

Experimental Evidence on the Endogenous Entry of Bidders in Internet Auctions

David H. Reiley¹

<reiley@eller.arizona.edu>

This Version: September 2004

Abstract

This paper tests the empirical predictions of recent theories of the endogenous entry of bidders in auctions. Data come from a field experiment, involving sealed-bid auctions for collectible trading cards over the Internet. Manipulating the reserve prices in the auctions as an experimental treatment variable generates several results. First, observed participation behavior indicates that bidders consider their bid submission to be costly, and that bidder participation is indeed an endogenous decision. Second, the participation is more consistent with a mixed-strategy entry equilibrium than with a deterministic equilibrium. Third, the data reject the prediction that the profit-maximizing reserve price is greater than or equal to the auctioneer's salvage value for the good, showing instead that a zero reserve price provides higher expected profits in this case.

¹ Department of Economics, the University of Arizona. I wish to thank Mike Urbancic, Marius Hauser, and Mary Lucking for their research assistance, and Skaff Elias for product information about *Magic: the Gathering*. I would like to thank J.S. Butler, Rachel Croson, Glenn Ellison, Elton Hinshaw, Dan Levin, Kip King, Preston McAfee, Rob Porter, and Jennifer Reinganum for advice and constructive criticism.

1 Introduction

The earliest theoretical models of auctions assumed a fixed number N of participating bidders, with the number commonly known to the auctioneer and the participating bidders. More recent models have relaxed this assumption, considering the possibility of costly bidder participation, so that the actual number of participating bidders is an endogenous variable in the model. In this paper, I use a field experiment, auctioning several hundred collectible trading cards in an existing market on the Internet, to test the assumptions and the predictions of models of auctions with endogenous entry.

I concentrate on three empirical questions in this paper. First, can an experiment turn up evidence of endogenous entry behavior in a real-world market? The answer to this question appears to be yes. Second, given the existence of endogenous entry, does the entry equilibrium appear to be better modeled as stochastic, or as deterministic? Evidence from the experiment indicates that the stochastic equilibrium concept is a better model of behavior. Third, is it possible to verify the theory of McAfee, Quan, and Vincent (1998, henceforth, MQV), that even with endogenous bidder entry, the optimal reserve price for the auctioneer to set is at least as great as the auctioneer's salvage value? The answer to this question is "no," as a reserve price of zero appears to provide higher expected profits than a reserve price at the auctioneer's salvage value.

The field-experiment methodology of this study, that of auctioning real goods in a preexisting market, represents a hybrid between traditional laboratory experiments and traditional field research which takes the data as given. It shares with laboratory experiments the important advantage of allowing the researcher to control certain variables of interest, rather than leaving the researcher subject to the vagaries of the actual marketplace. (The key experimental treatment in this paper is the manipulation of the reserve price across auctions, to observe how

participants react in their entry and bidding decisions.) It shares with traditional field research the advantage of studying agents' behavior in a real-world environment, rather than in a more artificial laboratory setting.

Although the experimental literature on auctions is vast,² almost all of these studies have imposed an exogenous number of bidders (determined by the experimenter). Three exceptions are Smith and Levin (2001), Palfrey and Pevnitskaya (2003), and Cox, Dinkin, and Swarthout (2001). Smith and Levin (2001) and Palfrey and Pevnitskaya (2003) design their experiments to determine whether the entry equilibrium which obtains is deterministic or stochastic, a question I also investigate in this paper. Cox, Dinkin, and Swarthout (2001) show that when participation in a common-value auction is costly, winner's-curse effects are attenuated.

In the empirical literature on auctions in the field,³ one recent study considers endogenous entry. Bajari and Hortacsu (2003) note that in eBay auctions for coin proof sets, the number of observed bidders is positively correlated with the book value of the item and negatively correlated with the minimum bid for the item. From this they infer that bidding is costly, and they therefore provide a structural econometric model of bidding that includes an endogenous entry decision. The present paper adds to the empirical and experimental literatures on the endogenous entry of bidders by conducting a controlled experiment to gather evidence on the type of endogenous entry found in a real-world market.

The paper is organized as follows. The next section describes the relevant aspects of endogenous-entry auction theory, focusing on the testable implications. The third section describes the marketplace where the experiment took place. The following twin subsections explain the respective designs of the two sets of experiments. The fifth section presents the results, and a

² See Kagel (1995) for a review of auction experiments.

³ See Hendricks and Paarsch (1995) for a review of empirical work on auctions.

sixth section concludes.

2 Theoretical Background

Recently, there have been a number of important extensions to Vickrey's (1961) original model of auctions with a fixed, known number of bidders. The earliest examples of endogenous-entry bidding models include Samuelson (1985), Engelbrecht-Wiggans (1987), and McAfee and McMillan (1987). In these models, bidders have some cost to participating (either the research required to learn one's value for the good, or the effort required to decide on a bid and submit it). This cost causes some potential bidders to stay out of the auction entirely, and can cause other effects as well. For example, Samuelson (1985) and Engelbrecht-Wiggans (1987), making different modelling assumptions, both find that endogenous entry drives down the auctioneer's optimal reserve price relative to a model of costless entry. One of the goals of the present paper is to demonstrate the existence of entry costs in a real-world auction market.

McAfee and McMillan (1987) model bidder entry as a pure-strategy, asymmetric Nash equilibrium. In these models, exactly n bidders enter the auction (out of a total of $N > n$ potential bidders), and n is determined endogenously from the other parameters of the model (the auction format, the degree of affiliation of bidder values, the cost of entry, and so on). Alternatively, others have modeled a mixed-strategy, symmetric entry equilibrium (Engelbrecht-Wiggans (1987), Levin and Smith (1994), MQV (1998)). In the mixed-strategy models, bidders each enter with probability ρ , where ρ is determined endogenously.⁴

Levin and Smith (1994) point out that the difference between pure-strategy (deterministic)

⁴ These models find simple, symmetric solutions by assuming that bidders decide whether to participate before they learn their valuations. In my auctions, it is reasonable to assume that participants had information about their valuations before making the entry decision, so the entry outcome might be asymmetric. An example of such an asymmetric model is given by Samuelson (1985), where only those bidders with high valuations participate in the auction, and the entry equilibrium is in pure strategies.

models and mixed-strategy (stochastic) ones has implications for social welfare: if entry is stochastic, then expected social surplus is decreasing in the number N of potential bidders. The reason is that the variance of the number n of actual entrants is increasing in N , and such variance is costly. In common-value auctions, then, it turns out that auctioneers can increase both social welfare and their own profits by using reserve prices to discourage entry.

In a separate paper, Smith and Levin (2001) perform an experiment in which they attempt to determine whether entry by bidders is stochastic or deterministic: they find evidence in favor of their stochastic hypothesis. However, the experimental procedure doesn't actually involve any auctions; rather, it assigns simulated auction payoffs by a lottery procedure.⁵ Palfrey and Pevnitskaya (2003) modify this experimental design to conduct a first-price sealed-bid auction after the entry decision. They observe that the same bidders tend to enter repeated auctions, indicating a pure- rather than mixed-strategy equilibrium. Pevnitskaya (2004) provides a theoretical model of heterogeneously risk-averse bidders to explain this observation. When some bidders are more risk-averse than others, and all bidders know this fact, the more risk-averse bidders stay out of the auction deterministically in order to collect a fixed payoff. Only the relatively less risk-averse bidders enter the auction, also deterministically.⁶ mixed-strategy equilibrium disappears in favor of a pure-strategy equilibrium the more risk-averse bidders stay out of the auction in favor of a fixed payoff, while relatively less risk-averse bidders enter the auction. In this paper, I attempt to provide evidence on the question of stochastic versus deterministic entry equilibria in a field environment.

⁵ Subjects made the decision whether or not to incur the cost c to enter. After the entry outcome was observed, each of the n entrants had a $1/n$ chance of winning the payoff for that round of the experiment.

⁶ An interesting implication of Pevnitskaya's model is that an auctioneer can actually make himself worse off by advertising a sealed-bid auction heavily. An increase in the number of potential bidders increases the self-selection effect, causing less and less risk-averse bidders to enter the auction, and thereby causing less aggressive bidding, as risk aversion is well-known to increase bids in a first-price sealed-bid auction.

MQV (1998) examine the effects of reserve prices where valuations are where bidder entry is endogenous and bidder valuations may be either affiliated. In their model, the auctioneer chooses a reserve price and announces her auction, together with the level of her reserve price, to N potential bidders. Bidders then decide whether or not to incur the participation costs, making a stochastic (mixed-strategy) entry decision. Next the participating bidders find out their private information about the value of the good, they submit their bids, and finally the auctioneer awards the good to the highest bidder. If no bidder chooses to enter and to bid at least the reserve price, then the auctioneer keeps the good for herself and earns some outside option utility, or “salvage value.” The main prediction of MQV is that the optimal reserve price is at least as high as the salvage value of the good. This is a testable prediction; raising the reserve price from some lower value to the expected salvage value of the good should raise revenues for the auctioneer.

To summarize, this paper will attempt to answer three main questions. First, are entry costs relevant in the Internet auction market where I ran my experiments? Second, is the entry equilibrium a deterministic or a stochastic one? Third, is the optimal reserve price at least as high as the auctioneer’s salvage value? Note that the first question is about an assumption of endogenous-entry, the second attempts to distinguish between two rival theories, and the third is a test of the empirical prediction of a specific model.

3 Experimental Design

For this experiment, I auctioned trading cards via first-price, sealed-bid auctions, varying the reserve prices across treatments. The data in this paper are the same as in Lucking-Reiley (1999). The experiments took place in 1995 in a pre-eBay online market for collectible cards from *Magic: the Gathering*, a game which has enjoyed great success since its launch in August 1993. In the game, players assume the roles of dueling wizards, each with their own libraries of

magic spells (represented by decks of cards) that may potentially be used against opponents. Cards are sold in random assortments, just like baseball cards, at retail stores ranging from small game and hobby shops to large chain retailers. The game's maker, Wizards of the Coast (now a division of Hasbro) has developed and printed thousands of distinct card types, each of which plays a slightly different role in the game.

As discussed in Lucking-Reiley (1999), soon after the introduction of *Magic*, players and collectors interested in buying, selling, and trading game cards began to use the Internet to find each other and carry out transactions. In a Usenet newsgroup dedicated to this purpose, traders used a variety of trading institutions, including negotiated trades of one card for another, sales at posted prices, and auctions of various formats, typically lasting multiple days.

Scarcity was one major determinant of transaction prices for cards, as some cards were printed in relatively low quantities, and some cards had gone out of print. The most common in-print cards were not worth trading over the Internet; their values were pennies or less. Cards designated "uncommon" but not "rare" traded for prices of ten cents to two dollars. Cards designated "rare" but still in print typically ranged in price from one to fifteen dollars. Out-of-print cards, depending on their initial scarcity and on other attributes, traded for as much as three hundred dollars. In this research project, I dealt only in out-of-print cards.

In addition to data generated in my own auctions, I also make use of contemporaneous market data from the weekly Cloister price list in this marketplace. Cloister was a card trader who wrote a computer program that automatically searched the marketplace newsgroup for each instance of each card name (with some tolerance for misspellings) and gathered data on the prices posted next to each card name in the newsgroup messages. It then computed statistics for each card, and automatically archived these data on the Internet as a public service for other

interested traders.. Each card's reported list price is a trimmed mean over hundreds or thousands of different observations on the newsgroup. Despite some problems with these data, discussed in Lucking-Reiley (1999) many card traders adopted the Cloister price list as a standard measure of card market value, so I adopt it as a useful measure in my own analysis.

This marketplace represented an exciting opportunity to run auction field experiments. For the experiments, I purchased several thousand dollars' worth of cards (also via the Internet), and auctioned them off while systematically manipulating the reserve prices in order to observe their effects on bidder participation and bidding behavior. Because in any given week there were dozens of auctioneers holding *Magic* auctions on the Internet, as an experimenter I was able to be a "small player" who did not significantly perturb the overall market.

I employed two distinct experimental designs to collect the data. The first design examines the effects of a binary variable: whether or not minimum bids were used. By auctioning the same cards twice, once with and once without minimum bids, it exploits within-card variation to find the effects of the treatment variable on bidding and entry behavior. The second design investigates the effects of a continuous variable: the reserve price level (expressed as a fraction of the Cloister reference price). The between-card variation provides information that can be used to test the MQV prediction about the optimal reserve price level.

4.1 Within-Card Experiments

The first part of the data collection consisted of two pairs of auctions. Each of the four auctions was a sealed-bid, first-price auction of several dozen *Magic* cards auctioned off individually. This simultaneous auction of many different goods at once, although not common in

other economic environments,⁷ is the norm for auctions of *Magic* cards on the Internet. Running auctions in this simultaneous-auction format thus made the experiment as realistic and natural as possible for the bidders, who see many other similar auctions in the Internet marketplace for cards.

Each auction lasted for one week, from the time the auction was announced to the deadline by which all bids had to be received. I announced each auction to potential bidders via two channels. First, I posted three announcements to the appropriate Internet newsgroup. For each auction, I posted a total of three newsgroup messages spaced evenly over the course of the week of the auction. Second, I solicited some bidders directly via email messages to their personal electronic mailboxes. My mailing list for direct solicitation was comprised of people who had already demonstrated their interest in auctions for *Magic* cards by participation in previous ones.

The paired-auction experiment proceeded as follows. First, I held an absolute auction (no minimum bid) for 86 different cards (one of each card in the Antiquities expansion set). The subject line of the announcement read “Reiley's Auction #4: ANTIQUITIES, 5 Cent Minimum, Free Shipping!” so that potential bidders might be attracted by the unusually low minimum bid per card, essentially zero. (A 5-cent minimum is effectively no minimum, since the auction rules also required all bids to be in integer multiples of a nickel.) After the one-week deadline for submitting bids had passed, I computed the highest bid on each card. To each bidder who had won one or more cards, I mailed (electronically) a bill for the total amount owed.⁸ After

⁷ Although simultaneous auctions are not traditional for familiar auctions, such as those of art, estate goods, or tulip bulbs, such formats have been used for timber and offshore oil auctions. The advent of computerized bidding appears to be making the simultaneous auction format even more common. In addition to the card auctions in this newsgroup market, simultaneous Web-based auctions are becoming common at commercial sites such as eBay, and a simultaneous-auction format was used for the recent FCC auctions of spectrum rights (see McMillan (1994) for details).

⁸ Although the standard practice in this marketplace is for auctioneers and other card sellers to charge buyers for postage and/or handling, I chose not to do this. I wanted bidders to bid independently, as much as possible, on each of the cards in which they were interested. Someone seriously interested in one card might

receiving a winner's payment via check or money order, I mailed them their cards. Almost no one defaulted on their winning bids.⁹

I also mailed a list of the winning bids to each bidder who had participated in the auction, whether or not they had won cards. This represented an effort to maintain my reputation as a credible auctioneer, demonstrating my truthfulness to those who had participated. I did not, however, give the bidders any explicit information about the number of people who had participated in the auction, or about the number of people who had received email invitations to participate.

After one additional week of buffer time after the end of the first auction, I ran the second auction in the paired experiment, this time with reasonably high minimum bid levels on each of the same 86 cards as before. The minimum bid levels were determined by consulting the standard (trimmed-mean) Cloister price list of *Magic* cards cited in section 3 of this paper, and setting the minimum bid level for each card equal to 90% of the value of that card from the price list.

This contrast in minimum bid levels (zero versus 90% of the Cloister price list) was the only

decide to bid higher on a second card in the same auction than they would if the cards were auctioned independently, because they would like to spread out the postage costs per card by purchasing more than one card simultaneously from the same source. In addition, some of the cards I auctioned had rather low values, and I wanted to avoid having the card values be swamped by the cost of shipping. For example, if a bidder won a single card for 20 cents and then had to pay a fixed 50-cent shipping charge on top of that, the amount of useful information which could be derived from her bid would be rather suspect. Therefore, in the interests of keeping bid data as clean as possible, I decided to pay postage costs myself, and announced in advance that first-class shipping was included in the amount of each bid.

⁹ A small number of winning bidders failed to pay for the cards they had won. In all, I received payment for 90% of the cards sold, constituting 89% of the reported revenue in the within-card auctions. Almost all of the "deadbeat" bidders were those who won only a single card, and who explained that they had originally hoped to win more cards, and didn't feel it was worth it to complete the transaction. I discouraged such behavior, but was unable to eliminate it. Only one or two individuals won multiple cards but failed to pay for them. Since none of the unpaid cards seemed to have outlandishly high winning bid amounts, I have taken the point of view that all bids were made in good faith, and have not excluded any observations from my analysis.

economically significant difference between the two auctions.¹⁰ By keeping all other conditions identical between the two auctions, I attempted to isolate the effects of minimum bids on potential bidders' behavior. One condition that could not be kept identical, unfortunately, was the time period during which the auction took place. Because the two auctions took place two weeks apart, there were potential differences between the auctions that might have affected bidder behavior. First, the demands for the cards (or the supplies by other auctioneers) might have changed systematically over time, which is a realistic possibility in such a fast-changing market as this one.¹¹ Second, since the auctions shared many of the same bidders, the results of the first auction may have affected the demand for the cards sold in the second auction.¹²

To control for such potential variations in conditions over time, I simultaneously ran the same experiment in reverse order, using a different sample of cards. This second pair of auctions each featured the 78 cards in the Arabian Nights expansion set, with minimum bids present in the first auction but absent in the second. Just as before, minimum bids were set at ninety percent of the market price level from the Cloister price list. The first auction in this pair began three days after the start of the first auction in the previous pair, so that the auctions in the two experiments overlapped in time but were offset by three days. Also, I used a larger mailing list for my email announcement in this pair of auctions (232 people) than I had for the previous pair of auctions

¹⁰ Both auctions lasted exactly seven days. The same 86 cards were up for bid in each auction. Each auction announcement was posted exactly three times to the marketplace newsgroup, and was emailed to primarily the same list of potential bidders. Even the subject line of the announcements and mailings was kept identical, except that in the second auction, the words "5 Cent Minimum" were removed.

¹¹ For example, certain cards from the Arabian Nights expansion set increased in value by a factor of ten during their first year out of print. It turns out that market prices for cards were actually rather stable during the month in which this experiment was conducted, but I did not know *a priori* what was going to happen to card prices.

¹² For example, suppose that a particular bidder is anxious to obtain a single Guardian Beast card for her deck, so that her valuation of the card is higher than that of any of the other bidders in the experiment. She may win the card in the first auction, and then have zero demand for that same card in the second auction. If this is generally the case for most cards, that the highest-value bidders in the sample are screened out in the first auction, then we might expect to see systematically lower revenues in the second auction.

(50 people), with the first mailing list being a subset of the second mailing list. Otherwise, all other conditions were identical between the two pairs of auctions. A sample auction announcement, as it appeared to the potential bidders both in electronic mail and in the marketplace newsgroup, is displayed in the Appendix.

Table 1 shows a set of summary statistics for each of the four auctions in the within-card experiments.¹³ The auctions are displayed in two pairs: first Auctions AA and AR, for the 86 Antiquities cards, and then Auctions BA and BR, for the 78 Arabian Nights cards.¹⁴ Auctions AA and BA were with no minimum bids, while Auctions AR and BR had sizable minimums (equal to 90% of the market price).

The table contains quite a bit of descriptive information about the auctions, including the number of participating bidders, the number of bids received, and the total payments received from winning bidders. Note two key points. First, “real money” was involved in the auction transactions. Of the 73 different bills I sent to winning bidders over the course of the experiment, the median payment amount for each auction was between \$10 and \$24. A few individual payments even exceeded \$100.

Second, in each auction there are multiple winners. The number of winners in each auction ranges from 6 to 27, and the fraction of bidders who win at least one card is between 40 percent

¹³ A note on mnemonics. The first letter represents the card set: A for Antiquities, B for Arabian Nights. The second letter is A for an absolute auction (reserve prices equal to zero), and R for an auction with positive reserve prices.

These auctions were part of a series of auctions run for a larger research program, so participating bidders saw me run several other auctions (not part of the research presented here) during the same time period. This had two advantages where the experimental design is concerned. First, it helped avoid drawing bidders' attention to the point of my research. (For example, during this period I also ran an English auction and a second-price auction and another first-price auction, with different sets of cards.) I feared that if they knew I was looking for the effects of reserve prices, it might distort their behavior (for example, they might consciously try to bid consistently from one auction to another). Second, it had the effect of making bidders unsure of what I would do next. In particular, I didn't want bidders to expect that I would always auction the same card twice, for it might distort their behavior if they knew they would have a second chance to bid on the same card.

¹⁴ A few of the auction items I denote as “cards” were actually groups of cards: either a sealed packet of out-of-print cards, or a set of common cards bundled together.

and 86 percent. In each auction, the median number of cards won by each winner is between 2 and 3.5, while the maximum number of cards won by a single bidder ranges from 12 to 26. Except in Auction AR, no winner won more than 29 percent of the cards sold in any single auction. (In Auction AR, participation was very low: only 7 people submitted bids, 6 of whom won at least one card, and 39 of the cards went unsold.) The biggest spender in any of the auctions won cards totalling \$316.50 of the total revenue of \$774.75 in Auction BA, generating 41 percent of the revenue despite winning no more than 15 percent of the cards_evidently, she was particularly interested in high-value cards. Thus, it is not the case that some people are the highest bidders on all cards in an auction, which suggests that a given bidder's valuations for different cards are at least somewhat independent. This gives some justification for reporting regression results in which each individual card bid is assumed to be an independent observation.

4.2 Between-Card Experiments

A second set of experiments was designed to examine the effects of changes in the *level* of the reserve price, rather than merely changes in the *existence* of reserve prices. Five first-price, sealed-bid auctions took place, each with a one-week timeframe for the submission of bids. Each was a simultaneous auction of many different items, this time with no overlap of items between auctions. Each card in the first four auctions (R1 through R4) had a posted reserve price. The fifth auction (R0) used a zero reserve price on every card, in order to provide a basis for comparison.¹⁵ Just as before, I announced each auction via three posts to the relevant newsgroup, as well as via email to a list of bidders.¹⁶

In the first four auctions, I auctioned 99 different cards each time, setting a reserve price for each card as a particular fraction of the current Cloister price of that card. In each of the first two auctions, nine cards were auctioned at a minimum bid of 10 percent of the Cloister price, nine at 20 percent, nine at 30 percent, and so on, up to a maximum of 110 percent of the Cloister price. For each reserve-price level, I chose an assortment of different cards with widely different Cloister prices, and scattered the group randomly across the complete list of cards. After an analysis of the data from those auctions, I chose to collect more data both at very low and at very high reserve price levels. Therefore, the third and fourth auctions were designed to have equal numbers of cards auctioned at reserve levels of 10, 20, 30, 40, 50, 100, 110, 120, 130, 140, and 150 percent of the Cloister price.¹⁷

¹⁵ It was necessary to do another absolute auction, rather than just reusing those of the previous section, because those took place in a substantially different time period, with a very different number of invited bidders, thus making their results incomparable to those of the within-card experiments.

¹⁶ For this series of auctions, the bidder pool was quite a bit larger than before. 531 individuals were emailed to participate in Auction R1, and as some people specifically requested to be removed from my auction announcement mailing list, the list dropped to 489 individuals by the time Auction R4 began.

¹⁷ In practice, the number of cards at each reserve-price level ended up not being precisely equal. Because I required bids to be in multiples of \$0.05, I always took the computed reserve price and rounded it down to the nearest acceptable bid amount. In the analysis below, I take the ratio of the actual reserve price used to the

This variation in reserve price levels was designed to investigate how both bidder behavior and expected auction revenue would react to changes in the reserve price, and to calculate the optimal reserve price level. Normalizing by the Cloister price, since this is a standard reference price computed in the same way for all *Magic* cards, makes cross-card comparisons feasible. Besides the exceptions noted above, all experimental protocols and bidder instructions were kept identical to those used in the auctions with reserve prices in the experimental design described in section 4.1.

Summary statistics for the between-card auctions are given in Table 3. In auctions R1 to R4, reserve prices ranged from 0% to 150% of each individual card's Cloister value, and the average reserve price level varied slightly from auction to auction, from 60% to 85%. In auction R0, of course, the average reserve price level was zero.

As can be seen in the table, each auction had dozens of bidders and hundreds of bids on individual cards. The number of people receiving email invitations to participate declined with each successive auction, but only due to recipients asking to be removed from my mailing list, so the changes in the mailing list should not have affected the number of potential participants. Note that the data from the between-card auctions is not directly comparable to that from the within-card auctions, because the size and composition of the pool of participating bidders changed considerably during the intervening six months. Very few bidders overlapped between the two experiments; most of the bidders in the between-card experiment were new recruits.

The table also displays aggregate statistics on revenue, including the total Cloister value of all the cards in each auction, the total revenue earned on cards which were sold, and a grand-total revenue figure which also includes the salvage value of the unsold cards. The auction

Cloister price, and round to the nearest 10% level in order to examine the effects of the reserve price on expected revenue. This results in unequal numbers of cards at the different levels of reserve prices.

revenue in each case was reasonably close to the total Cloister value of the cards; in Auction R2 I earned revenue greater than the total Cloister value, while in the three others I earned slightly less.

5 Results

I now present the results from the experiment, separately for each of the three empirical questions outlined above. Are entry costs relevant? Is the entry equilibrium stochastic or deterministic? Do the auctioneer's profits improve as he raises the reserve price to be at least as high as his salvage value?

5.1 Entry costs are relevant

The within-card experiments demonstrate that endogenous bidder entry appears to be the right model for this market. Statistics on the number of card bids per participating bidder are shown in Table 2. As expected, individual bidders tend to submit fewer bids in the presence of minimums than they do in the absence of minimums. This does not in itself demonstrate the existence of bidding costs; a bidder who contemplates how much to bid and then decides that the reserve price exceeds his maximum willingness could still be counted as having “participated,” because the decision cost would already have been incurred even though the reserve price prevents me from observing a low bid. In the auctions with minimums, no single bidder submitted bids on even half of the cards; the maximum number of bids by a single bidder was 30. By contrast, there were bidders in both of the no-minimum auctions who submitted individual bids on every single card.

Interestingly, relatively few bidders followed this strategy of bidding on every single card in the absolute (no-minimum) auctions. Only one out of 19 bidders bid on every single item in Auction AA, and only six of 62 bidders bid on every single item in Auction BA. These statistics indicate that the cost of submitting a bid (the participation cost) is high enough to affect bidder behavior, and thus this experimental environment is appropriate for exploring endogenous-entry bidding models such as MQV. If there were no cost to submitting a bid, then one would expect to see all of the participating bidders submitting bids on every card (as low as a nickel, say), since every card does have some positive resale value even to people who get no consumption utility from it.¹⁸ I conclude

¹⁸ Because of the time and transaction costs involved in selling it off, it is conceivable that for some bidders, the net resale value of a card might be less than five cents. However, most cards had gross resale values of over a dollar, and many bidders in this market could be assumed to take some pleasure in trading cards with others, as trading is a big part of the game culture.

that bidders deem the probability of getting a bargain (and thus a resale profit) on such a card is low enough that the expected profit from bidding does not always outweigh the cost of having to decide on a bid amount and to type the approximately ten characters required to submit another card bid. Indeed, the median number of card bids submitted by a single bidder was only 13 (of a possible 87) in Auction AA, and 14 (of a possible 78) in Auction BA, even though these auctions had no minimum bids.

Thus, bidders do appear to make a participation decision consistent with the existence of small entry costs; the number of participating bidders in each auction is not exogenous. The classical theory makes some accurate predictions about the effects of reserve prices, as shown earlier, despite this violation of its assumptions.

5.2 Is the entry equilibrium stochastic or deterministic?

Given the existence of endogenous entry, I now ask: is the entry equilibrium deterministic or stochastic? Very few bidders bid on a card both times it was offered, despite the fact that the same people were invited each time. Nineteen and seven bidders, respectively, bid in the two Antiquities auctions, but only 4 people overlapped between the two auctions. In the Arabian Nights auctions, there were 42 and 62 bidders, but only 17 of the bidders overlapped between the two. Thus, in each pair of auctions, there were a proportionally large number of people who entered the first auction but not the second, and other people who entered the second auction but not the first. This argues in favor of a stochastic equilibrium, as the most natural kind of deterministic equilibrium is one in which the same bidders enter each time.

Two objections might be raised to the result just presented. First, it might be the case that people enter one auction but not the other because the latter auction has reserve

prices which are higher than they are willing to pay. However, this screening-out explanation cannot account for the bidders who bid in the presence of reserve prices but fail to bid in the absence of reserve prices; there were 3 such bidders in the Antiquities auctions, and 25 such bidders in the Arabian Nights auctions. The second potential objection is that bidders may have bid in the chronologically first auction, but not the second, in a pair because they had already bought the cards by the time the second auction occurred. This complaint potentially affects the 25 Arabian Nights bidders just cited, who bid in Auction BR but not in Auction BA. Indeed, three of these bidders each placed a bid on a single card in Auction BR and won it, so there would be no reason to expect them to bid in the second auction. However, none of the remaining 22 bidders won all the cards they bid on in Auction BR: ten did not win any cards at all, while the remaining twelve won an average of 50 percent of the cards they bid on. It is still possible that these bidders managed to purchase the rest of the cards they were interested in from someone else during the week that passed between my two auctions, but I can at least say that they did not buy them from me. Thus, the evidence is fairly strong that bidders in these auctions made stochastic entry decisions: faced with the same auction opportunity, the same person might sometimes enter and sometimes fail to enter. This contrasts with the laboratory findings of Palfrey and Pevnitskaya (2003), who find .

The stochastic entry decision might not be due to conscious randomization by an bidder trying to follow a “mixed strategy” in the textbook sense. Perhaps bidders enter “randomly” because of other things happening in their lives: a college student had too much homework one week, or a computer programmer had a family emergency. Lots of random events could cause bidders to show up to one auction but not another. However,

in terms of auction design and welfare considerations, what matters is whether the entry decisions in a real-world auction are deterministically predictable by the auctioneer and by the rival bidders. My evidence shows that at least in this market, bidder entry decisions are stochastic, so the model of Levin and Smith (1994) has empirical relevance.

5.3 Optimal Reserve Price with Endogenous Entry

Recall that the main prediction of the MQV paper is that raising the reserve price from some low value to the salvage value of the good will increase expected auction profits, even in an endogenous-entry context. In order to understand the effect of the reserve price on expected revenues, I turn to the between-card experimental data. Recall that these data provide samples of auction revenues at differing reserve price levels (normalized by Cloister price for each card).

Table 4 summarizes the results of the experiment separately for each reserve-price decile, from reserve prices of 0% of the Cloister price to reserve prices of 150% of the Cloister price. The table displays the total number of cards I auctioned at each reserve price, the number of those which went unsold, and the mean and standard deviation of the revenues at each reserve price level. The revenues are also normalized by the Cloister price of each card, and an unsold card counts as an observation of zero revenue. The data are displayed graphically in Figure 1, with the mean revenues plotted against the reserve prices. The error bars show one standard error in each direction (where the standard error equals the standard deviation in revenues for that reserve price level divided by the square root of the number of observations at that reserve price level). We see that the revenues are quite high at a reserve price of zero, then drop off sharply at a reserve price of 10% of Cloister price. Revenues seem to rise again, generally, between 50% and

100% of the Cloister price, then fall again at higher reserve price levels. There are surprisingly high revenues observed at 140% to 150% of the Cloister price, albeit with high standard errors.

To test the MQV prediction also requires an estimate of the salvage value for the unsold cards. I asked my local card dealer what he would pay me for my unsold cards; he responded with an offer that was 20 percent of their Cloister price. He further indicated that 20% of Cloister price would be his average offer price for cards of this quality and quantity, so I adopt a salvage value of 20% percent of the Cloister price for each card.¹⁹

Now the question is whether a reserve price equal to the salvage value yields expected profits at least as high as a reserve price less than the salvage value (0% or 10%) of salvage value. The point estimates of revenues certainly indicate that the opposite is the case. In order to perform a formal hypothesis test, first I calculate expected profits rather than expected revenues. For the 0% reserve price, all cards sold, so profits remain the same as revenues: 1.192. For the 20% reserve price, two cards went unsold; when I count salvage profits of 20% of Cloister price for each of these cards, the estimate of expected profit rises slightly, from 0.857 to 0.870. Using the calculated standard deviations, I conduct a test of the null hypothesis of equality between expected profits at 0% reserve price and expected profits at 20% reserve price. The resulting standard normal test statistic is 2.18, with a p-value of 0.029. Thus, I reject the null hypothesis of equality at the 5% level of significance, and conclude that expected profits are actually *higher* for a zero reserve price than they are for a reserve price equal to the

¹⁹ I might have been able to shop around for a better price with a different card dealer, but this represents my best estimate of a salvage value, which by definition should be net of all administrative costs, including search costs.

salvage value.²⁰

This is a violation of the theoretical prediction, an example of a case in which the auctioneer does better to hold an auction with a zero reserve price than to set the reserve price equal to the salvage value. One possible explanation is that an auction with no reserve price generates more enthusiasm among bidders, causing higher levels of participation. In other words, modest minimum bids may eliminate some high valuation-bidders, who would have bid high if they had participated, but decide not to participate unless their attention is attracted by an auction with zero minimum bids. Although a few items may end up being sold at very low prices, they might serve as “loss leaders,” similar to the goods advertised at deep discounts by supermarkets, enabling the auctioneer to collect higher revenues overall. This proposed effect involves increased entry through attracting bidders’ *attention*, with the absolute auction as a type of promotion, rather than assuming the bidders will make a careful calculation of the costs versus the benefits of bidding. Note in Table 4 that the total number of bidders in Auction R0 is actually lower than in the other auctions, which might seem to be evidence against this effect, although I should also note that the number of cards in auction R0 is also lower than in the auctions with reserve prices. One caveat about this finding is that most of the zero-reserve-price cards were sold in the same auction (R0). Although I did attempt to keep all other variables constant across auctions, the anomaly might be due to some uncontrolled factor which was different between R0 and the earlier auctions.

²⁰ The reader might wonder about robustness to alternative assumptions about the salvage value. In particular is possible that I may have overstated auction profits, because my revenue figures are not discounted for the labor and postage I spent in order to ship the cards to the winning bidders, and therefore non-auction salvage might actually be more attractive, relative to auction revenues, than I initially assumed. Assuming salvage values of 30%, 40%, or even 50% of Cloister price still yields statistically significantly higher profits for a zero reserve price than for a reserve price equal to salvage value, so the result is quite robust. (The difference is no longer statistically significant for assumed salvage values of 60% or higher, as there are fewer observations at these higher reserve price levels.)

6 Conclusions

This study presents the results of controlled experimental auctions performed in a field environment. By auctioning real goods in a preexisting, natural auction market, I have obtained data in a manner that is intermediate between laboratory experiments and traditional studies of field data. Some variables were unfortunately unobservable and uncontrolled—for example, I could not assign “valuations” for each good to each bidder, as a laboratory experimentalist might. On the other hand, I have the opportunity to hold constant most of the relevant variables in the environment, and to manipulate the treatment variable, which in this case was the existence and level of reserve prices. By giving up the ability to observe and manipulate some of variables that laboratory experimenters can control, I gained a realistic environment. The participants had previous experience bidding for the types of real goods I auctioned, and the auctions took place in an Internet-based market where bidder entry decisions seemed potentially important.

The first result is that entry costs are an important feature of this real-world auction markets, thus confirming the central assumption of endogenous-entry auction theory. The costs in the *Magic*-card market are probably not nearly as dramatic as those postulated in other markets (for example, in the market for offshore oil rights the bidders typically hire geologists to perform extensive analysis of the potential for oil in a given tract). Here, the cost of acquiring information about individual cards is quite small, but even the cost of typing in a bid amount appears to have observable effects.

Second, when the same cards were auctioned twice in rapid succession, very different sets of people decided to submit bids, despite the fact that the same superset of people

were invited to participate both times. This can be interpreted as evidence in favor of the stochastic (mixed strategy) entry equilibrium model, where the number of participating bidders varies unpredictably.

Third, I found that, contrary to the theory of McAfee, Quan, and Vincent (1998), a zero reserve price can earn higher expected profits than a reserve price equal to the auctioneer's salvage value. Perhaps an absolute auction attracts significantly more bidder attention than an auction with even modest reserve prices, causing more additional entry than might be suggested by a model of rationally calculated bidder entry decisions. It will be interesting to see whether this finding can be replicated in other auction markets.

References

- Bajari, Patrick and Ali Hortacsu. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, Summer 2003, vol. 34, no. 2, pp. 329-355.
- Black, Jason. *Cloister's Magic Card Price List*, <<http://www.hhhh.org/cloister/pricelists/>>, various weekly issues.
- Cox, James C., Samuel H. Dinkin, and James T. Swarthout. "Endogenous Entry and Exit in Common Value Auctions," *Experimental Economics*, October 2001, vol. 4, no. 2, pp. 163-181.
- Cox, James C., Bruce Roberson, and Vernon L. Smith. "Theory and Behavior of Single Object Auctions," in *Research in Experimental Economics*, Vernon L. Smith, ed., Greenwich, Conn.: JAI Press, 1982.
- Engelbrecht-Wiggans, Richard. "On Optimal Reservation Prices in Auctions," *Management Science*, June 1987, vol. 33, no. 6, pp.763-770.
- Engelbrecht-Wiggans, Richard. "Optimal Auctions Revisited," *Games and Economic Behavior*, 1992, vol. 5, pp. 227-239.
- Hendricks, Kenneth, and Harry J. Paarsch. "A Survey of Recent Empirical Work Concerning Auctions," *Canadian Journal of Economics*, 1995, vol. 28, no. 2, pp. 315-338.
- Kagel, John H. "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, J. Kagel and A. Roth, eds. Princeton: Princeton University Press, 1995, pp.501- 585.
- Levin, Dan, and James L. Smith. "Equilibrium in Auctions with Entry," *American Economic Review*, 1994, vol. 84, no. 3, pp. 585-599.
- Levin, Dan, and James L. Smith. "Optimal Reservation Prices in Auctions," *Economic Journal*, 1996, vol. 106, pp. 1271-1282.
- Lucking-Reiley, David. "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, 1999, vol. 89, no. 5, pp.1063-1080.
- Palfrey, Thomas R., and Svetlana Pevnitskaya, "Endogenous Entry and Self-Selection in Private-Value Auctions: An Experimental Study," Caltech Social Science working paper #1172, August 2003.
- Pevnitskaya, Svetlana, "Endogenous Entry in First-Price Private-Value Auctions: The Selection Effect," Ohio State University working paper, January 2004.
- Reiley, David H. "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet," University of Arizona working paper, August 2004.
- McAfee, R. Preston, and John McMillan. "Auctions and Bidding," *Journal of Economic Literature*, 1987a, vol. 25, no. 2, pp. 699-738.
- McAfee, R. Preston, and John McMillan. "Auctions with Entry." *Economics Letters*, 1987b, vol. 23, no. 4, pp. 343-347.
- McAfee, R. Preston, Daniel Quan, and Daniel Vincent. "How to Set Optimal Minimum Bids, with an Application to Real Estate Auctions," *Journal of Industrial Economics*, 2002, vol. 50, no. 4, pp. 391-416.
- McMillan, John, "Selling Spectrum Rights," *Journal of Economic Perspectives*, 1994, vol. 8, no. 3, pp. 145-162.
- Milgrom, Paul R., and Robert J. Weber. "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, vol. 50, pp. 1089-1122.
- Riley, John G., and William Samuelson, "Optimal Auctions," *American Economic*

- Review*, 1981, vol. 71, no. 3, pp. 381-392.
- Samuelson, William F. "Competitive Bidding with Entry Costs," *Economics Letters*, 1985, vol. 17, pp. 53-57.
- Smith, James L, and Dan Levin, "Entry Coordination, Market Thickness, and Social Welfare," *International Journal of Game Theory*, 2001, vol. 30, pp. 321-350.
- Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 1961, vol. 16, no.1, pp. 8-37.
- Wilson, Robert. "Strategic Analysis of Auctions," in *The Handbook of Game Theory*, R.J. Aumann and S. Hart, eds. New York: North-Holland, 1992, pp. 227-279.

Figure 1: Mean Revenue as a Function of Reserve Price

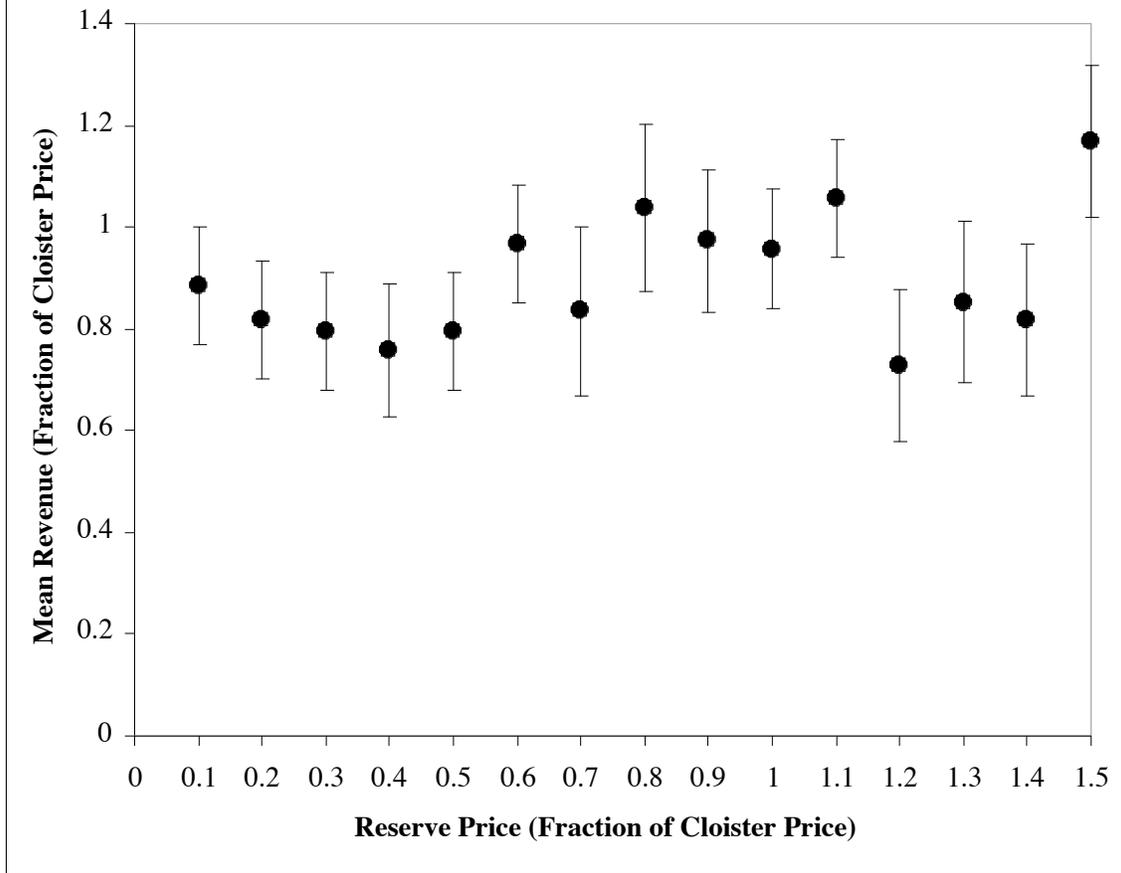


Table 1: Summary statistics for within-card experiments.

	Auction AA	Auction AR	Auction BA	Auction BR
Minimum bids?	No	Yes	No	Yes
Card set	Antiquities	Antiquities	Arabian Nights	Arabian Nights
Start date	Fri, 24 Feb	Fri, 10 Mar	Tue, 14 Mar	Tue, 28 Feb
End date	Fri, 3 Mar	Fri, 17 Mar	Tue, 21 Mar	Tue, 7 Mar
Number of items for auction	86	86	78	78
Number of items sold	86	47	78	74
Revenue from twice-sold cards	\$189.90	\$234.75	\$758.25	\$783.80
Total auction revenue	\$292.40	\$234.75	\$774.75	\$783.80
Total number of bids	565	71	1583	238
Total number of bidders	19	7	62	42
from email invitations	12	5	46	35
from newsgroup announcements	7	2	18	7
Number of email invitations sent	52	50	232	234
Number of winners	15	6	25	27
Winner/bidder ratio	78.9%	85.7%	40.3%	64.3%
Cards per winner:				
Max	25	26	12	18
as share of total	29.1%	55.3%	15.4%	24.3%
Min	1	1	1	1
Mean	5.7	7.8	3.1	2.7
Median	3	3.5	2	2
Payment per winner:				
Max	\$70.00	\$129.40	\$316.50	\$128.00
as share of total	23.9%	55.1%	40.9%	16.3%
Min	\$3.00	\$0.70	\$1.05	\$2.55
Mean	\$19.49	\$39.13	\$30.99	\$29.03
Median	\$10.50	\$23.68	\$13.15	\$13.00

Table 2: Bids received in the within-card auctions.

	Auction AA	Auction AR	Auction BA	Auction BR
Minimum bids?	No	Yes	No	Yes
Card set	Antiquities	Antiquities	Arabian Nights	Arabian Nights
Number of bidders	19	7	62	42
Number of items for auction	86	86	78	78
Number of bids per bidder:				
Mean	29.7	10.1	25.5	5.7
Median	13.0	4.0	14.0	4.0
Max	86.0	29.0	78.0	30.0
Min	1.0	1.0	1.0	1.0

Table 3: Summary statistics for the between-card experiments.

	Auction R1	Auction R2	Auction R3	Auction R4	Auction R0
Card set	Artifacts	Black	White	Blue	Red/Green
Start date	Tue, 3 Oct	Fri, 6 Oct	Fri, 20 Oct	Mon, 23 Oct	Tue, 31 Oct
End date	Tue, 10 Oct	Fri, 13 Oct	Fri, 27 Oct	Mon, 30 Oct	Tue, 6 Nov
Number of items for auction	99	99	99	99	86
Number of items sold	98	92	77	78	86
Mean reserve level	60%	60%	85%	81%	0%
Total number of bids	798	652	366	401	1069
Total number of bidders	57	55	46	38	42
Number of email invitations sent	532	523	512	489	472
Total Cloister value	345.83	271.55	285.87	224.89	327.05
Total auction revenue	338.45	282.65	260.95	219.25	316.70
Revenue plus salvage	343.94	283.65	269.48	224.52	316.70

Table 4: Cards and revenues at each reserve price.
(Reserve prices and revenues normalized by the Cloister price of each card).

Reserve price	Total cards	Unsold cards	Mean revenue	Std dev of revenue
0.0	96	0	1.192	1.071
0.1	33	0	0.847	0.549
0.2	36	2	0.857	0.594
0.3	34	0	0.823	0.599
0.4	27	2	0.775	0.500
0.5	32	2	0.945	0.517
0.6	20	1	0.977	0.480
0.7	16	0	0.965	0.259
0.8	23	1	1.093	0.469
0.9	31	4	0.983	0.500
1.0	35	6	1.055	0.674
1.1	32	4	1.113	0.491
1.2	15	7	0.804	0.861
1.3	21	10	0.760	0.772
1.4	17	6	0.967	0.867
1.5	14	5	1.104	0.893