have varied the model in a number of (potentially complementary) ways. We considered three possible types of changes to the theory of Bulow and Klemperer (2009). First, Bulow and Klemperer (2009) assume that bidders play a particular perfect sequential signaling equilibrium that is uniquely identified via a standard equilibrium refinement.[22] Without this refinement, there are a continuum of potential perfect sequential equilibria. One possibility in our data might be that player behavior is more closely approximated by some other equilibrium. There is a substantial literature examining whether signaling equilibria develop experimentally and the efficacy of various equilibrium refinements. For example, in the context of limit pricing, Cooper et al. (1997b) find that signaling equilibrium behavior will often develop. On the other hand, in the same context, equilibrium selections predicted by seemingly plausible refinements do not always present themselves in the data (Cooper et al. 1997a). In these experiments, the signaling environment is typically more simple than the one studied here due to finiteness of both the type space and strategy spaces of the players. The complexity of the type and strategy spaces in the sequential mechanism and the ensuing continuum of potential other equilibria makes our experiment not amenable to a rigorous examination of whether other equilibria are chosen.

A second approach might be to abandon the perfect sequential equilibrium approach altogether in favor of another equilibrium concept that allows for, potentially, more realistic behavior. Some possibilities might include an adaptive-learning-type model proposed by Cooper et al. (1997b) or the increasingly popular quantal-response-equilibrium model of McKelvey and Palfrey (1995) and extended to extensive form games with the agent quantal-response-equilibrium (AQRE) model in McKelvey and Palfrey (1998). Although these models have proven remarkably successful in explaining experimental data and, as we discuss below, there is a similarity between our proposed model and a simplified AQRE model, a full model of either adaptive learning or quantal response across all stages of the game has not proven to be readily tractable. In addition, an equilibrium concept that allows for noisy behavior in all stages of the game would, therefore, predict noisy behavior in the auction-bidding phase, whereas our data indicate that bidding behavior in the auction stage is remarkably consistent with standard theory.

The third approach, which we ultimately selected, is to propose changes to the underlying payoffs or structure of the game and to retain the perfect sequential equilibrium concept (with a similar refinement). This is a common approach taken by many behavioral models that seek to explain experimental data. For example, models of equity and reciprocity (Bolton and Ockenfels 2000) have proven successful at explaining behavior in ultimatum, public good, and dictator games, among others. Models that allow for regret (Engelbrecht-Wiggans and Katok 2008) are also consistent with experimentally observed bidding behavior in first-price auctions. In this context, Roberts and Sweeting (2013) demonstrate that a model that retains the same equilibrium concepts but assumes that players get a noisy signal of their valuation prior to entry is sufficient to substantially change predicted bidder behavior away from the partially pooling equilibrium identified by Bulow and Klemperer (2009) in favor an equilibrium where the preemptive bid function reveals the first bidder's valuation. The theory model we develop here shares a number of similarities with the approach concurrently developed by Roberts and Sweeting (2013).

The fact that we have chosen this third approach is not meant to exclude the other two approaches as potential explanations of our data. Indeed, as is shown by Goeree et al. (2002), it is often the case that many different modeling approaches can arrive at similar conclusions. It is also possible that a hybrid model that includes features such as learning would better organize our experimental data. However, the exercise here is not to identify the exact model of behavior but to look for a simple and tractable model that generates the observed behavior. The fact that so many models might arrive at similar conclusions further highlights the fragility of the Bulow and Klemperer (2009) normative result.

4.1. The Model

We develop a model of noisy bidder-entry decisions. In particular, we assume that in addition to paying a cost c to enter and learn their values, each bidder (i) perceives an additional benefit/cost of ϵ_i for entry into the mechanism where ϵ_i is privately known by the bidder at the time of entry. We assume that ϵ_i is drawn independently from the normal distribution $N(\mu, \sigma)$. As is the case of the entry cost c, we assume the additional cost factor ϵ_i to be sunk at the time of entry decision so that it does not directly impact

[22] See footnote 11 in Bulow and Klemperer (2009) for a description of the equilibrium refinement utilized.

future decisions such as auction-bidding strategies (for both bidders) or preemptive bidding strategies (for bidder 1).[23]

There are a number of justifications for the inclusion of such a term. Formally, these errors might be generated as some type of idiosyncratic cost or benefit element. In practice, we feel it is reasonable that in high-value auctions, of the type where this model is probably most appropriate, such as mergers and acquisitions and procurement settings, that, in addition to the commonly known cost element, there might be idiosyncratic cost or benefit elements that are not known by the other participants. For example, a firm considering bidding on a procurement contract might decide not to spend the considerable effort required to put together a cost estimate (e.g., decide not to enter) because of internal issues within the firm. In the laboratory setting, this idiosyncratic cost element might come from more *psychic* benefits or costs perceived by the experimental subjects. For example, a subject might prefer to avoid the cognitively difficult task of determining a proper bid and therefore decide not to enter the auction. On the other hand, a subject might perceive some benefit from "getting in the game" and decide to enter despite potentially negative monetary rewards. Previous work by Katok and Kwasnica (2008) and Kwasnica and Katok (2007) has shown that bidders in auctions will often respond to other costs/benefits not directly induced via monetary incentives. Modeling noise as a random cost or benefit element may serve as a useful approximation to capture other behavioral issues, such as regret, social preferences, or errors in calculations of expected profitability. While these models might invite somewhat different function formulations, exploratory attempts to formally model these behavior yielded largely consistent results. The advantage of the approach we chose here is the fact that the additional cost term ϵ_i enters into the bidders' calculations in an additively separable manner that makes the theory substantially more tractable. The fact that the term is sunk at the time of bidding meant that auction stage behavior would conform to theory as it does in the experiment.

We continue to assume that values are distributed uniformly on 0 to 1. Since the payoff from nonentry is 0, a bidder will decide to enter only if

$$\pi_i^a - c + \epsilon_i \geq 0,$$

where π_i^a is bidder i's expected profits from the auction. This, of course, means that a bidder will only enter if this extra term is sufficiently large where $\epsilon_i^* = c - \pi_i^a$ represents the cutoff between entry and not. The (ex ante) entry probability for a bidder is then given by

$$\beta_i = \Pr(\epsilon_i \geq \epsilon_i^*) = 1 - \Pr(\epsilon_i \leq \epsilon_i^*) = 1 - \Phi\left(\frac{\epsilon_i^* - \mu}{\sigma}\right), \quad (7)$$

where $\Phi(\cdot)$ is the cumulative distribution function (cdf) of the standard normal distribution.

Let us first consider the impact of a noisy cost of entry on the equilibrium entry decisions of both players in the auction. Because both players are now entering with less than probability 1, each bidder must consider the fact that they may be the sole entrant into the auction and, therefore, obtain a greater profit. Given the (ex ante) entry probability of the other bidder β_j, bidder i's expected payoff from the auction is

$$\pi_i^a = \beta_j \tfrac{1}{6} + (1-\beta_j)\tfrac{1}{2} \quad (8)$$

$$= \tfrac{1}{2} - \beta_j \tfrac{1}{3}, \quad (9)$$

where the first term in Equation (8) is the expected profits to a bidder in a two-person auction, and the second term is the expected profits in the event of nonentry by the other bidder so that the auction price is 0. Since the payoff from nonentry is 0, bidder i will enter only if

$$\tfrac{1}{2} - \beta_j \tfrac{1}{3} - c + \epsilon_i \geq 0$$

or

$$\epsilon_i \geq \beta_j \tfrac{1}{3} + c - \tfrac{1}{2},$$

which results in the following entry probability:

$$\beta_i = 1 - \Phi\left(\frac{\frac{1}{3}\beta_j + c - \frac{1}{2} - \mu}{\sigma}\right). \quad (10)$$

In a symmetric equilibrium in the auction, each bidder will enter if $\epsilon_i \geq \frac{1}{3}\beta^* + c - \frac{1}{2}$, resulting in expected entry probability $\beta^* = \beta_1 = \beta_2$ that is the solution to the previous equation for both bidders.

Next, consider the sequential mechanism. The key strategic variable is now the preemptive bid. We proceed by characterizing the necessary conditions for a revealing equilibrium with noisy entry decisions. In Appendix B, we show, following a similar approach to Roberts and Sweeting (2013), that the equilibrium identified is indeed the unique perfect sequential equilibrium under the D1 refinement (Banks and Sobel 1987, Cho and Kreps 1987), which is a common restriction placed on out-of-equilibrium beliefs in signaling games. Let us suppose there exists a revealing preemptive bid function $p(v_1)$ with $p(v_1) \leq v_1$ for all v_1. Suppose that the bid function is differentiable and increasing everywhere so that

[23] Our model can also be considered to be a restricted AQRE model of McKelvey and Palfrey (1998) where the noisy behavior is restricted to only occur in the entry decisions and not in the other stages of the game, and the random unobserved error term is normally distributed. Although most implementations of QRE models assume a logistic distribution of the error term, the general theory allows for many error distributions, including the normal distribution.

$p'(v_1) > 0$. The boundary condition is that $p(0) = 0$. Let $v^{-1}(p)$ be the inverse preemptive bid function. Then, the second bidder's expected profit from the auction having observed a preemptive bid p is given by

$$\pi_2^a(p) = \frac{(1 - v^{-1}(p))^2}{2}.$$

The cutoff value for entry is therefore given by $\epsilon_2^*(p) = c - \pi_2^a(p)$, and the ex ante entry probability of the second bidder, denoted now by $\beta_2(p)$, is given by Equation (7), with this expected term substituted into the equation. The first bidder's expected profit from the auction contingent upon entry by the second bidder is still given by Equation (5). The first bidder's expected profits from a particular preemptive bid level is therefore given by[24]

$$\begin{aligned} \pi_1(p, v_1) &= \beta_2(p)\pi_1^a(p, v_1) + (1 - \beta_2(p))(v_1 - p) \qquad (11) \\ &= \pi_1^a(p, v_1) + \Phi\left(\frac{[c - \pi_2^a(p)] - \mu}{\sigma}\right) \\ &\quad \cdot [v_1 - p - \pi_1^a(p, v_1)]. \qquad (12) \end{aligned}$$

For the preemptive bid strategy to be an equilibrium, it must be that the prescribed preemptive bid maximizes expected profits for a first bidder with that valuation. The necessary first-order condition is given by

$$\frac{\partial \pi_1(p, v_1)}{\partial p} = 0, \qquad (13)$$

$$\begin{aligned} &\frac{\partial \pi_1^a(p, v_1)}{\partial p} - \phi(\gamma(p)) \frac{\partial \pi_2^a(p)/\partial p}{\sigma}[v_1 - p - \pi_1^a(p, v_1)] \\ &\quad - \Phi(\gamma(p))\left(1 + \frac{\partial \pi_1^a(p, v_1)}{\partial p}\right) = 0, \qquad (14) \end{aligned}$$

where $\gamma(p) = (\epsilon_2^* - \mu)/\sigma$.

Since

$$\frac{\partial \pi_2^a(p)}{\partial p} = -(1 - v^{-1}(p))\frac{\partial v^{-1}(p)}{\partial p}$$

and

$$\frac{\partial \pi_1^a(p, v_1)}{\partial p} = -p,$$

Equation (14) can be rewritten as follows:

$$\begin{aligned} &-p + \phi(\gamma(p))\frac{(1 - v^{-1}(p))(\partial v^{-1}(p)/\partial p)}{\sigma} \\ &\quad \cdot [v_1 - p - \pi_1^a(p, v_1)] - \Phi(\gamma(p))(1 - p) = 0. \end{aligned}$$

If this is in equilibrium, then it must be that $v^{-1}(p) = v_1$, utilizing the fact that $\partial v^{-1}(p)/\partial p = 1/p'(v_1)$, and solving for $p'(v_1)$, we arrive at the following differential equation:

$$\begin{aligned} &p'(v_1) \\ &= \frac{\phi(\gamma(p(v_1)))(1/\sigma)(1 - v_1)[(v_1 - p(v_1))(1 - (v_1 + p(v_1))/2)]}{p(v_1) + \Phi(\gamma(p(v_1)))(1 - p(v_1))}. \qquad (15) \end{aligned}$$

Although this differential equation does not readily admit an analytic solution, it can be solved for numerically. Let $p^*(v_1)$ be the solution to the differential Equation (15). Given this solution, we can now move to the earlier stage, where bidder one makes her entry decision. To treat both players symmetrically, we assume this decision to be noisy as well. Therefore, the first bidder's expected payoff from entry (e_1) is given by

$$\pi_1(e_1) = \int_0^1 \pi_1(p^*(v_1), v_1)\, dv_1 - c + \epsilon_1,$$

where $\pi_1(p, v)$ is given by Equation (12).

Because the payoff from nonentry is 0, bidder one will decide to enter only if

$$\int_0^1 \pi_1(p^*(v_1), v_1)\, dv_1 - c + \epsilon_1 \geq 0.$$

This, of course, means that bidder one will only enter if

$$\epsilon_1 \geq c - \int_0^1 \pi_1(p^*(v_1), v_1)\, dv_1.$$

The (ex ante) entry probability for bidder one is then given by

$$\begin{aligned} \beta_1 &= \Pr\left(\epsilon_1 \geq c - \int_0^1 \pi_1(p^*(v_1), v_1)\, dv_1\right) \\ &= 1 - \Phi\left(\frac{[c - \int_0^1 \pi_1(p^*(v_1), v_1)\, dv_1] - \mu}{\sigma}\right). \qquad (16) \end{aligned}$$

The entry probabilities of the two bidders β_1 Equation (16) and β_2 Equation (7) given the preemptive bid $p^*(v_1)$ characterize the equilibrium under noisy costly entry and can be utilized to calculate expected revenue and profit results for the seller and both bidders.

Note that this is in contrast to the result of the standard theory of Bulow and Klemperer (2009) where there exists a partially pooling equilibria. The reason that such an equilibrium fails to exist in our setting is that increases in the preemptive bid by the first bidder will always have a measurable impact on the likelihood of entry by the second bidder (by changing the cutoff level ϵ_2^*). This provides sufficient incentive for high-valuing first bidders to attempt to differentiate themselves by placing a higher preemptive bid. In contrast, under the standard theory, any increase of bids beyond the one specified in the equilibrium will only have a negative impact for first bidders since the second bidder is already not entering for sure, so a

[24] Note that the c and ϵ_1 terms are dropped from these equations because they are sunk at the time of preemptive bid decision making. This is primarily done for notational simplicity when considering the first-bidder-entry decision.

higher preemptive bid only increases the price that the first bidder will pay.

The model of noisy bidder cost of entry replicates many of the features observed in the experimental data. In the auction, bidders fail to enter all the time due to high idiosyncratic cost draws in our model. In the sequential mechanism, first bidders place preemptive bids that are positively correlated with their own value, and second bidders, having observed any preemptive bid, still enter with a positive (ex ante) probability. Next, we proceed by using maximum-likelihood estimation to identify parameters (distributions of ϵ_i) that best fit the observed experimental data.

4.2. Parameter Estimation

In this section, we estimate the parameters that define the distribution of ϵ, (μ, σ), that best fit our experimental data. We use maximum-likelihood estimation (MLE) for this purpose. We take a progressive approach to the estimation process, first fitting a common set of (μ, σ) across both institutions, the auction and the sequential mechanism, and then allowing (μ, σ) to vary between the auction and the sequential mechanism. Then we allow for the noise terms for the two bidders to come from different distributions. One could argue that the first bidder's entry decision is simpler than the second bidder's, because he does not have to think about the preemptive bid, so that the first bidder's term may come from a distribution with μ closer to 0, and a smaller σ. In contrast, the second bidder must interpret the first bidder's preemptive bid, and this complexity may cause the second bidder's σ to be large. Knowing that failure to enter is sure to result in the first bidder earning higher profit may trigger some inequality aversion, and $\mu > 0$ may be a reasonable approximation for modeling it.[25]

Let t denote a single decision period, $t = 1, \ldots, \mathcal{T}$, where $\mathcal{T}$ represents the total number of entry decision periods. If we assume a common (μ, σ) across institutions, then the joint likelihood function is given by

$$L(\mu, \sigma) = \prod_{t \in \mathcal{T}} (\beta_2^{e_{2t}} (1 - \beta_2)^{(1 - e_{2t})})(\beta_1^{e_{1t}} (1 - \beta_1)^{(1 - e_{1t})}),$$

where

$$e_{it} = \begin{cases} 1 & \text{if bidder } i \text{ enters in decision } t, \\ 0 & \text{otherwise.} \end{cases}$$

[25] Huck et al. (2001) report results of Stackelberg duopolies experiments to the literature survey. They find that leaders behave less aggressively and followers behave more aggressively than they should, and attribute it to aversion to disadvantageous inequality. It is not clear how the fact that our more complex setting affects these results, but it is likely that the presence of private information dampens down disadvantageous inequality aversion. Nevertheless, modeling social preferences is beyond the scope of this paper.

Table 5 MLE and LL Results for the Noisy Entry Model

Treatment	MLEs (μ, σ)	LL	Likelihood-ratio test
Both institutions	(0.405, 0.360)	−2,017	$\chi^2 = 114$ $p < 0.001$
Auction	(−0.031, 0.066)	−1,986	$\chi^2 = 52$ $p < 0.001$
Seqmech	(0.490, 0.405)		
Auction	(−0.031, 0.066)	−1,960	—
Seqmech: First bidder	(0.102, 0.115)		
Seqmech: Second bidder	(0.590, 0.532)		

Table 5 presents the MLE and log-likelihood (LL) results for our noisy entry model when μ and σ are fixed across both institutions, and when they are allowed to vary between the two mechanisms and first and second bidders. We have also provided the results from likelihood-ratio tests against the model with the most free parameters in an effort to identify the most efficient fit.

As one can see from the likelihood-ratio tests in Table 5, the estimation allowing parameters to vary between bidders is the best fit. For this estimation, starting from the bottom, we see that μ and σ are quite large for the second bidder, agreeing with our data that second bidders overenter with considerable noise. The first bidder, however, has a relatively low μ and σ, indicating a smaller benefit of entry along with less variability. This too matches up with our data, where first bidders consistently entered (and standard theory assumes an entry rate of 100%). The MLEs for the auction also seem to coincide with our data, where the fixed benefit of entry is actually slightly negative, $\mu = -0.031$, with $\sigma = 0.066$ accounting for the noise we see, pushing entry rates slightly below 100%. In short, the MLEs, overall, agree with our experimental results.

We show the predicted seller revenue levels, based on the MLEs that vary between bidders, for the noisy entry model in Table 6. We also show what the corresponding preemptive bid function looks like for these same estimates in Figures 6(a) and 6(b).

We see that the predicted revenues from our model closely match our data and that the preemptive bids fit reasonably well, certainly an improvement over the predictions of the Bulow and Klemperer (2009) model. The predicted revenues are not only close to the actual revenues in terms of the point predictions but also in terms of the qualitative comparison between the two institutions. Specifically, in line with our data and contrary to the Bulow and Klemperer (2009) theory, the bidder behavior model predicts higher revenues under the sequential mechanism than under the auction.

Figure 6 **Predicted Preemptive Bids for $\mu = 0.102$ and $\sigma = 0.115$ and Average Preemptive Bids**

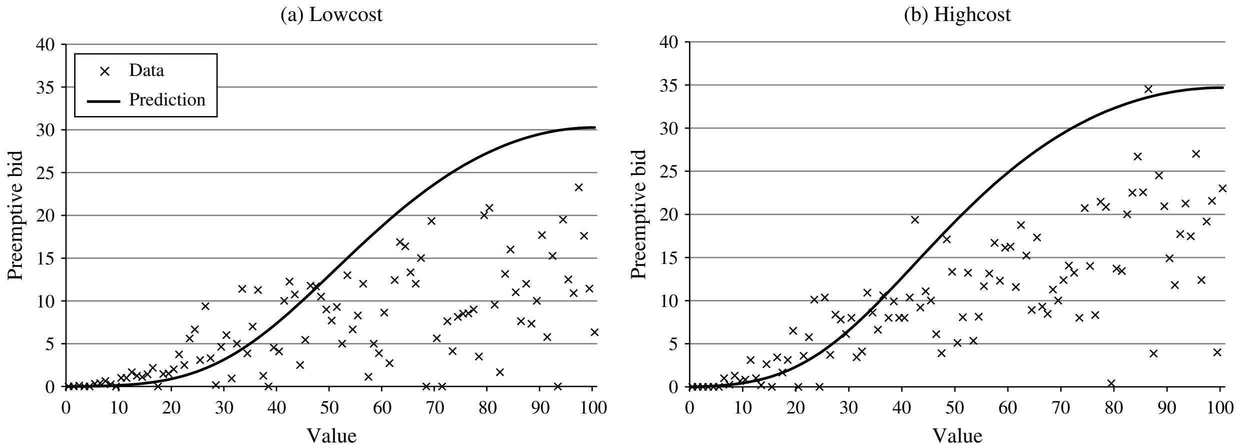

Table 6 **Comparison of Institution Revenue Between the Experimental Data and MLE Predictions for the Noisy Entry Model, Where the MLEs Are Allowed to Vary Across Bidders**

	Actual		Predicted	
	Auction	Seqmech	Auction	Seqmech
Lowcost	31.57	34.25	30.90	32.23
Highcost	26.61	30.30	25.63	31.07

5. Conclusion

To the best of our knowledge, our paper is the first to compare the performance of an auction and a sequential mechanism in a controlled laboratory setting. We design our experiments to closely match the Bulow and Klemperer (2009) model. We find that the average seller revenue in the auction is slightly lower than what the theory predicts, and the average sequential mechanism revenue is significantly higher, especially in the treatment with high entry costs. More importantly, when comparing the two institutions to each other, these results generate seller revenue that is higher in the sequential mechanism than the auction, contrary to standard theory.

Our experiments also demonstrate that individual behavior can vary significantly from the strong predictions of standard game theory. Although individual variations from theoretical predictions are certainly not surprising, especially given results reported in prior experimental literature on signaling, we demonstrate that those variations can be sufficient to reverse the normative prescriptions of the theory.

The behavior we observe in our experiment differs from the model predictions in three ways. First, bidders do not enter the auction 100% of the time, causing auction revenues to be somewhat lower than predicted by the model. Second, in the sequential mechanism, we find that the first bidders do not set preemptive bids according to the threshold strategy. Instead, both the probabilities of setting positive preemptive bids and the magnitudes of these bids increase with the first bidders' values. Third, second bidders in the sequential mechanism tend to over-enter in response to high preemptive first-bidder bids.

Our main contribution is a new model we developed that incorporates the noisy entry behavior and uses MLE techniques to estimate model parameters for our data. We find that the model organizes our data reasonably well, in that it matches revenues fairly closely. Although our model is quite consistent with the data, we recognize that there might be other behavioral factors that we have not accounted for (such as limited rationality regarding the informational content of the preemptive bid) that might also be playing a role in the experiments. Rather, our results are a warning that a mechanism designer might want to consider the robustness of their results to many possible behavioral phenomena. The formal incorporation of nonstandard behavior into the design and selection of mechanism is, in our opinion, an exciting and challenging avenue for future research.

One limitation of our study is that it does not directly incorporate reserve prices (which was done to closely follow the Bulow and Klemperer 2009 model). Davis et al. (2011) show that subjects in auctions do not set reserve prices correctly in the laboratory; however, our experimental findings do highlight the potential importance of reserve prices in auctions. Specifically, in our setting, less than full entry in the auction means that the seller may receive zero revenue without a reserve price, whereas in the sequential mechanism, any irrational behavior by a first bidder results in a reserve price in essence.

The main managerial implication that comes from our study is that sequential mechanisms may well represent a viable alternative to auctions for a variety of applications. Not only do bidders prefer them, but they may actually be better for the sellers as well. Therefore, for powerful sellers, sequential mechanisms may well represent a viable alternative to auctions, one that both sellers and bidders prefer.

Additionally, in a setting with costly entry, sequential mechanisms are also more efficient than auctions, because fewer potential bidders end up paying the entry fees unnecessarily. Thus, further theoretical and empirical work is called for to better understand sequential mechanisms.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2013.1800.

Acknowledgments

The authors thank three anonymous referees, an associate editor, and the department editor for their feedback. The authors especially recognize one referee whose input greatly improved the theoretical aspects of the paper. The authors thank seminar participants at the University of Tennessee, the University of Arizona, Chapman University, and Harvard University. This research was funded by the National Science Foundation [Award SES 0849054], the Deutsche Forschungsgemeinschaft through the DFG-Research Group "Design and Behavior," and the Small Grants Program of Pennsylvania State University, Smeal College of Business.

Appendix A. Risk Aversion

In this section, we provide a number of results related to the incorporation of risk aversion to the standard model of behavior in both the sequential mechanism and the auction. For notational simplicity, we refer to random variables with capital letters so that the bidder's values can be described as V_i, which is assumed to be distributed uniformly on the support $[0, 1]$. Importantly, a bidder's profit in the auction can be expressed as a random variable denoted by Π_i^a, which is defined based on V_1 and V_2 as follows:

$$\Pi_i^a = \begin{cases} v_i - v_j & v_i > v_j, \\ 0 & v_i \leq v_j. \end{cases}$$

Also, we assume that bidders have common Bernoulli utility functions denoted by u, where u is assumed to be strictly increasing $u' > 0$ and concave $u'' \leq 0$. Since expected utility functions are unique up to an affine transformation, we set $u(-c) = 0$ without loss of generality.

The equilibrium conditions for the sequential mechanism expressed in §2.1 can be similarly expressed for risk-averse preferences. The deterring value is determined by the amount that makes the second bidder indifferent between entry or not:

$$\mathrm{Eu}(\Pi_2^a - c \mid V_1 \geq v_s) = u(0). \tag{A1}$$

Given this deterring value, the equilibrium preemptive bid is the amount that makes a first bidder with a value of v_s indifferent between placing a preemptive bid or not, or

$$\mathrm{Eu}(\Pi_1^a - c \mid V_1 = v_s) = u(v_s - p^* - c). \tag{A2}$$

Note that, since the right-hand side of both these equations represents certain amounts, an examination of varying risk preferences boils down to an examination of differing certainty equivalents.

Under the auction mechanism, consider the simultaneous entry game originally depicted in Figure 1 and modified in Figure A.1 to allow for risk aversion. Since $V_i - c$ first-order stochastically dominates $\Pi_i^a - c$, then for all risk-averse preferences we know that $\mathrm{Eu}(V_i - c) \geq \mathrm{Eu}(\Pi_i^a - c)$. If $\mathrm{Eu}(V_i - c) \geq \mathrm{Eu}(\Pi_i^a - c) > u(0)$, then, as in the risk-neutral case, it is a dominant strategy for each player to enter the auction, and, therefore, revenue is unaffected by risk aversion. If $u(0) > \mathrm{Eu}(V_i - c) \geq \mathrm{Eu}(\Pi_i^a - c)$, then each player has a dominant strategy to not enter and auction revenue is 0. Only if $\mathrm{Eu}(V_i - c) \geq u(0) \geq \mathrm{Eu}(\Pi_i^a - c)$ can there exist a purely mixed strategy between entry and not. In this case, there are two pure-strategy Nash equilibria where one bidder enters and the other does not.

Figure A.1 **Payoff Table from Auction Entry Stage with Risk-Averse Preferences**

		Bidder 2: Enter	Bidder 2: Not enter
Bidder 1	Enter	$\mathrm{Eu}(\Pi_1^a - c)$, $\mathrm{Eu}(\Pi_2^a - c)$	$\mathrm{Eu}(V_1 - c)$, $u(0)$
	Not enter	$u(0)$, $\mathrm{Eu}(V_2 - c)$	$u(0)$, $u(0)$

Since in any mixed-strategy equilibrium, each of the pure strategies in the support of the mixed strategy must be a best response to the mixed strategy of the other player, it must be in this case that

$$u(0) = q\,\mathrm{Eu}(\Pi_i^a - c) + (1 - q)\,\mathrm{Eu}(V_i - c),$$

where $q > 0$ is the probability the other bidder enters. This implies a symmetric mixed-strategy equilibrium is given by solving for q above to yield

$$q = \frac{u(0) - \mathrm{Eu}(V_i - c)}{\mathrm{Eu}(\Pi_i^a - c) - \mathrm{Eu}(V_i - c)},$$

which is only satisfied by $q > 0$ if $\mathrm{Eu}(V_i - c) > u(0) > \mathrm{Eu}(\Pi_i^a - c)$.

To see that this implies a deterring value of $v_s = 0$ in the sequential mechanism, note that $u(0) > \mathrm{Eu}(\Pi_i^a - c)$ obviously implies $u(0) > \mathrm{Eu}(\Pi_2^a - c \mid V_1 \geq v_s)$ for all v_s so that the second bidder would prefer to not enter under the sequential mechanism upon knowing that first bidder had entered. In turn, the first bidder would have no incentive to place a preemptive bid.

Changes in the relative value of the zero payoff needed to induce this behavior can easily be accomplished by increases in risk aversion, as defined by Pratt (1964). Let $\mathrm{ce}(X; u)$ be the certainty equivalent for the gamble described by the random variable X for a player with risk-averse utility function u, or the amount such that $\mathrm{Eu}(X) = u(\mathrm{ce}(X; u))$. If another risk-averse utility function v is "more risk averse" than u, then we know that for all X $\mathrm{ce}(X; v) \leq \mathrm{ce}(X; u)$, or certainty equivalents shift down as risk aversion increases. Since all the shifts are relative to the certain amount of 0, increases in risk aversion will accomplish changes in the relevant inequalities as expressed above.

Appendix B. Uniqueness of Equilibrium

In this section, we proceed as do Roberts and Sweeting (2013), and use existing results on sequential equilibria of signaling games to verify that the revealing equilibrium identified in the §4.1 is indeed the unique perfect sequential equilibrium under the D1 refinement concerning out of equilibrium beliefs.

Proposition 1. *There exists a unique equilibrium preemptive bid function and cutoff values in the sequential mechanism under the D1 refinement.*

To prove this proposition, consider the four main "stages" of the described equilibrium.

Stage 1. The first bidder arrives and observes her private cost of entry ϵ_1 and decides to enter or not.

Stage 2. The first bidder learns her private value (v_1) and may place preemptive bid p, which might reveal information about v_1.

Stage 3. The second bidder arrives and observes the preemptive bid (p) and his private cost of entry ϵ_2 and decides to enter or not.

Stage 4. If the second bidder enters, he learns his value (v_2) and bids in an English auction with a starting price of p. Note that it is only the actions in Stages 2 and 3 that may be impacted by the signaling nature of the game. However, the other stage strategies must obviously satisfy sequential rationality. Therefore, we work backward to establish our equilibrium result.

In Stage 4, given that each player has a dominant strategy to bid his or her value, the payoffs from the auction are determined accordingly. Namely, the ex post auction profits in equilibrium are given by

$$\pi_i^a(p, v_i, v_j) = \begin{cases} v_i - v_j & v_i \ge v_j \ge p, \\ v_i - p & v_i \ge p > v_j, \\ 0 & \text{otherwise.} \end{cases}$$

In Stage 3, given a revealing preemptive bid equilibrium (or any beliefs that the second bidder has about the first bidder's value as a result of Stage 2), the second bidder's expected utility from entry is given by $\pi_2^a + \epsilon_2 - c$, which is strictly increasing in ϵ_2 so a cutoff strategy will be the unique equilibrium given beliefs about auction profits defined by π_2^a.

In Stage 2, we need to show that given the cutoff strategy of the second bidder in Stage 3 and dominant strategy bidding strategies in Stage 4, the differential equation defined by Equation (15) is the unique sequentially rational equilibrium bid function satisfying the D1 refinement. We utilize the results of Mailath and von Thadden (2013), which generalizes the earlier results of Mailath (1987) to the situation examined in this paper. Namely, they provide sufficient conditions for the incentive-compatible, separating strategy (e.g., preemptive bid function) to be differentiable and therefore characterized by the differential equation. In particular, they examine the case where both the private information type space and domain of the strategy space are compact, as is the case in our situation since the private information types space is the values of the first bidder and is the given by the unit interval $[0,1]$ and the set of admissible preemptive bids is also the unit interval $[0,1]$.

We redefine the interim expected profit function (Equation (12)) of the first bidder to be given by

$$\Pi_1(p, v_1, \hat{v}_1) = \pi_1^a(p, v_1) + \Phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right) \cdot [v_1 - p - \pi_1^a(p, v_1)], \tag{B1}$$

where v_1 is the bidder's actual value, $\hat{v}_1$ is the value the second bidder believes the first bidder to have, and p is the preemptive bid. As defined in Equation (5), we have

$$\pi_1^a(p, v_1) = \frac{v_1^2 - p^2}{2}$$

and

$$\pi_2^a(\hat{v}_1) = \frac{(1-\hat{v}_1)^2}{2}.$$

The function Π_1 defined by Equation (B1) is obviously twice continuous and differentiable in each argument. We begin by checking that the two assumptions of Mailath and von Thadden (2013) are satisfied in our problem.

1. The first-best (full-information) contracting problem has a unique solution. In our case, if the second bidder is fully informed of v_1, then it is always optimal for the first bidder to place a preemptive bid of 0 so that the unique solution is $p(v_1) = 0$ for all $v_1 \in [0,1]$.
2. The second-order condition is satisfied, or

$$\frac{\partial^2 \Pi_1}{\partial p^2} < 0$$

for all v_1 evaluated at the first-best optimal solution. To see that this is true, note that

$$\begin{aligned} \frac{\partial \Pi_1}{\partial p} &= \frac{\partial \pi_1^a(p, v_1)}{\partial p} + \Phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right)\left[-1 - \frac{\partial \pi_1^a(p, v_1)}{\partial p}\right] \\ &= -p + \Phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right)[p-1] \\ &= p\left(\Phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right) - 1\right) - \Phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right), \end{aligned} \tag{B2}$$

so the second partial derivative is given by

$$\frac{\partial^2 \Pi_1}{\partial p^2} = \Phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right) - 1,$$

which is strictly negative for any value of $\hat{v}_1$ since $\Phi(\cdot) < 1$ for all values because it is the cdf of the standard normal distribution.

Given that the assumption of Mailath and von Thadden (2013) are satisfied, we can then apply Theorem 3 for a sufficient condition for differentiability of $p(v_1)$. In particular, we need that $\partial \Pi_1(p, v_1, v_1)/\partial p \neq 0$ for all $p \in [0,1]$. Examining Equation (B2), this is clearly satisfied for all p since this value is strictly negative. Since $p(v_1)$ is differentiable, it is continuous. Theorem 6 of Mailath and von Thadden (2013) shows that a standard single-crossing property implies incentive compatibility. In particular, since $p'(v_1) > 0$ and $\partial \Pi_1/\partial \hat{v}_1 \le 0$, for all $v_1 \in (0,1)$ and $\hat{v}_1 \in (0,1)$, it must be that

$$\frac{\partial}{\partial v_1}\left\{\frac{(\partial \Pi_1(p(\hat{v}_1), v_1, \hat{v}_1))/\partial p}{(\partial \Pi_1(p(\hat{v}_1), v_1, \hat{v}_1))/\partial \hat{v}_1}\right\} \le 0.$$

Note that

$$\frac{\partial \Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)}{\partial \hat{v}_1} = -\frac{1}{\sigma}\phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right)[v_1 - p - \pi_1^a(p, v_1)](1-\hat{v}_1)$$

so that

$$\frac{\partial^2 \Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)}{\partial \hat{v}_1 \partial v_1} = -\frac{1}{\sigma}\phi\left(\frac{\pi_2^a(\hat{v}_1) - \mu}{\sigma}\right)(1-v_1)(1-\hat{v}_1),$$

which is strictly negative since $\phi(\cdot)$ is the strictly positive density of the normal distribution. Also, note from Equation (B2) that this derivative is not a function of v_1, so $\partial^2\Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)/\partial p \partial v_1 = 0$. Taking these results together, we have that

$$\frac{\partial}{\partial v_1}\left\{\frac{\partial\Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)/\partial p}{\partial\Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)/\partial\hat{v}_1}\right\} = \frac{-((\partial\Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)/\partial p)(\partial^2\Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)/\partial\hat{v}_1\partial v_1))}{(\partial\Pi_1(p(\hat{v}_1), v_1, \hat{v}_1)/\partial\hat{v}_1)^2},$$

which is strictly negative.

Although this result has established that the separating equilibrium characterized by the differential equation is the unique separating equilibrium to this problem, we have not ruled out the possible coexistence of pooling equilibria. Ramey (1996) shows that all equilibria satisfying the D1 refinement will be separating equilibria as long as two conditions are satisfied. The first is that the single-crossing property established above is satisfied. The second is that no first bidder can want to play the highest possible preemptive bid. It is clear in our case that no bidder would want to place a preemptive bid of 1 since it guarantees payoffs of 0 (or less). Thus, it follows that the revealing preemptive bidding strategy described is the unique equilibrium strategy satisfying the D1 requirement at this stage.

Finally, in Stage 1, The expected profit from entry for the first bidder is given by

$$\pi_1(e_1) = \int_0^1 \pi_1(p^*(v_1), v_1)\, dv_1 - c + \epsilon_1,$$

which is strictly increasing in ϵ_1 so the cutoff strategy described will be the unique equilibrium of this stage.

Thus, we have identified a unique perfect sequential equilibrium under the D1 refinement.

Appendix C. Cohort-Level Results for Revenue

Treatment	Cohort	Auction	Seqmech
Lowcost ($c = 3$)	1	33.63	29.86
	2	32.07	35.72
	3	29.99	36.31
	4	32.19	29.97
	5	34.18	39.98
	6	30.64	36.02
	7	25.27	38.53
	8	34.66	36.02
	9	30.87	27.30
	10	32.18	32.78
	Average	31.57	34.25
Highcost ($c = 10$)	1	25.65	30.73
	2	27.83	36.84
	3	25.61	34.78
	4	30.20	27.87
	5	26.38	22.99
	6	23.78	32.12
	7	30.70	28.14
	8	24.90	29.51
	9	23.90	25.39
	10	27.17	34.64
	Average	26.61	30.30

References

Arnold MA, Lippman SA (1995) Selecting a selling institution: Auctions versus sequential search. *Econom. Inquiry* 33(1):1–23.

Avery C (1998) Strategic jump bidding in English auctions. *Rev. Econom. Stud.* 65(2):185–210.

Banks JS, Sobel J (1987) Equilibrium selection in signaling games. *Econometrica* 55(3):647–661.

Berkshire Hathaway (2009) Berkshire Hathaway 2008 annual report. Berkshire Hathaway, Omaha, NE. http://www.berkshirehathaway.com/2008ar/2008ar.pdf.

Bernhardt D, Scoones D (1993) Promotion, turnover, and preemptive wage offers. *Amer. Econom. Rev.* 83(4):771–791.

Binswanger HP (1980) Attitudes toward risk: Experimental measurement in rural India. *Amer. J. Agricultural Econom.* 62(3): 395–407.

Bolton G, Ockenfels A (2000) ERC: A theory of equity, reciprocity, and competition. *Amer. Econom. Rev.* 90(1):166–193.

Bulow J, Klemperer P (2009) Why do sellers usually prefer auctions? *Amer. Econom. Rev.* 99(4):1544–1575.

Campo S, Guerre E, Perrigne I, Vuong Q (2011) Semiparametric estimation of first-price auctions with risk-averse bidders. *Rev. Econom. Stud.* 78(1):112–147.

Chen K-Y, Plott CR (1998) Nonlinear behavior in sealed bid first-price auctions. *Games Econom. Behav.* 25(1):34–78.

Cho I-K, Kreps DM (1987) Signaling games and stable equilibria. *Quart. J. Econom.* 102(2):179–222.

Conlisk J (1996) Why bounded rationality? *J. Econom. Literature* 34(2):669–700.

Cooper DJ, Garvin S, Kagel JH (1997a) Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *Econom. J.* 107(442):553–575.

Cooper DJ, Garvin S, Kagel JH (1997b) Signalling and adaptive learning in an entry limit pricing game. *RAND J. Econom.* 28(4):662–683.

Cox JC, Oaxaca RL (1996) Is bidding behavior consistent with bidding theory for private value auctions? *Res. Experiment. Econom.* 6:131–148.

Cox JC, Roberson B, Smith VL (1982) Theory and behavior of single object auctions. *Res. Experiment. Econom.* 2:1–43.

Davis AM, Katok E, Kwasnica AM (2011) Do auctioneers pick optimal reserve prices? *Management Sci.* 57(1):177–192.

Denton JR (2008) Stacked deck: Go-shops and auction theory. *Stanford Law Rev.* 60(5):1529–1554.

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

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

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

Fishman MJ (1988) A theory of preemptive takeover bidding. *RAND J. Econom.* 19(1):88–101.

Goeree JK, Holt CA (2004) A model of noisy introspection. *Games Econom. Behav.* 46(2):365–382.

Goeree JK, Holt CA, Palfrey TR (2002) Quantal response equilibrium and overbidding in private-value auctions. *J. Econom. Theory* 104(1):247–272.

Hirshleifer D, Png IPL (1989) Facilitation of competing bids and the price of a takeover target. *Rev. Financial Stud.* 2(4):587–606.

Holt CA, Laury SK (2002) Risk aversion and incentive effects. *Amer. Econom. Rev.* 92(5):1644–1655.

Huck S, Müller W, Normann H-T (2001) Stackelberg beats Cournot: On collusion and efficiency in experimental markets. *Econom. J.* 111(474):749–765.

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

Kagel JH (1995) A survey of experimental research. Kagel JH, Roth AE, eds. *Handbook of Experimental Economics* (Princeton University Press, Princeton, NJ), 501–585.

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

Kwasnica AM, Katok E (2007) The effect of timing on bid increments in ascending auctions. *Production Oper. Management* 16(4):483–494.

Mailath GJ (1987) Incentive compatibility in signaling games with a continuum of types. *Econometrica* 55(6):1349–1365.

Mailath GJ, von Thadden E-L (2013) Incentive compatibility and differentiability: New results and classic applications. *J. Econom. Theory* 148(5):1841–1861.

McKelvey RR, Palfrey TR (1995) Quantal response equilibria for normal form games. *Games Econom. Behav.* 10(1):6–38.

McKelvey RR, Palfrey TR (1998) Quantal response equilibria for extensive form games. *Experiment. Econom.* 1(1):9–41.

Pratt JW (1964) Risk aversion in the small and in the large. *Econometrica* 32(1/2):122–136.

Ramey G (1996) D1 signaling equilibria with multiple signals and a continuum of types. *J. Econom. Theory* 69(2):508–531.

Roberts JW, Sweeting A (2013) When should sellers use auctions? *Amer. Econom. Rev.* 103(5):1830–1861.

Simon HA (1984) *Models of Bounded Rationality* (MIT Press, Cambridge, MA).

Stephenson R, Jones M, Di Lapigio S (2006) Surviving M&A auction rigors. *Mergers and Acquisitions: The Dealermaker's J.* 41(12):28–36.

Subramanian G (2008) Go-shops vs. no-shops in private equity deals: Evidence and implications. *Bus. Lawyer* 63(3):729–760.