

This work is distributed as a Discussion Paper by the
STANFORD INSTITUTE FOR ECONOMIC POLICY RESEARCH

SIEPR Discussion Paper No. 99-23

**Winner Curse, Reserve Prices and Endogenous
Entry: Empirical Insights from eBay Auctions**
by

Patrick Bajari
Ali Hortacsu

Stanford University

March 2000

Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
(650) 725-1874

The Stanford Institute for Economic Policy Research at Stanford University supports research bearing on economic and public policy issues. The SIEPR Discussion Paper Series reports on research and policy analysis conducted by researchers affiliated with the Institute. Working papers in this series reflect the views of the authors and not necessarily those of the Stanford Institute for Economic Policy Research or Stanford University.

Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions

Patrick Bajari and Ali Hortacısu*

Department of Economics
Stanford University

March 27, 2000

Abstract

Internet auctions have recently gained widespread popularity and are one of the most successful forms of electronic commerce. We examine a dataset of eBay coin auctions to explore features of online bidding and selling behavior. We find that profit margins for bidders appear to be slim on average, that previous bidding experience does not play a significant role in observed bidding behavior and that bidding activity is concentrated at the end of the auction. We then develop new econometric techniques to estimate a structural model of bidding to address three main issues. First, we measure the extent of the winner's curse. We find that for a representative auction in our sample, a bidder's expected profits fall by 3.2 percent when the expected number of bidders increases by one. Second, we document that costly entry is a key component in understanding observed bidding behavior. For a representative auction in our sample, a bidder requires \$3.20 of expected profit to enter the auction. Third, we study the seller's choice of reserve prices. We find that items with higher book value tend to be sold using a secret reserve price with a low minimum bid. We find that this is, to a first approximation, consistent with maximizing behavior.

1 Introduction

Auctions have found their way into millions of homes with the recent proliferation of auction sites on the Internet. The opportunity to study such online markets, especially the Web behemoth eBay,

*E-mail: bajari@stanford.edu, hortacsu@stanford.edu. We would like to thank Susan Athey, Philip Haile, David Lucking-Reiley, and participants at the Winter 2000 NBER Industrial Organization Conference, Caltech, University of Iowa, USC, Stanford, Federal Reserve Bank of Minneapolis and SITE seminars for their valuable comments. Several of the subroutines we employed in this paper can be downloaded from <http://www.stanford.edu/~bajari/wp/auction/code/code.html>

is doubtlessly a golden one for economists.¹ eBay is a source for a large amount of high quality data and serves as a natural testing ground for existing theories of bidding and for market design issues in auctions.

Auctions are clearly one of the leading innovations of Internet commerce. *The Economist* reports that on a single day, eBay boasts some 2.4 million items for sale in over 1,500 unique categories ranging from beanie babies, to art, to used computers. On the surface, online auction markets, such as eBay, appear to be a close approximation to an economist's idealization of a frictionless, competitive market where a large number of buyers and sellers engage in trade. The eBay community alone boasts some 3.8 million registered members.²

Online auctions and electronic commerce more generally have informational asymmetries that make these market arrangements distinct from shopping at "brick and mortar" stores. First, it is not possible for the buyer to directly inspect the good she is purchasing. On eBay, where primarily used goods are bought and sold, the inability of buyers to directly inspect a good might inhibit her ability to exactly what product she is buying and its quality. In some, but not all auctions, a picture of the object being bought and sold is displayed on the web page listing the auction and a brief description is presented. Bidders also can e-mail the seller to ask questions about the object being sold. In "brick and mortar" stores, on the other hand, the buyer has considerably more freedom to inspect the item being sold. Second, many of the sellers are individuals with a limited track record. Bidder's have to worry whether the seller has provided accurate information about the object up for auction and will send the object undamaged if the bidder wins the auction. The reliability of "brick and mortar" stores that is located near the home of the buyer may be subject to considerably less uncertainty.

In online auctions there is clearly a possibility of a winner's curse- bidders only win items when they pay too much. The winner's curse is more severe as the number of bidders becomes larger and, in a community of 3.8 million bidders, might be a factor that inhibits trade. Recently, leading theorists have suggested that information asymmetries and the winner's curse may inhibit efficient trade in online markets.³

One aim of this paper is to quantify the extent of the "winner's curse" on eBay. To do this, we chose the market for U.S. mint/proof coin sets. These coin sets are regarded as investment-grade collectibles. The winner's curse may be a factor in bidding on these coins for two reasons. First, coin collectors care deeply about the appearance of the coin that they are buying. Coins that are scratched and that lack brilliance have diminished worth. Second, many of these coins are traded in liquid resale markets and therefore assessing the value of the coin for resale may introduce common value elements.

Another striking feature of auctions on eBay is that sellers are allowed some freedom in choosing

¹For a broad survey article covering 142 online auction sites, see Reiley (1999)

²See *The Economist*, July 22, 1999, "The Heyday of the Auction."

³"The Heyday of the Auction"

the rules of the auction. Economic theorists have spent a great deal of time analyzing incentives and equilibrium under alternative auction rules. One form of rules we will focus on quite carefully is the choice of reserve price policy.

In an auction, a reserve price is the price below which the auctioneer refuses to sell the item. On eBay, we study two alternative forms of reserve prices. This first is that the seller can set a minimum bid for the auction, which is equivalent to what theorists typically call a posted reserve. Second, for an additional fee, sellers can use a hidden reserve price so that bidders are aware that a reserve is present, but do not know the value of the reserve at the time of bidding.

A large number of papers in economic theory, starting with Myerson (1981) and Milgrom and Weber (1982), have studied equilibrium in auctions with reserve prices and characterized the optimal reserve price policy. It can be shown that in the case of independent private values, the revenue maximizing mechanism involves a non-zero reserve price. In many real world auctions, including procurement, treasury bills, real estate and auction houses, both posted and hidden reserve prices are commonly used.

The large volume of bidding data present on eBay allows us to study the sellers choice of reserve price policy. We find in the data, somewhat surprisingly, that high value items tend to use a secret reserve price with a low minimum bid while items with a low book value use only a posted reserve. To our knowledge, there is no existing theoretical model that has this prediction.

We believe the solution to the above puzzle relies on the fact that bidders *choose* which auction to attend. Our data set indicates quite clearly that setting a high minimum bid (public reserve price) deters bidders from participating in the auction. Studying the relationship between reserve prices and entry, however, comes at a cost: we have to build a realistic model of bidding with endogenous entry that we can take to the data.

The theoretical motivation for incorporating endogenous entry into models of bidding have previously been pointed out by Harstad (1990), and Levin and Smith (1994): Most celebrated results in auction theory, including the revenue ranking/equivalence theorems of Vickrey (1961), Milgrom and Weber (1982), and the optimal mechanism design literature assume that the number of bidders is fixed across the mechanisms they compare. However, in a setting like eBay, auctions of similar or even identical objects are listed side-by-side with differing bidding rules. Therefore if one set of auction rules promises higher expected profit to bidders than another, we would expect more bidders to enter the high-expected profit auction than the other, to the point that bidder rents are dissipated to a point of indifference. If there are enough potential bidders (millions on eBay), this point of indifference will be the cost of entry, which might be the cost of determining a sensible bid and the opportunity cost of time spent browsing the Internet. This means that not all bidder surplus is appropriated by the seller in terms of bidding proceeds, and hence the seller's revenue might depend on the amount of entry induced by the mechanism.

Empirical literature on auctions also takes the number of bidders at an auction as exogenously

given. When, in the spirit of Paarsch (1992), Laffont, Ossard and Vuong (1995), or Elyakime, Laffont, Loisel and Vuong (1994), we estimate the distribution of valuations of a set of bidders from their observed bids, we implicitly condition on the fact that these bidders have rationally chosen to enter the auction while others from a larger population of potential bidders have not. And when we make policy recommendations based on these estimates, we typically assume that the same set of bidders will participate in the modified mechanism. If, however, modifying the mechanism affects the entry decisions, this will no longer be true.

Estimating a model with common values and endogenous entry poses several technical difficulties. First, no previous estimation procedures can clearly be applied to the modeling framework we wish to study. One recent approach to estimation, suggested by Elyakime et al. (1994), requires that in the first order conditions, bidders private information can be represented as a function of the observed bid and the empirical probability distribution of bids. In our model, this condition fails because the valuation of the highest bidder cannot be recovered from the data when only the second highest price is observed. Also, when common values are present, the first order condition can be used to identify a function of the underlying model primitives, as has been noticed for instance by Hendricks, Pinske and Porter (1999) in first price auctions. However, in many cases additional steps are needed to recover the structural parameters of the model.

Another recent estimation procedure, suggested by Laffont et al. (1995) uses the fact that the expected revenue in a first price auction and second price auction are the same to formulate a method of moments estimation procedure. Laffont et al. (1995) argue that their procedure is desirable since it avoids the difficulty of computing equilibrium bid functions. In the case of eBay, however, the result that expected revenue is identical under alternative auction formats will fail because of the presence of common values.

The approach to estimation we take is most closely related to the work of Paarsch (1992) and Bajari (1997). An important step in the estimation procedure is to calculate the equilibrium bid functions efficiently. With common values, this may involve solving high dimensional integration problems that are computationally burdensome. An innovation of this paper is to realize, that in many cases, equilibrium bid functions are “linearly scalable.” That is if we make a linear transformation of the underlying valuations, a linear transformation of the bid function will be an equilibrium to the new game. This insight makes structural estimation of common value models computationally feasible in some settings. It is worth noting, that this insight can be employed in structural estimation of many other auction games, such as first price sealed bid auctions with private or common values, where computing the bid functions can be very expensive computationally.

Another problem for estimation of our model, as was first recognized by Donald and Paarsch (1993) is that the asymptotics for maximum likelihood estimation will not be straightforward because the support of the likelihood function depends on the parameter values. This is a generic problem in auction models. This paper will utilize Bayesian methods proposed first by Bajari (1997) to

overcome this difficulty. While the problem we study in this paper is bidding in eBay auctions, the methods we propose are useful for estimation of other models that are difficult to compute.

The auction model we estimate generalizes previous estimation strategies in two ways. First, we are able to study a more general information structure in common value auctions. In Paarsch (1992), for example, the author is not able to recover the distribution of the unobserved common value component. This is because of the high dimensional integrals involved in computing the equilibrium. Our approach allows for the recovery of this common value component. Second, we estimate the unobserved entry process and its relation to the choice of minimum bid and reserve price policy.

The outline of the paper is as follows: we first give an overview of the market we're studying and present our data. We document several reduced form findings: bidding activity seems to be concentrated at the end of the auction, there is significant variation in the number of bidders who attend the auction, reduced form bid functions are decreasing in the number of bidders and reserve prices seem to have a strong impact on the number of bidders who participate in the auction. Also, sellers use secret reserves with low minimums for "big ticket" items and posted reserves for more mundane objects. We also note that reserve prices, as well as the number of bidders are likely to be endogenous and that compelling instruments do not exist. Therefore, we argue that a structural approach will be required to answer certain questions.

Second, after examining the reduced form evidence, we argue that a common values model with endogenous entry is the best description of the market that is computationally tractable. We derive the equilibrium to the model and discuss its properties. Third, we specify a structural model and describe how to estimate the model parameters.

Next, we summarize the parameter estimates from the structural model and discuss their economic interpretation. We then use our structural model to study a set of policy questions. First we measure the extent of the winner's curse in the market. This allows us to assess the claims that have been made by theorists that informational asymmetries may inhibit efficient trade in online auctions. Then we document that it is costly for bidders to enter the auction and measure the implicit costs of entry in this model. Finally, we simulate the seller's expected profits under alternative reserve price policies. We find that, to a first approximation, the observed behavior of the sellers appears to be optimal.

2 Description of the market and data

eBay is the premier auction site on the Internet. Since it was launched in 1995, it has captured a good share of the second-hand and collectible goods trade, and by the meteoric rise in its share price, has become one of the most successful and popular Internet companies. Furthermore, with thousands of browsable auctions closing each day, it is a goldmine for empirical research on auctions.

However, since most items on eBay are second-hand goods or collectibles, they are very difficult to appraise by a non-specialist economist, and hence empirical testing of results from auction theory run into the difficulty of not being able to account for item-specific differences. Therefore, in this research, we focus on a market segment where a little reading about the market and consulting published price lists enable us to control for item-specific differences. This is the collectible coins market, which we describe briefly in the next section.

2.1 The collectible coins market

Coin collecting is a hobby with a long history and tradition. Aside from very rare coins from antiquity that can command hundreds of thousands or millions of dollars, even relatively modern coins find a market among the millions of coin collectors in the U.S. Recently, the trade in collectible coins has become an investment channel for many, especially after coins began to fetch astronomic prices in the 80s. Although these prices have come down in the last decade, many coin traders can be classified as speculators or at least collectors with strong investment motives. As such, the coin market is an interesting cross between durable goods markets and securities markets, which, from an auction theorists' perspective, boils down to a cross between private values and common values.

The traditional channels for trade in coins are retail collectors' shops or mail-order outlets. Many dealers and experienced collectors also attend regional coin shows or private auctions. On-line auction sites such as eBay are a promising new venue for coins trade; however, they still capture only a small part of the coin trade.

For this study, we focus on collectible mint/proof sets. Mint sets have been prepared by the Treasury for sale to collectors since 1947, and contain uncirculated specimens of each year's coins for every denomination issued from each mint. Proof sets are manufactured to have sharper details and a more-than-ordinary brilliance than mint sets. They also contain one of each denomination of the coins minted that year, and are once again sold directly from the government to collectors. These sets have been sold since 1936, and usually command a higher price than mint sets.

Since both mint and proof coins are uncirculated, the proper grading of the coins becomes less of an issue⁴, making them "well-defined products" for which a non-coin collecting econometrician can find price lists, unlike some specialty coins that would require an expert to appraise. A drawback to focusing on mint/proof sets is that post-1960 issues of these sets are fairly mundanely priced (some as low as \$5-\$10), so the assumption of "bidder sophistication" may become somewhat suspect.

⁴According to the Official Red Book, since 1959, mint sets have been sealed in a protective plastic envelope. Proof sets also come in special plastic cases that protect them against tarnish. Some of the auction listings give additional information about the condition of the original casing, sometimes the original envelope or the casing is damaged or not available. But unless otherwise noted, the condition of the coins is taken to be "mint-state," the highest grade a coin can attain. We have also excluded reportedly self-assembled mint or proof sets, since these sets usually fetch much lower prices than original sets.

2.2 eBay auctions

Our data set contains the on-line auctions of U.S. mint/proof sets that were conducted between September 28 and October 2, 1998 at the eBay on-line auction site. In that period, about 100 auctions closed every day, though the volume has since increased to about 400 auctions per day. The auctions end 3, 5 or 7 days after they are listed. eBay only acts as an intermediary in these trades and does not handle the physical delivery of the goods. eBay's revenue comes from a fixed listing fee and a percentage of the transaction value if the auction is successful.

The bidding format used at eBay is a dynamic variant of the second-price sealed-bid auction, called "proxy bidding." Each bidder is asked to enter the maximum amount she is willing to pay. The eBay computer compares this value to the maximum bids submitted by other bidders. In case it is higher than other maximum bids, the new bidder is designated as the "leading bidder." However, instead of the value entered by the bidder, only the amount that's necessary to outbid the bidder with the next highest maximum bid is displayed as the "leading bid." This also becomes the amount the bidder ends up paying if there are no other challengers. If the value is submitted by a new bidder is not high enough to lead the auction, or if the current auction leader is outbid, he's notified by e-mail and may revise his bid. Once an auction has concluded, the winner is notified by e-mail and it is up to him to contact the seller to arrange shipment and payment details. Most of the auctions that we study have shipping and handling fees preset by the seller, which is to be paid on top of the bid amount.

A seller on eBay is constrained to use the proxy bidding method⁵, but has a few options to customize the sale mechanism. She can set a minimum bid, or a secret reserve price⁶. The minimum bid is observed by all bidders; however, in a secret reserve price auction, bidders are only told whether the reserve price is met or not. The reserve price is not announced even after the auction ends.

Aside from choosing whether to use a secret reserve price, the seller can also pay eBay to advertise her auction as a "featured auction" for \$99.95 or a "category-featured auction" for \$14.95. None of the auctions in our dataset were featured auctions, which is understandable since the average value of the coins were about \$25.

Sellers can also differentiate themselves through eBay's innovative buyer/seller feedback system. eBay allows the buyer and the seller to rate their counterparts in transactions. The rating is in the form of a positive, negative or neutral response. Next to each buyer or seller's ID (which is usually a pseudonym or an e-mail address), the number of net positive responses is displayed. In order to prevent planted responses, only unique responses are factored into the reputation rating.

In addition to data downloadable from eBay pages, we found the book values for the auctioned items from the November 1998 issue of *Coins* magazine, which collects value estimates from coin

⁵ "Dutch auctions" for the sale of multiple objects are also allowed, but we exclude such auctions from our study.

⁶ The seller can still set a minimum bid in a secret reserve price auction.

dealers and coin auctions around the U.S.⁷

For each auction we collected the following variables:

1. SECRET: 1 for a secret reserve price auction, 0 if not.
2. SALE: 1 if the auction was successful, 0 if not.
3. BOOKVAL: bookvalue of object according to Coins magazine.
4. SHIPHAND: Shipping and handling fee (if reported by seller).
5. BLEMISH: 1 if the seller reports that there's a blemish in the object.
6. OVERALL: Overall reputation of seller (number of unique positive responses - number of unique negative responses).
7. NEGATIVE: Number of unique negative responses to seller.
8. NBIDDER: Number of bidders.
9. AVGNBID is the average of bids in an auction as a percentage of bookvalue (conditional on there being at least one bidder).
10. DISPERSION is the difference between the maximum and the minimum bid as a percentage of bookvalue.
11. AVGFRAC is the average of bid submission times (as a fraction of auction duration).
12. AVGEXP is the average experience of bidders as measured by the number of previous transactions conducted on eBay.
13. MINBID: The minimum bid as set by the seller.
14. BID1N: The winning bid as a percentage of bookvalue (conditional on sale).
15. REVENUE: Revenue as a percentage of bookvalue (not conditional on sale).
16. BID1-BID25: Reported bids - which for every bidder except the winner corresponds to the submitted bid. BID1 is the amount paid by the winner, which is the bid increment plus the second bid. The bids are ranked in descending order.
17. FRAC1-FRAC25: Timing of final bid updates⁸ - as a fraction of the total auction duration.

⁷The November issue of the magazine was bought by one of the authors on October 22nd. We confirmed by e-mailing the magazine that the prices quoted were market prices of mid-October. This is a reasonable time frame for resale, taking into account that it usually takes a week for the buyer to get the coins from a seller on Ebay.

⁸Bidders have the ability to update their bids during the auction. What we have available in the "bid-history" pages on Ebay are the final bids by each buyer and the last time they updated their bids.

18. PREWIN1- PREWIN25: Previous wins of the bidder in eBay auctions - we use this variable as a proxy for bidder experience.

There were a total of 516 auctions completed in the mint-proof category in the five day sample we consider. We could not find reliable book values for 48 of the auctions and since the book value serves as our primary controlling variable for cross-object heterogeneity, we had to discard those observations from our data set (we checked whether this truncation makes any difference in the observed entry/reserve price setting behavior and could not find a noticeable pattern). 11 auctions were conducted as Dutch auctions and since those auctions are subject to a different set of rules, were omitted from our analysis. Detailed bid histories for 39 auctions were lost due to a technical error in data transfer. Aside from those, our dataset contains the entire set of mint-proof set auctions conducted during the dates considered.

We report summary statistics on the above variables in Table 1. We see that 85% of the auctions resulted in a sale, with an average of 3 bidders per auction. The average value of the traded coins is \$47, with a \$2 shipping and handling charge. The coins range from \$3 to \$3700 in value, reflecting the wide dispersion of prices even within this relative narrow section of the collectible coin market. 10% of the coin sets had reported blemishes.

Conditional on sale, revenues are 5% lower than the reported book values of the sample. However, not all items ended up being sold, so the expected average revenue from an auction, considering the possibility of no sale, was 81% of the bookvalue. The average bid in an auction was 83% of the bookvalue and the mean bid spread was 32% of the book value. There were quite a few outliers in the data: in 20 auctions, the winning bid was 50% above the bookvalue and in 5 auctions (all were posted reserve auctions), the winning bid was less than half the bookvalue.⁹

As for bidder and seller characteristics, we see that the average bidder has previously won 41 auctions and the average seller has done at least 203 transactions on eBay. Negative feedback points are very low compared to the positive feedback points of the sellers.¹⁰ We conclude from these observations that eBay is not just a marketplace for occasional hobbyists and that serious collectors and coin dealers are also participants.

3 Empirical regularities

In this section we report empirical regularities we observe in the data that we deem significant in formulating a model of bidding on eBay. We summarize the main findings of this section as the following:

1. We find that bidders update their bids frequently and that bidding activity is concentrated

⁹We checked these outliers to make sure there was not an error in coding the bookvalue. It's also interesting to note that 3 of the 5 low revenue auctions had more than one bidder.

¹⁰The maximum negative rating of 21 corresponds to the seller with the seller with 973 overall feedback points.

Table 1: Summary Statistics

Variable	Mean	Std.	Min	Max
SECRET	.1606	.3676	0	1
SALE	.8513	.3571	0	1
BOOKVAL	47.51	212.63	3	3700
SHIPHAND	2.18	.92	1	5
BLEMISH	.10	.32	0	1
OVERALL	203.06	208.43	0	973
NEGATIVE	.43	1.63	0	21
NBIDDER	3.05	2.46	0	14
AVGNBID	0.8354	0.2616	0.16	1.2824
DISPERSION	0.3215	0.2826	0	1.2333
AVGFRAC	.70	.25	0	1
AVGEXP	41.22	38.40	0	261.5
MINBID	16.28	28.59	.01	230
BID1N	0.9556	0.2820	0.1600	1.4767
REVENUE	0.8135	0.4284	0	1.4767
BID1	36.70	168.94	0	3010

at the end of the auction. We also find that bid levels decline in the number of bidders in an auction. We conclude that a private values framework can not explain this kind of behavior and that a common values framework, with a potential “winner’s curse” effect is better suited to model the auctions in our sample.

2. There is significant variation in the number of bidders in an auction. We find that a low minimum bid, a high bookvalue and low negative ratings increase entry to an auction.
3. Profit margins for bidders are slim and do not deviate significantly from book values of the coins once shipping and handling costs are accounted for.
4. The number of bids submitted in the auction is decreasing in the ratio of the minimum bid to the book value.
5. Sets with a high book value tend to be auctioned using a secret reserve price with a minimum bid that is significantly lower than the book value. Less expensive sets are auctioned using a posted reserve price with a higher minimum bid to book ratio.
6. Experience does not appear to play a significant role in bidding.

3.1 Private vs. common values and the winner’s curse

How do bidders on eBay value the coin sets they are bidding on? One answer to this question is to say that each bidder (indexed by i) gets utility v_i from winning the coin set. If the bidder i knows v_i (but not necessarily the v ’s of other bidders) we have the case of “private values.”

“Private values” may make sense if the bidder is buying the coin set as a part of a collection. However, the existence of an active secondary market in collectible coins makes it quite plausible for bidders to have a resale incentive. The frequent activity of individual bidders on eBay suggests that quite a few of them are dealers who make a living from buying and selling coins. If resale is the primary incentive, we can argue that all bidders get the same utility from winning the coin set; namely the resale value of the coin set. Here, $v_i = v$, the same number across all bidders. This is the case of (symmetric) “common values.”

Another source of common values besides resale comes from the fact that it is difficult for buyers to assess the condition of the coin that they are buying. Unlike shopping in “brick and mortar” stores, the buyer cannot directly inspect the coin they are buying online. Bidders must rely on information that is disclosed to them by the seller. The problem with the “common values” case is that, typically the value v is not known by the bidders. Rather than observing v , each bidder instead comes up with an “estimate” x_i , and bases her bid on this estimate. We can model this estimate as having the form $x_i = v + \varepsilon_i$, with ε_i as a noise term. Now let’s think whether a bidder would care to know about other bidders’ estimates. If a bidder believes that ε_i is truly an idiosyncratic noise term, then she can reduce the uncertainty of her estimate by averaging across all estimates. Through this averaging, a high estimate of the value of the object will be adjusted downwards, and a low estimate would be adjusted upwards.

Unfortunately, bidders rarely share their estimates. Here’s where the “winner’s curse” comes in: a bidder realizes that, conditional on winning, she must have had the highest estimate.¹¹ Therefore, with some probability, she will have paid too much for the item. Observe that there is no room for a “winner’s curse” in the private values setting. Each bidder is sure of the utility she will get from winning the auction, and does not have worry about whether she has understated or overstated this number¹².

The “winner’s curse” becomes worse as the number of bidders increases. The highest estimate among 15 bidders is very likely to be even higher than the highest estimate among 2 bidders.

¹¹To see why this is so, think of the case when each bidders estimate is unbiased. The unconditional mean of each estimate will be v , but the mean of the estimate *conditional* on the fact that it is the “maximum among n ” estimates will be greater than v .

¹²One could argue that even in a private values setting, bidders do not observe v_i until they receive the coin set and hence bid on the basis of an estimate $x_i = v_i + \varepsilon_i$. If the noise term ε_i is independent across bidders, then bidders have nothing to learn from other people’s estimates, and there is no room for a “winner’s curse.” If, however, the noise term is somehow correlated across bidders, then a “winner’s curse” still exists. Observe that this is observationally equivalent to a setting where the bidders valuation for the object has both a private and a common value component, where the private value component is u_i and the bidder makes an estimate, $\chi_i = v + \xi_i$ about the unknown common value v with noise ξ_i (i.e. $x_i = u_i + v + \xi_i$, with the correlation captured by v).

Table 2: Regression of bid/bookval on the number of bidders

Specification	Linear	Quadratic	Logarithmic
	$b = \beta_0 + \beta_1 N$	$b = \beta_0 + \beta_1 N + \beta_2 N^2$	$\ln(b) = \beta_0 + \beta_1 \ln N$
β_0	.9836 (.0202)	1.0932 (.0354)	-.0420 (.033)
β_1	-.0342 (.0035)	-.0804 (.0128)	-.0572 (.005)
β_2	N/A	.0038 (.0010)	N/A
R^2	0.0625	0.072	0.065

Therefore, we would expect bidders to be more afraid of a “winner’s curse in a 15 person auction than in a 2 person auction; hence we would expect lower bids in a 15 person auction than in a 2 bidder auction¹³.

As pointed out by Paarsch (1992), this last observation allows us to distinguish between a common values vs. a private values setting. Within the context of a second price auction or an ascending bid auction, such as the eBay auction, the number of bidders would have no effect on bids in a private values setting, since the dominant strategy in these auctions is to bid one’s valuation. However, as will be noted more formally in the theory section, the common value model of eBay auctions predicts that bids will unambiguously decrease with an increase in the number of bidders.

To perform this test on our data set, we report in table 2 the regression of bids normalized by book value (b) on the realized number of bidders (N). We use the linear, quadratic and logarithmic specifications of Paarsch (1992) to check the robustness of our results. We see that for each specification, there is a significant negative relationship between the number of bidders and the normalized bid. This suggests the presence of common values over private values.

We also define the following measure for the winner’s curse: *the “winner’s curse effect” is the percentage decline in bids in response to one additional bidder in the auction.* If we take the linear model at face value, the “winner’s curse effect” can be quantified at about 3.4% of the bookvalue per additional bidder.

An important criticism that could be made of the regressions in Table 2 is that the number of bidders who participate in an auction is very likely to be endogenous. If the number of bidders is correlated with attributes of the auction that we have failed to model in our simple specification, then the estimate of the slope coefficient will be biased, making it difficult to assess the extent of the winner’s curse from the above specification. Unfortunately, it is hard to think of natural instruments for the number of bidders, since factors correlated with the number of bidders should also enter into the distribution of valuations. We therefore take the evidence in table 2 as suggestive, as opposed to conclusive, about the impact of the number of bidders on the bid/bookval ratio. Also, in eBay,

¹³Though we could also expect that bidders would increase their bids in the face of increased competition, and that the effect of increasing the number of bidders on bid levels would be ambiguous. We will resolve this ambiguity in our model of bidding.

it is very hard for a bidder to know exactly how many competitors she is facing (since almost all the action transpires at the end of the auction); so the number of bidders we observe might not be the number of bidders that the bidder had in mind when forming her bid. We will attempt to address these issues more completely in sections 4 and 5 where we build and estimate a structural model of price formation in eBay.

We should note that on eBay there is an even simpler test to distinguish between a private values and a common values environment. Recall that in the “proxy-bidding” system, bidders can observe the bids of others at least until the very end of the auction and have the option to update their bids. The (weakly) dominant strategy in this auction in a private values environment is to bid one’s valuation, therefore it does not matter when a bid is placed. Furthermore, once a bid, presumably equal to the valuation of the bidder, is placed, it should not be necessary to change it.

Our data rejects both predictions of the private values setting. First, as we will show below, the bids are lumped disproportionately at the end of the auction. So it does matter when the bids are placed. In fact, we will show that bidding at the last minute turns out to be an equilibrium strategy under the common values model. Second, we find that bidders routinely update their bids. In our data, we can also observe how many total bids were submitted throughout the auction. In the set of auctions we consider, there are on average 5.2 bids. With an average 3 bidders per auction, this means that almost every bidder updates her bid once.

It is worth noting, however, that there are other possible descriptions of the information structure besides common and private values. For instance, Milgrom and Weber (1982) have developed an affiliated values models that allows for both private and common values as special cases. Also, Haile (1999a), Haile (1999b), Haile (1997) and Bikhchandani and Huang (1989) develop models where resale of the object is strategically anticipated. Table 2 and the submission of multiple bids is clearly inconsistent with a pure private values story but does not rule models with affiliation or where sellers strategically anticipate resale.

3.2 Timing of bids

Figure 1 is a histogram of final bid submission times.¹⁴ As one can see, more than 50% of final bids are submitted after 90% of the auction duration has passed. This corresponds to the last 7 hours of the auction in a 3 day auction. About 32% of the bids are submitted after 97% of the auction has passed, that is, the last 2 hours of a 3 day auction.

The winning bids tend to come even later.¹⁵ The median winning bid arrives after 98.3% of the auction time has elapsed (within the last 73 minutes of a 3 day auction), and 25% of the winning bids arrived after 99.8% of the auction time elapsed (the last 8 minutes of a 3 day auction).

The evidence shows that bidding activity is concentrated at the very end of the auction. As we

¹⁴Recall that we track only the final bids of each bidder.

¹⁵The winning bid is not necessarily the latest.

Table 3: Regression of the number of bidders on auction specific variables

NBIDDER	Coef.	Std. Err.	t-stat
(MINBID/BOOK)*SECRET	-3.7875	1.8886	-3.187
(MINBID/BOOK)*(1-SECRET)	-3.6767	.2378	-15.457
SECRET	.3840	.4752	0.808
LN(BOOK)	.4576	.0934	4.894
BLEMISH	-.1354	.3077	-0.440
NEGATIVE	-.1562	.0542	-2.879
CONSTANT	4.1218	.3368	12.239
R^2	0.5009		

explained in the section regarding common vs. private values, this finding can not be consistent with the presence of private values.

It is interesting to note that in eBay’s explanation of the proxy bidding system, the bidders are urged to decide on their private value, submit it online and wait until the conclusion of the auction. It is quite clear from the above evidence that bidders do not abide by eBay’s recommendation, and, as we will show in section 4, with good reason.

3.3 Entry

Since the previous section revealed that the number of competitors in an auction is an important determinant of bid levels, we ask the following question: what attracts bidders into an auction? In Table 3, we run a regression of the number of bidders on the minimum bid (normalized over book value), the reserve price policy (SECRET is a dummy for secret reserve price policy), the bookvalue of the item, and the negative reputation of the seller.

The minimum bid seems to be the single most important determinant of entry in our sample, regardless of there being a secret reserve price or not. The difference between setting the minimum bid equal to the bookvalue and at zero is almost 4 more bidders in the auction!

This finding is driven by the fact that the minimum bid censors the distribution of bidder valuations we observe. As we shall discuss below, in a common values setting, we will not observe bids from bidders whose signals are below a “screening-level,” which is strictly above the minimum bid. The negative coefficient on the minbid/book could be due to two reasons. The first is we observe fewer bidders in the data because the reserve exceeded the bidder’s maximum willingness to pay. A second is that high reserves discourage bidders from even attending the auction and investigating the value of the object. While this distinction may seem trivial at first, a number of theorists, including Harstad (1990) and Levin and Smith (1994) have argued that correctly accounting for the *potential* number of bidders is key in assessing the relationship between reserve prices and seller

Table 4: Regression of minimum bid levels on auction specific regressors

MINBID/BOOK	Coef.	Std. Err.	t-stat
LN(OVERALL)	-.000376	.0158	-0.024
LN(NEGATIVE)	-.03754	.0408	0.920
LN(BOOK)*SECRET	-.01376	0.028	-0.479
LN(BOOK)*(1-SECRET)	-.1102	.02611	-4.221
SECRET	-.6531	.1406	-4.644
BLEMISH	-.1033	.0646	-1.598
CONSTANT	.9960	.1078	9.085
R^2	0.1668		

revenue. We will emphasize this distinction while building our structural model in sections 4 and 5, where we find correctly accounting for the potential number of participants is key in understanding how sellers set their reserve prices.

We also find that the existence of a secret reserve price does not affect entry significantly. All things being equal, one might expect that a secret reserve price should deter entry, for rational bidders should realize that their chance of winning an auction with a secret reserve price is lower. However, as we point out below, secret reserve price auctions are almost exclusively used for auctioning high value items, hence even if the probability of actually winning the auction might be low, the expected profit of the bidder can be quite high.

Aside from reserve price/minimum bid policy, bookvalue also appears as a significant determinant of entry; however, the parameter estimate in Table 3 means that a ten-fold increase in bookvalue is needed to add another bidder to the auction. Also, a seller with negative feedback attracts fewer bidders (7 or 8 negative comments drive away one bidder).

Given that the minimum bid is by far the most important determinant of entry, what does the seller's choice of a minimum bid depend on? In table 4, we regress the minimum bid (normalized over book value) on SECRET, BLEMISH and the natural logs of OVERALL (which we take as a proxy for seller experience), NEGATIVE and BOOKVAL . We find that the seller's experience or reputation do not affect the choice of minimum bid. However, the minimum bid is decidedly lower if the seller is running a secret reserve price auction, or if the seller is running an ordinary (no secret reserve price) auction with a high book value.

One way to interpret the results in Table 4 is to say that sellers of high book value items are willing to risk selling the item for a low price in order to induce entry to their auctions. Their hope is that after at least two bidders have entered the auction, the bidding will reach a competitive level, provided that both bidders have valuations that are close to the book value of the item.

This argument is only true to a certain extent, however. For high book value items, sellers

Table 5: Quantiles of Book Value Across Auction Formats

Quantile	Secret Reserve	Posted Reserve
5%	4	4.75
10%	6.75	6
25%	9	8.5
50%	42.5	13.5
75%	148.5	22.25
90%	620	41.25
95%	910	56
99%	3700	200

overwhelmingly favor a secret reserve price auction over an ordinary auction, as seen in Table 5. Conducting a secret reserve price auction allows the seller to set a low minimum bid to attract entry¹⁶ at the same time ensuring a high price conditional upon sale¹⁷. Unfortunately eBay does not divulge the secret reserve prices at the close of the auction, therefore we can not compare the secret reserve price to minimum bid levels in ordinary auctions; however, the fact that far fewer secret reserve price auctions result in a sale (50% vs. 84%) indicate that secret reserve price levels are higher than the minimum bids in ordinary auctions.

The fact that we see the use of secret reserve prices for items with higher bookvalues suggests that there might be a threshold level of the seller's value for the coin after which using a secret reserve price dominates making the reserve price public. We should note that this would run counter to the only result we have been able to find in the auction theory literature regarding the use of secret reserve prices: Elyakime et al. (1994) prove that in a private values framework, making the reserve price public always yields the seller higher revenue than when it is secret. However, Elyakime et al. (1994) do not incorporate the effect of endogenous entry in their results. It is also conceivable that in a common value setting, a different conclusion might arise. Although we do not know of any general result concerning the benefits/costs of using secret reserve prices, the structural model that we build in the sections 4 and 5 allows us to simulate the effects of different reserve price policies on seller revenues.

3.4 Revenues, profit margins and values

The summary statistics in Table 1 reveal that the average revenue (BID1) in the auctions are lower than the average book value of the object being sold. The average revenue (as a percentage of book

¹⁶The average minimum bid in secret reserve price auctions is set at 25% of bookvalue, with several equal to 0. In auctions with no secret reserve price the minimum bid is, on average, set to 70% of bookvalue.

¹⁷Kaiser and Kaiser (1999) confirms this intuition, so do some conversations we overheard in bidder discussion groups on eBay.

Table 6: Regression of revenue/bookvalue on auction specific variables

REVN	Coef.	Std. Err.	t-stat
NBIDDER	.0356	.0071	4.988
MINBID/BOOK	.1833	.0339	5.396
SECRET	.0078	.0446	0.176
BLEMISH	-.1286	.0493	-2.605
NEGATIVE	-.0021	.0080	-0.267
AVGEXP	-.0004	.0004	1.001
PREWIN1	-0.0002	.0003	-0.570
CONSTANT	.7431	.0438	16.932
R^2	0.3327		

value) conditional on sale is 95.8%. Keeping in mind that some auctions did not result in a sale, the unconditional revenue is 81.3%.

In table 6, we define REVN to be the normalized revenue of the item and regress it on the same set of regressors as in the entry regression. We find that an increase in the number of bidders and in the minimum bid increase the revenue and that reported blemishes decrease revenue. Auction type, seller's negative reputation, the average experience of bidders, and the experience level of the winning bidder are not significant determinants of revenue.

The average profit margin of the bidders (taking the book value as the true value of the item) is 4.2% (with a standard deviation of 30.1%). However, this does not take the shipping and handling costs in to account. As we noted in the data section, we have data on only about a third of the sample on the shipping and handling terms. The average profit margin net of shipping and handling costs on this sub-sample is -16.2% (std. 37.3%). Using the average shipping and handling cost of \$2.18 in Table 1 for the entire sample, we find that the net profit margin is -15.4% (std. 36.4%). However, since the standard deviations are very large and the estimates are highly dependent on our measurement of the book value, these figures should be taken with a great deal of skepticism.¹⁸

What values do bidders assign to the coins on auction? The revenue regression above gives a biased answer to this question, since revenue is the second highest bid in the auction. Even if we believe a private values framework, in which bids are equal to the valuation, we can not answer general questions about the valuation of the objects just by looking at the revenue, since revenue is the second-order statistic of the valuation distribution.

Table 7 is the regression of individual bids (normalized over book value) on the regressors in Table 4. We see that once again the minimum bid enters as a significant determinant of overall bids. A reported blemish in the item reduces the bids by 10%. Bidder experience also enters as a weak,

¹⁸For example, if we use the median value of the net profit margin as our summary statistic, we get -1.5% profit net of shipping and handling.

Table 7: Regression of normalized bids on auction specific variables

BIDN	Coef.	Std. Err.	t-stat
MINBID/BOOK	.44316	.0239	18.434
SECRET	-.0294	.0185	-1.584
BLEMISH	-.0956	.0249	-3.828
NEGATIVE	-.0047	.0055	0.851
PREWIN	0.00028	.00012	2.365
CONSTANT	.6214	.0149	41.581
R^2	0.2569		

but significant factor in determining bid levels.

A straightforward but misleading conclusion from the regressions in Table 6 and Table 7 is to say that revenue and bids increase with the minimum bid. It is easy to see why the conclusion is misleading for the regression table 6: the bids in this regression are observed conditional on the fact that they exceed the minimum bid. Hence, the coefficient on MINBID is subject to selection bias because only bidders with a high willingness to pay for the object will choose to participate in the auction. While the results of this regression are suggestive, we will need to build a structural model in order to consistently estimate the relationship between the reserve price and the bids submitted in the auction.

3.5 Bidder experience

A breakdown of bidder experience measured by number of auctions previously won at eBay reveals a very heterogeneous bidder population (ranging from first time participants to a bidder who has won 557 auctions) with the median bidder having 21 previous wins. However, as was noted in the last section, neither the average experience level of the bidders in an auction, nor the experience of the winning bidder seem to have a significant effect on the revenue of the seller. Also, as seen on Table 5, experience actually has a very weak positive effect on individual bids, i.e. experienced bidders tend to bid higher than inexperienced bidders.

These results point out that despite the heterogeneity of the bidder population, informational asymmetries or differing degrees of bidder sophistication are not likely to be significant issues in this setting. Therefore we will refrain from incorporating ex-ante informational asymmetries in our model.

4 A model of bidding in eBay coin auctions

In this section, we develop a model of bidding in eBay coin auctions that we believe captures the main findings of last section and that will allow us to further our ability to understand observed bidding behavior. In the next section, we embed this model in an econometric procedure to estimate its exogenous parameters. This requires that we impose structure onto the bidding process. We understand that it is difficult for a theoretical model to capture the complexity and idiosyncrasy of actual bidder behavior. However, the imposed structure is necessary to answer the following questions without ambiguity:

1. What is the informational environment that rationalizes observed bidder behavior? We suggested in the previous section that there is evidence against a private values framework; however, even in a common values framework, we still need to answer the following questions: What does the common value depend on? How diffuse are the bidders' information? It is very difficult to interpret the regression results in the previous section to get good answers to these questions, since the data suffers from both an endogeneity bias due to the endogenous nature of entry, and a truncation due to the minimum bid level. What is needed is a specification that accounts for the fact that entry is endogenous, and that some potential bids are unobserved due to the minimum bid.
2. Another question of use to users of eBay that emerged in the previous section is the following: How does the choice of a minimum bid and/or a secret reserve price affect the revenues of a seller? Although our results in section 3.2 seem to indicate that, at least beyond a certain book value, a secret reserve price auction might be more advantageous to the seller, we also pointed out that there is no known theoretical justification for why this should be true. Our structural estimates allow us to simulate the revenue effects using a secret reserve over a public reserve price and to answer this question unambiguously.
3. Statements about welfare require that we use a structural model. For instance, if we wish to measure the change in surplus due to the entry of an additional bidder, we must compute expected utility which in turn requires knowledge of the model's structural parameters.

Our model assumptions are based on the empirical findings from the previous sections. A pure private values model would be inappropriate given the observations of the previous section. We will therefore focus on a common values model. We also account for the fact that the number of bidders in an auction is a random variable whose realization depends strongly on auction specific observables, most importantly, the minimum bid level. In doing so, we assume that the bidders are ex-ante symmetric, as bidder experience was not found to be a significant determinant of bidding behavior.

It is not possible, a priori, to rule out the importance of affiliation or a strategic model of with resale in the econometric specification. A specification with common values and endogenous entry was chosen among other competing hypotheses for two reasons. First, most auction models without pure private values are difficult to estimate since evaluating a bidder's expected profits typically requires solving a high dimensional numerical integration problem. In this paper, we have developed new techniques that can be applied to common value models with endogenous participation that allow us to estimate a structural model that heretofore has proven intractable. The econometric tools needed to solve our problem will be somewhat specific to the information structure of the model. We view structural estimation of Internet auctions under these alternative estimation structures as laudable but beyond the scope of this project since an entirely new and quite complex econometric machinery will have to be developed.

Second, we have found, in work leading up to this paper that models where the number of bidders are known *before* bidding begins have unsavory econometric properties in our application. These unsavory properties come from the fact that the bid functions when the number of bidders becomes ten or more become very unstable computationally. This instability arises from the simple fact that many valuation have a very low ex ante probability of winning the auction. For instance, take the case where there are 20 bidders. A bidder who draws the median valuation (conditional on available all publicly available information) wins only if all other bidders draw a lower valuation. The probability of this event is $.5^{19}$ or roughly, two parts in one million. In our application, where the average value of objects for sale is about \$40 a bidder's expected profits conditional upon having the median valuation will be zero to several significant digits. The set of bid functions that are approximate equilibria is extremely large. We believe that relying on a model where optimization takes place on a much finer scale than the minimum bid increment (a penny) makes little sense. Also, such models are very difficult to study computationally. (This problem is discussed more completely in Bajari (1999)).

Models with endogenous entry are not as susceptible to this problem since what matters to the bidder is the distribution over the number of bidders in the auction and not the number who actually show up. The bid functions in the model with endogenous entry are much more robust computationally and the distinction between equilibrium and approximate equilibrium is much less important since, as we will show in our results section, bidders have non-trivial ex-ante expected profits.

In what follows, we first specify a stochastic entry process motivated by the regression in Table 3. We then focus on the bidding game conditional on entry. We argue that the eBay auction, despite its ascending nature, differs significantly from the ascending auction models investigated in the theoretical literature. Instead of capturing the full dynamic nature of the game, we point out that there exists no symmetric equilibrium of the game in which the observable actions of bidders can fully reveal their private information. Moreover, we demonstrate that there exists a symmetric

equilibrium in which bidders choose to reveal no information at all, i.e. they all wait until the very end of the auction¹⁹. In this equilibrium, bids can be modeled as coming from a sealed-bid second price auction, for which the well-known results of Milgrom and Weber (1982) apply. Our equilibrium analysis also incorporates the stochastic nature of the entry process.

4.1 Model primitives

We assume that the number of bidders in auction t is a Poisson random variable, with rate λ_t , which depends on the regressors in Table 3. The Poisson random variable arises as the limiting distribution of a sum of independent Bernoulli random variables, each with small probability of success. Since there are millions of users on eBay and thousands of coin collectors browsing auction listings everyday, of which only a few decide to bid in a given auction, we believe the Poisson specification is appropriate in this setting.

We assume that bidders are risk neutral expected utility maximizers. We eschew any dynamic considerations that bidders might have and assume that they focus entirely on the auction at hand²⁰. To model the information structure, we assume the symmetric common value setting of Wilson (1977). That is, there is an unknown common value, v of the coin being auctioned, which is distributed according to $f_v(v)$ over the support $[\underline{v}, \bar{v}]$. Since the collectible coins market is characterized by frequent buying and selling, v can represent the unknown resale value of the item. We assume that the prior belief about the distribution of the common value is shared by all bidders.

Given the particular realization of the common value, which is unobserved, bidder i receives a signal x_i which is distributed according to $f_x(x_i|v)$ over the support $[\underline{x}, \bar{x}]$. We assume that conditional on v , x_i are identically and independently distributed and the distribution of bidder signals is common knowledge to all bidders.²¹ We also assume that $f_x(x_i|v)$ satisfies the monotone likelihood ratio property of Milgrom and Weber (1982) and hence the signals are affiliated²².

4.2 Bidding

4.2.1 Auction format

eBay uses an open ascending auction format. However, the exact format does not fit neatly under the category of “open-exit English auctions” studied by Milgrom and Weber (1982), in which the

¹⁹This mode of bidding is also called “sniping” Kaiser and Kaiser (1999). It is interesting to note that there exist programs downloadable from the Web to facilitate such last-minute bids.

²⁰We also recognize, but do not model the fact that on eBay similar objects are often being sold side-by-side and that someone in the market for a certain coin can “smooth” her strategies across different auctions.

²¹Observe that this specification of the common value environment is entirely analogous to that in Section 3.1.

²²Briefly, affiliation implies that if $v' > v$ and $x_i > x_j$, then $f_x(x_i|v')f_x(x_j|v) > f_x(x_i|v)f_x(x_j|v')$. Roughly, this means that signals are positively correlated: a high signal for a bidder makes high values for other bidders’ signals more likely.

auctioneer raises the price of the object until all but one of the bidders has dropped out.²³ A bidder can not rejoin the auction once she has dropped out. The key to the analysis in Milgrom and Weber (1982) is that the drop-out decision of a particular bidder conveys information to other bidders, who will update their estimate of the item's true worth and will change their drop-out points. Hence bidding strategy in this type of an ascending auction differs from that in a sealed-bid auction.

However, in eBay agents have the opportunity to update their bids and rejoin the race anytime before the auction ends. With the option to revise one's bid, bidders might not be able to infer others' valuations by observing their drop-out decisions, since "drop-outs" can be insincere. In fact, we will show that there is a Nash equilibrium in which the eBay auction is equivalent to a sealed-bid second-price auction.

To make the above argument more formal, let us view the eBay auction as a two stage auction. Taking the total auction time to be T , the first stage auction is an open-exit ascending auction played until $T - \varepsilon$, where $\varepsilon \ll T$ ²⁴. The drop-out points in this auction are openly observed by all bidders, who will be ordered by their drop-out points in the initial auction, $\theta_1 \leq \theta_2 \leq \dots \leq \theta_n$, with only θ_n unobservable (bidders can only infer that it is higher than θ_{n-1}). The second stage of the auction transpires from $T - \varepsilon$ to T , and is conducted as a sealed bid second price auction in which every bidder, including those who dropped out in the first stage, are given the option to submit a bid, b . The highest bidder in the second stage auction wins the object.

This model of the eBay auction captures the ability of bidders to revise their bids. It also is consistent with the empirical finding that quite a few final bids are submitted within the last minutes of the auction. As our experience as bidders on eBay show, the last minutes of the auction do resemble a sealed bid auction more than an ascending auction, since delays in e-mail transmission and on the eBay server may make it difficult to follow the bidding in real time²⁵.

Given the above setup, we make the following claim:

Proposition 1 *The first stage drop-out points $\theta_i, i = 1, \dots, n$ can not be of the form $\theta(x_i)$, a monotonic function in bidder i 's signal.*

Proof 1 *Suppose $\theta_i = \theta(x_i)$ a monotonic function in x_i . Then, since $\theta(\cdot)$ is invertible, at the beginning of the second stage of the auction, the signals of bidders $i = 1, \dots, n - 1$ become common knowledge. Then, the second stage bid functions will be the expected value of the object given the*

²³The eBay auction format is also different from traditional oral-ascending auctions, since the proxy-bidding system precludes an initial "jump" bid, through which the first bidder can raise the auction price significantly more than the bid increment.

²⁴ ε is the time frame in which bidders on Ebay can not update their bids in response to others. This can be on the order of 5-10 minutes, or even longer; as one of the authors had the unpleasant experience of receiving the notification that he was outbid about 2 days after an auction closed.

²⁵Visitors to eBay can attest to the fact that navigation can be quite slow at times.

information available to the bidders²⁶:

$$\begin{aligned} b_{i \neq n} &= \max\{\theta_i, E[v|x_1 = \theta^{-1}(\theta_1), \dots, x_{n-1} = \theta^{-1}(\theta_{n-1}), x_n \geq \theta^{-1}(\theta_{n-1})]\} \\ b_n &= \max\{\theta_n, E[v|x_1 = \theta^{-1}(\theta_1), \dots, x_{n-1} = \theta^{-1}(\theta_{n-1}), x_n = x_n]\} \end{aligned}$$

Observe that this results in identical bids for bidders $i \neq n$, provided that their signals are high enough, since they all form the same expectation of the common value. The expectation formed by the highest bidder in the first stage is a little different, since she knows her own value with certainty.

In this situation, there is a profitable deviation. If bidder j decreases her drop-out point in the first stage to $\hat{\theta}_j < \theta_j$, then in the second stage

$$b_{i \neq j}(\hat{\theta}_j) \leq b_{i \neq j}(\theta_j)$$

since the conditional expectations are decreasing in θ_j by Milgrom and Weber (1982) 's Theorem 5. But since all other bidders are bidding sincerely in the first stage, j 's bid will not change, and she will unilaterally increase her probability of winning the auction.

If the bidders can not bid sincerely in the first stage of the auction, then the ascending auction aspect of eBay leads to less information revelation than in the models of ascending auctions in Milgrom and Weber (1982), Wilson (1995) and the structural econometric model of Hong and Shum (1997). In fact, the following result reveals that the eBay auction format might not have any advantages over a sealed bid second price auction:

Proposition 2 *Bidding 0 (or not bidding at all) in the first stage of the auction and participating only in the second stage of the auction is a symmetric Nash equilibrium of the eBay auction. In this case the eBay auction is equivalent to a sealed bid second price auction.*

Proof 2 *Bidding 0 in the first stage reveals nothing about a bidder's signal in the first stage. Therefore the second stage is just a sealed bid second price auction, where each bidder knows only her own signal. In this case, the symmetric equilibrium bid function, as derived by Milgrom and Weber (1982), is $b(x) = v(x, x)$, where $v(x, y) = E[v|x_i = x, y_i = y]$, with $y_i = \max_{j \in \{1, \dots, n\} \setminus i} \{x_j\}$. To rule out profitable deviations, observe the following: if bidder j were to bid $\theta_j > 0$, then it would be evident that $x_j > 0$, and since signals are affiliated, $E[v|x_i = x, y_i = x, x_j > 0] \geq E[v|x_i = x, y_i = x]$, so bidder j would unilaterally decrease her probability of winning²⁷.*

²⁶Since bidders are risk neutral, this follows from exactly the same argument behind sincere bidding in private value second price auctions: if, given the information available, I bid higher than my conditional expectation of the value of the item, then on average, I make a loss. If, on the other hand, I bid lower than my conditional expectation, then I lower the probability I will win, but since my payment does not change (since it is the bid of the second highest bidder), my expected profit is lower.

²⁷In fact, we can generalize the above equilibrium to a continuum of "symmetric pooling equilibria" where bidders with signal $x \geq x^*$ bid θ^* in the first stage. However, this cutoff point is very difficult to coordinate on, and zero seems to be the most likely focal point for such symmetric equilibria.

The main drawback of the above analysis is that our data reveals that bidders on eBay do not necessarily follow this equilibrium. This is not very surprising, because aside from this one symmetric equilibrium, there might exist many asymmetric equilibria. However, the main point of our logic is clear: there is no *symmetric* equilibrium of the eBay auction where bidders can learn much from the drop-out decisions of others. Therefore if we are going to focus on symmetric equilibria in our econometric evaluation, we have no choice but to model the bids as if they were submitted in a sealed-bid auction.

As an interesting aside, we should note that one of the main competitors of eBay, Amazon's auction site, extends its auctions beyond the prespecified closing time until bidding activity ceases²⁸. This feature changes the analysis above, since the "sealed-bid second-stage" of the auction no longer exists (T is no longer set, so bidders can not coordinate on the "blind-bidding" period $[T - \varepsilon, T]$). In this case, "drop-outs" are final, since the end of the auction is dependent on bidder activity. The open-exit ascending auction framework of Milgrom and Weber (1982) can be applied here. Moreover, under the assumption of affiliated signals, it can be shown that the Amazon auction results in higher expected revenue to the seller than in the symmetric equilibrium of the eBay auction analyzed above²⁹. Hence Amazon has a more "seller-friendly" mechanism than eBay.

There are other reasons why auction sites might introduce features like this to prevent the type of "sniping" we observe on eBay. First of all, it would be very difficult to regulate transactions traffic on the Web server if all bids were to come in at the last few minutes of the auction. This would cause higher congestion and more frequent service outages, driving customers away. Also, sniping could discourage occasional bidders from entering the auction, since they might not find it worthwhile to wait patiently until the very end of an auction.

So why does eBay stick with its hard-deadline policy? One reason could simply be resistance to change: since quite a few customers have invested time into deciphering eBay's system, they might protest attempts to change the auction mechanism (which was quite recently the case, when eBay tried, unsuccessfully, to modify its reserve price rules). Another is that hard-deadlines maybe seen beneficial by bidders and sellers with time constraints, in that they prevent auctions from dragging on. Yet eBay does try to convince bidders that "sniping" is not a strategy they should use. As we mentioned above, eBay's bidding tutorial assumes a private values framework and explains that with private values, the dominant strategy is to bid one's valuation and sit back.

4.2.2 Bidding in a sealed bid second price auction with a random number of bidders:

Since we are using a sealed bid second price auction as our model of bidding conditional upon entry, we will develop the bidding strategies in this auction format. If the number of bidders is common knowledge, then the bid function derived by Milgrom and Weber (1982) applies (recall that

²⁸More precisely, the auction duration is extended until there are no more bids within a 10 minute period.

²⁹The revenue comparison is between an English auction and a sealed-bid second price auction. See Milgrom and Weber (1982), Theorem 11.

$y_i = \max_{j \in \{1, \dots, n\} \setminus i} \{x_j\}$:

$$\begin{aligned} b(x) &= E[v|x_i = x, y_i = y] \\ &= \frac{\int_{\underline{v}}^{\bar{v}} v f_x^2(x|v) F_x^{n-2}(x|v) f_v(v) dv}{\int_{\underline{v}}^{\bar{v}} f_x^2(x|v) F_x^{n-2}(x|v) f_v(v) dv} \end{aligned}$$

where $f_v(v)$ is the density function of the common value and $f_x(x|v)$ and $F_x(x|v)$ are the conditional density and distribution functions of the signal.

However, as we noted in the section on entry, the number of bidders is not common knowledge on eBay. The best we can hope for is that bidders have rational expectations about the number of competitors they are going to face. Therefore we have to introduce this uncertainty into our model of bidding.

Proposition 3 *Let $p(n, \lambda)$ be the pdf of the number of bidders in the auction conditional upon entry of particular bidder to the auction. Also define $y_i^{(n)} = \max_{j \in \{1, \dots, n\} \setminus i} \{x_j\}$ and $v(x, y, n) = E[v|x_i = x, y_i^{(n)} = y]$. Then if the function*

$$b^*(x, \lambda) = \frac{\sum_{n=2}^{\bar{N}} v(x, x, n) f_{y_i^{(n)}}(x|x) p(n, \lambda)}{\sum_{n=2}^{\bar{N}} f_{y_i^{(n)}}(x|x) p(n, \lambda)} \quad (1)$$

is increasing in x , then it is a symmetric Nash equilibrium of the common value second price auction game with a random number of bidders.

Proof 3 *See Appendix A.*

As can be seen by inspection, the bid function in the case of stochastic bidders is a weighted average of bid functions with a given number of bidders, where the weighting factor is $\frac{f_{y_i^{(n)}}(x|x) p(n, \lambda)}{\sum_{n=2}^{\bar{N}} f_{y_i^{(n)}}(x|x) p(n, \lambda)}$. So, once the bid functions in the deterministic case are solved for, we can easily compute the bid function in the stochastic case, using $p(n, \lambda)$.

4.2.3 Properties of the bid function

Although we have the bid function in closed form, to analyze its properties and to use it in a computational procedure, we still have to evaluate the integrals in the numerator and denominator numerically. Fortunately, if we assume that $v \sim N(\mu, \sigma_v^2)$ and $x|v \sim N(v, k\sigma_v^2)$, we can compute this bid function with very high accuracy using the Gauss-Hermite quadrature method (see Judd (1998) and Stroud and Secrest (1968) for details).

Figure 2 plots the bid functions in an auction where are $\mu = 1, \sigma = 0.6$ and $k = 0.3$. We see that the functions look quite linear (with a slope of about 0.85) and are indeed increasing in x .

More interestingly, these plots also show that bids are also decreasing in λ , the expected number of bidders in the auction. So the “winner’s curse” does increase with the number of competitors.

The apparent linearity of the bid functions might raise questions about the validity of the computational procedure or whether there is a theoretical result remaining to be proved about the behavior of bid functions under the distributional we have made. To answer the first question, we have repeated our calculations for an increasing number of quadrature points, and, instead of using the Gauss-Hermite quadrature evaluation points, have also tried simulating the integrals. Both methods gave similar results. As to whether it can theoretically be shown whether the bid functions are linear, we have only been able to verify that linearity holds only in the (deterministic number of bidders) case when $n = 2$. In this case, the bid function is:

$$\begin{aligned} v(x, x) &= E[v|x_1 = x, x_2 = x] \\ &= \frac{k\mu + 2x}{2 + k} \end{aligned}$$

It could be conjectured, then, that as the number of bidders gets larger, the departure from linearity remains slight.

Another “linearity” property that we did succeed in verifying theoretically is that bid functions scale linearly when the signal distribution undergoes a linear transformation. That is:

Proposition 4 *Let $b(x|\mu^\circ, \sigma^\circ, k, \lambda)$ be the bid function in an auction where the common value v is distributed normally with mean μ° and standard deviation σ° . We claim that in an auction where the common value is distributed with mean $r_1\mu^\circ + r_2$ and standard deviation $r_1\sigma^\circ$, the equilibrium bid function is $r_1b(r_1x + r_2|\mu^\circ, \sigma^\circ, k, \lambda) + r_2$. That is, if it was optimal to bid b with a signal x in the original auction with the common value distributed $N(\mu^\circ, \sigma^\circ)$, it is now optimal to bid $r_1b + r_2$ when your signal is $r_1x + r_2$, with the common value distributed $N(r_1\mu^\circ + r_2, r_1\sigma^\circ)$.*

Proof 4 *See Appendix A.*

Observe that the above argument depends only on the linearity of the expectation of the common value given other bidder’s signals, and the existence of monotonic bid functions. Therefore, we can extend this result to any auction where bidder signals are distributed according to a location-scale family, and where the common value or the expectation of the common value is a linear function of the bidders’ signals.

This is a key result in our estimation procedure, because it allows us to compute the bid function only once for a base-case auction. We can then apply an affine transformation to this “pre-computed” function to find the bid function in an auction where the distribution of the common value has a different mean and variance.

It is worthwhile to note that bid functions will be likewise linearly scalable in many commonly studied auctions, such as the first price sealed bid auction. Elyakime et al. (1994) note that efficient

evaluation of the bid function is a key problem in structural estimation of auction models. The simplification we propose can be used to for structural estimation in a wide variety of models where the bid function is linearly scalable and the distribution of valuations is characterized by its first two moments.

4.2.4 Minimum bids

eBay allows the seller to set a minimum bid. Since minimum bids induce truncation on the set of bids (and hence valuations) we observe, it is necessary to incorporate the truncation region in our likelihood function.

However, in the case of a minimum bid r , equilibrium strategies are derived by Milgrom and Weber to be:

$$\begin{aligned} b(x) &= v(x, x) & \text{for } x \geq x^* \\ b(x) &< r & \text{for } x < x^{*30} \end{aligned}$$

where $x^* = x^*(r)$ is called the screening level and is given by

$$x^*(r) = \inf\{E[v|x_i = x, y_i < x] \geq r\}$$

Since the above expectation is monotonic in v (due to affiliation), the screening level is the solution of:

$$\frac{\int_v^v v f_x^2(x^*|v) F_x^{n-1}(x^*|v) f_v(v) dv}{\int_v^v f_x^2(x^*|v) F_x^{n-1}(x^*|v) f_v(v) dv} = r$$

We calculate x^* for a string of values of r and use interpolation to find x^* in our likelihood function evaluations.

For the case of a stochastic number of bidders, the above equilibrium generalizes to:

$$\begin{aligned} b(x, \lambda) &= \frac{\sum_{n=2}^{\bar{N}} v(x, x, n) f_{y_i^{(n)}}(x|x) p(n, \lambda)}{\sum_{n=2}^{\bar{N}} f_{y_i^{(n)}}(x|x) p(n, \lambda)} & \text{for } x \geq x^* \\ b(x, \lambda) &= 0 & \text{for } x < x^* \end{aligned} \quad (2)$$

where x^* is the expected screening level

$$x^*(r, \lambda) = \inf\{E_N E[v|x_i = x, x_i < x, N = n] \geq r\}$$

which can be calculated as the solution of:

$$p(1, \lambda)E[v|x] + \sum_{n=2}^{\bar{N}} p(n, \lambda) \frac{\int_v^v v f_x^2(x^*|v) F_x^{n-1}(x^*|v) f_v(v) dv}{\int_v^v f_x^2(x^*|v) F_x^{n-1}(x^*|v) f_v(v) dv} = r \quad (3)$$

Once again, normality assumptions allow us to compute the integrals in the above expressions with very high accuracy using the Gaussian quadrature method. However, we also need to solve for x^* in expression (3) using a numerical procedure. We use a simple Newton-Raphson algorithm to find this solution.

Figure 3 plots the screening level as a function of the minimum bid for the auction with $\mu = 1, \sigma = 0.6$ and $k = 0.3$. The screening level appears to depend on the minimum bid in a similar near-linear fashion. From the two cases, $\lambda = 3$ and $\lambda = 5$, that are plotted in this figure, we also see that the screening level is decreasing in λ and, as shown by Milgrom and Weber (1982), is above the minimum bid.

4.2.5 Secret Reserve Prices

We will use a shortcut in modelling the effect of a secret reserve price on bidding behavior: we will assume that bidders believe that the seller's valuation for the item comes from the same information structure as theirs, i.e. the seller is yet another bidder. This makes sense if the seller has good enough outside options, i.e. she's not constrained to sell the item immediately, or on eBay. Since quite a few sellers on eBay are professional coin traders, this condition is likely to be satisfied.

With this assumption, the only difference between the secret reserve price auction and a regular auction is that now potential bidders know that they face at least one competitor, the seller, so now:

$$b^{SR}(x, \lambda) = \frac{\sum_{n=2}^{\bar{N}} v(x, x, n) f_{y_i^{(n)}}(x|x) p(n-1, \lambda)}{\sum_{n=2}^{\bar{N}} f_{y_i^{(n)}}(x|x) p(n-1, \lambda)} \quad (4)$$

A minimum bid is similarly incorporated into the secret reserve price setting. The screening level, $x^{*SR}(r, \lambda)$ in this case is the solution of:

$$\sum_{n=2}^{\bar{N}} p(n-1, \lambda) \frac{\int_v^v v f_x^2(x^*|v) F_x^{n-1}(x^*|v) f_v(v) dv}{\int_v^v f_x^2(x^*|v) F_x^{n-1}(x^*|v) f_v(v) dv} = r \quad (5)$$

5 Structural Estimation of the Model

In order to make predictions using the above model, we need to compute the bid functions for a set of "reasonable" parameter values. The problem is that the "structural parameters" of our model, i.e. the parameters of the prior density of the common value $f_v(v)$ and the conditional density of bidders' signals, $f_x(x|v)$, are not observable from data. The next section outlines a structural estimation method to uncover these unobservables from data.

5.1 Specification of the information structure and entry process

We will assume that the prior distribution on the common value and the conditional distributions of bidders' signals are normally distributed. The normality assumption is motivated not only by its computational convenience, but also from an inspection of the bid distribution (normalized by book value) in Figure 4, which resembles a normal distribution fairly well. We specify the prior distribution of the common value in auction t to be $N(\mu_t, \sigma_t^2)$, where:

$$\begin{aligned}\mu_t &= \mu_t(\beta, BOOKVAL_t, BLEMISH_t) \\ &= \beta_1 BOOKVAL_t + \beta_2 BLEMISH_t BOOKVAL_t - 2.18\end{aligned}\tag{6}$$

$$\begin{aligned}\sigma_t &= \sigma_t(\beta, BOOKVAL_t, BLEMISH_t) \\ &= \beta_3 BOOKVAL_t + \beta_4 BLEMISH_t BOOKVAL_t\end{aligned}\tag{7}$$

This choice for modeling the common value is motivated in part by our reduced form finds that objects with a blemish are sold at a discount in the auction. Also, since agents must pay shipping and handling fees, we subtract \$2.18 in computing μ_t , the average shipping and handling fee for the mint/proof sets³¹. The parameter vector $\beta = (\beta_1, \beta_2, \beta_3, \beta_4)$ is held constant across all auctions and all bidders so that identification of these parameters is possible.

The distribution of individual signals conditional on the common value is $N(v_t, k\sigma_t^2)$ where $k = \beta_5$ is a parameter to be estimated. This parameter models the uncertainty in the private signals. In the case where $k = 0$, we have the private values situation. We hold k constant across auctions for identification.

We use a reduced form specification for λ_t , the Poisson parameter for entry, motivated by the regression in Table 3. In order to avoid negative values of λ_t , we use its log as the dependent variable:

$$\begin{aligned}\log \lambda_t &= \beta_6 + \beta_7 SECRET_t + \beta_8 LN(BOOK_t) + \beta_9 NEGATIVE_t \\ &\quad + \beta_{10} SECRET_t \frac{MINBID_t}{BOOKVAL_t} + \beta_{11} (1 - SECRET_t) \frac{MINBID_t}{BOOKVAL_t}\end{aligned}\tag{8}$$

5.2 Likelihood function

The equilibrium bid functions we derived in the previous sections generate a likelihood function in a natural way. Let Ω_t be shorthand for $\{\lambda_t, SECRET_t, r_t\}$; the Poisson parameter for auction t , the auction type, and, the minimum bid. Then $b(x|v, \sigma_t, k, \Omega_t)$ denotes the equilibrium bid function

³¹Our data on shipping and handling has a lot of missing values, so we chose not to use the individual values.

given a particular (unobserved) realization of the common value, the variance of the conditional distribution of the bidder's signal and the information contained in Ω_t .

Assuming that the bids in our dataset are generated by equilibrium bidding, we can calculate the distribution of the unobserved signals given the observed distribution of bids. Let $\phi(b|v, \sigma_t, k, \Omega_t)$ denote the inverse bid function, mapping the observed bid to the unobserved individual signal x . Then $f_b(b_i|v, \sigma_t, k, \Omega_t)$, the probability that player i submits a bid b_i in auction t is:

$$f_b(b_i|v, \sigma_t, k, \Omega_t) = f_x(\phi(b_i|v, \sigma_t, k, \Omega_t)) \phi'(b_i|v, \sigma_t, k, \Omega_t) \quad (10)$$

Unfortunately, neither the inverse bid function $\phi(\cdot)$ nor its derivative $\phi'(\cdot)$ are available in closed form for the auctions we consider. In Appendix B, we outline a computational procedure to evaluate $\phi(b_i|v, \sigma_t, k, \Omega_t)$, $\phi'(b_i|v, \sigma_t, k, \Omega_t)$ efficiently.

To find the likelihood for the observed bids $b = (b_1, \dots, b_{n_t})$ it is necessary to account for the fact that the b_i are correlated through the realization of v . Also, remember that for the winning bidder, we do not observe the bid directly, but his payment. Then in an auction where there is no minimum bid, the probability of observing the bids $b = (b_1, \dots, b_{n_t})$ is:

$$\begin{aligned} & f(b_1, b_2, \dots, b_{n_t} | \mu_t, \sigma_t, \Omega_t) \\ &= p(n_t | \lambda_t) * \int_{-\infty}^{\infty} (1 - F_x(\phi(b_1|v, \sigma_t, k, \Omega_t))) * \\ & \quad \prod_{i=2}^{n_t} f_b(b_i|v, \sigma_t, k, \Omega_t) f(v | \mu_t, \sigma_t) dv \end{aligned}$$

Observe that we also to took into account that n_t bidders showed up as a result of a stochastic process, by adding in $p(n_t | \lambda_t)$ ³².

In an auction with a minimum bid, we have to take into account that some entrants might have drawn signals that are lower than the screening level $x^*(\lambda_t, r_t)$ and hence have not submitted a bid.

$$\begin{aligned} & f(b_1, b_2, \dots, b_n | \mu_t, \sigma_t, \Omega_t) \quad (11) \\ &= \sum_{j=n_t}^{\bar{N}} p(j | \lambda_t) * \int_{-\infty}^{\infty} \{F_x(x^*|v, \sigma_t, k, \Omega_t)^{j-n_t} (1 - F_x(\phi(b_1|v, \sigma_t, k, \Omega_t)))^{1_{\{n_t \geq 1\}}}\} * \\ & \quad \prod_{i=2}^{n_t} f_b(b_i|v, \sigma_t, k, \Omega_t) f(v | \mu_t, \sigma_t) \} dv \quad (12) \end{aligned}$$

³²We assume that with a zero reserve price the probability of not bidding conditional upon entering the auction is zero. That is, bidders always draw high enough valuations after they've incurred the bid preparation cost to make bidding worthwhile. Correcting for truncation at zero slows down the computation of the likelihood function, and does not change the results discernibly.

where $1\{n_t \geq 1\}$ is an indicator variable for at least one bidder in the auction and \bar{N} is an upper bound on the number of entrants.³³ Observe that this specification allows us to assign positive likelihood to auctions with no bidders.

For auctions with a secret reserve price, the likelihood function in (11) should reflect the fact that bidders assume that seller is another competitor (who is certain to be there) and that they use a different bid function because of this assumption:

$$\begin{aligned} & f(b_1, b_2, \dots, b_n | \mu_t, \sigma_t, \Omega_t) \\ = & \sum_{j=n_t}^{\bar{N}} p(j | \lambda_t) * \int_{-\infty}^{\infty} \{F_x(x^{*SR} | v, \sigma_t, k, \Omega_t)^{j-n_t} (1 - F_x(\phi^{SR}(b_1 | v, \sigma_t, k, \Omega_t))^{1\{n_t \geq 1\}} * \\ & \prod_{i=2}^{n_t} f_b(b_i^{SR} | v, \sigma_t, k, \Omega_t) f(v | \mu_t, \sigma_t)\} dv \end{aligned} \quad (13)$$

The full likelihood function can then be constructed as follows:

$$L(BIDS_t | \beta) = \prod_{t=1}^T f(b_1, b_2, \dots, b_n | \mu_t, \sigma_t, \Omega_t)$$

5.3 Posterior simulation

Our econometric approach will be Bayesian. As Hendricks and Paarsch (1995) point out, many auction models generate likelihood functions that do not have standard asymptotics and therefore large sample approaches to inference are difficult to implement econometrically. We will counter the classical objection to the arbitrariness of Bayesian priors with the following: there are over 1000 observed bids in the data set and the likelihood function will therefore have considerably more weight than the prior in forming the posterior.

For the purposes of this analysis we will impose flat priors with bounded support on the vector of parameters $\{\beta_1, \dots, \beta_{10}\}$. Let

$$X_t = \{SECRET_t, BOOKVAL_t, BLEMISH_t, NEGATIVE_t, MINBID_t\}$$

be the auction-specific data vector and $p(\beta | BIDS_t, X_t)$ denote the posterior distribution of the vector of parameters β conditional upon the observed data. By Bayes theorem:

$$p(\beta | BIDS_t, X_t) \propto p(\beta) L(BID | \beta, X_t)$$

³³With a Poisson process, $\bar{N} = \infty$. But this makes it impossible to account for the truncation in bids due to the reserve price, hence we take \bar{N} to be a large enough number so that the difference between the truncated Poisson distribution and the full distribution is negligible. We found that $\bar{N} = 30$ is good enough for the range of λ 's we consider.

That is, the posterior is proportional to the prior times the likelihood function. In general, we are interested in moments of β or functions of β ; and these can be computed by averaging over the posterior. For example, the expectation of β is given by:

$$\begin{aligned} E[\beta|BIDS_t, X_t] &= \int \beta p(\beta|BIDS_t, X_t) d\beta \\ &= \frac{\int \beta p(\beta) L(BIDS_t|\beta, X_t) d\beta}{\int p(\beta) L(BIDS_t|\beta, X_t) d\beta} \end{aligned}$$

Unfortunately, the integrals in the numerator and denominator are difficult to evaluate analytically except in a very limited class of conjugate prior/likelihood families. Our likelihood function and priors certainly do not fall within this class. However, recent work in Bayesian econometrics has developed methods to simulate draws from the posterior distribution, so that the integrals in the numerator and the denominator can be evaluated using Monte-Carlo simulation. In this research we will use a variant of a Metropolis-Hastings algorithm, which is discussed in Geweke (1996). Details of the algorithm are discussed in Appendix C.

5.4 Discussion.

The econometric approach used in this paper was motivated by two fundamental concerns. First, censoring is present in our data set due to the fact that the minimum bid is not observed, bidders endogenously choose whether or not to participate and some bidders do not submit a bid because the minimum bid is too high. It is not clear how to implement the non-parametric approach taken by Elyakime et al. (1994) and Hendricks et al. (1999) in such settings.

Second, it is not clear how to use the non-parametric approach to directly recover the underlying primitives of the model which are the parameters of the distributions $f(v|\mu_t, \sigma_t)$ and $f_x(x|v)$. In a first price auction, for instance, Hendricks et al. (1999) have noted that the first order conditions can be used to recover the conditional expectation $E[v|x_i = x, y_i = y]$. However, it is important to note that this conditional expectation is a function of the model primitives, $f(v|\mu_t, \sigma_t)$ and $f_x(x|v)$. As the previous discussion illustrates, it may be a far from trivial task in applications to infer $f(v|\mu_t, \sigma_t)$ and $f_x(x|v)$ from knowledge of $E[v|x_i = x, y_i = y]$.

In addition to the approach we have outlined, there may be two other methods to infer $f(v|\mu_t, \sigma_t)$ and $f_x(x|v)$ from knowledge of $E[v|x_i = x, y_i = y]$. The first would rely on carefully chosen functional forms, such as the analysis of Paarsch (1992) or Hong and Shum (1999). We are unaware of any functional forms that would greatly simplify the problem for the model we are considering, however. The second, suggested by Hendricks et al. (1999) would rely on an assumption that it is possible to observe, ex post, the realization of the common value V . In our application, no obvious data source for this information is present.

Table 8: Posterior Means and Standard Deviations of Structural Parameters

Parameter	Mean	Std. Dev.
$\mu : BOOKVAL$	0.9911	0.0387
$\mu : BOOKVAL * BLEMISH$	-0.1446	0.1628
$\sigma : BOOKVAL$	0.5599	0.0260
$\sigma : BOOKVAL * BLEMISH$	0.0326	0.0260
k	0.2545	0.0259
$\lambda : CONSTANT$	1.4069	0.0897
$\lambda : SECRET$	-0.2857	0.1080
$\lambda : LN(BOOK)$	0.0883	0.0276
$\lambda : NEGATIVE$	-0.0414	0.0228
$\lambda : SECRET * \frac{MINBID}{BOOKVAL}$	-0.2514	0.3268
$\lambda : (1 - SECRET) * \frac{MINBID}{BOOKVAL}$	-0.7893	0.0864

The approach to econometric inference in this paper is Bayesian. As Donald and Paarsch (1993) and many others have noted, maximum likelihood estimation may not necessarily be consistent in many auction models. However, the Bayesian procedures we have proposed can be used to simulate the posterior distribution of model parameters in a wide variety of settings.

6 Results

Table 8 reports the results of our simulations of the posterior distribution of β for our structural model. All results are from a posterior simulation of 30,000 draws with an initial “burn-in” of 10,000 draws. Convergence to the posterior appeared to be rapid.

These parameter values translate into the following inferences:

- The information structure: Bidders have a mean valuation that is almost equal to the bookvalue of the item. Reported blemishes do not seem to have a significant effect on this valuation. The standard deviation of the prior distribution on the value of the item is 56% of the bookvalue. A blemish increases this variance by 3%. However, bidders’ signals (conditional on the realization of the common value) are more precise. Since $k = \beta_5 = 0.2545$, the standard deviation of the conditional distribution of bidders’ signals is $\sqrt{k}\sigma_v = 28.25\%$ of the bookvalue.
- Entry: Given equation (8), the reserve price mechanism, minimum bid level in an ordinary auction and the bookvalue of the coin appear as the main determinants of entry to the auction. Negative seller reputation also reduces entry. The minimum bid level in a secret reserve price auction does not affect the entry decision significantly.

To help interpret our estimates, we'll walk through the decision process of a typical bidder on eBay: a potential bidder will first peruse auction listings for a while, selecting the auction(s) in which she might be willing to participate. She will then read the item description and most likely consult a handy reference like the Official Red Book or Coins magazine to get a sense of the book value of the item. This process, we believe, will take an experienced coin collector no more than a minute or so. As a result the bidder will have formed a "prior" on the distribution of the value of the object. We find that the observed bids can be justified by our model of bidding if this "prior" distribution has a mean that is equal the bookvalue of the item and has a standard deviation that is 56% of the book value.

We infer from our estimates that published book values do indeed provide point estimates for how the coins are valued on the market³⁴. There is, however, considerable uncertainty about the value of the coin set about its book value. As we described in the data section, mint/proof sets comprise of uncirculated coins, and any adulteration by the seller would cause the set to lose all its value. Since bidders do not have the opportunity to inspect the items first hand, we believe such an a priori quality uncertainty is plausible³⁵.

After the prior is formed in the bidder's mind, she will decide whether to invest the time to submit a bid. This decision will depend on the amount of profit she expects to make, which, in turn depends on the bookvalue of the item, the number of competitors she expects to face, the reputation of the seller and the minimum-bid/reserve price policy set by the seller, as captured by the specification in equation (8). If she decides that it is worth the while to submit a bid in a particular auction, the bidder will enter what we deem a "costly bid preparation" stage. This is the stage where bidders refine their estimate of the value of the coin, by investigating relevant coin prices and market conditions further, and in almost all cases, augmenting any "objective estimate" by introspection. Of course, the amount of effort sunk into preparing a bid will depend on the stakes involved, and is likely to be different for a \$5 coin and a \$3000 coin.

We find that the refined estimate of bidders is sharper than the "prior," and the uncertainty about the resale value of the object is reduced to about one third of the book value. Given that bid functions for the estimated parameter values are nearly linear and have a slope of about 0.9, this estimate means that bids within an auction should be dispersed with a standard deviation of about 26% of book value.

To gauge the degree of fit our model achieves, we simulated the auctions in our data set using our parameter estimates. That is, given the auction specific variables, we generated n_t signals for each auction t , where n_t was generated by the estimated entry process. We then computed the bid functions corresponding to the signals. The top two plots in Figure 5 compares the simulated and

³⁴The fact that reported blemishes do not seem to change this valuation significantly stems from the fact that most of these "blemishes" were minor details concerning the condition of the container of the coin set, rather than the coins themselves.

³⁵In fact, our "profit" figures in section 3.4 only make sense if there is considerable uncertainty about the value of the object about its book value.

actual bid distributions.

Although the mean value of the bid distributions seem to match (the mean for the simulated distribution is 82.7% of the bookvalue, compared to 83% for the actual data), Figure 5 tells us that our estimate of the variance of cross-auction heterogeneity (β_3 in particular) is higher than what seems apparent in the data. The bottom plot in Figure 5 is the simulated distribution of bids when we set $\beta_3 = 0.3$. This seems to match the actual bid distribution better.

In interpreting the above plots, we should bear in mind that looking at Figure 5 alone does not tell anything about within-auction dispersion of bids. However, our empirical specification attempts to match the *within*-auction dispersion of bids alongside the *cross*-auction dispersion. Figure 6 is a histogram of the simulated and actual dispersion of bids. We find that the mean dispersion (difference between max and min) of the simulated bids is 23% of book value. This is lower than the actual dispersion of bids, which was found to be 32%. Hence we see that in trying to match all moments of the data, our specification has to make a tradeoff between overstating cross-auction heterogeneity and understating within-auction heterogeneity. We believe adding explicit consideration of unobserved heterogeneity into our econometric specification will help achieve a tighter fit; however, such an exercise would greatly complicate the estimation procedure, and is beyond the scope of this paper.

To evaluate our parameter estimates for the entry process, we compare them with the results in Table 3 in section 3.2. The signs of the coefficients on the independent variables seem to be in accord. The difference between the estimates in that regression and parameter estimates for specification (8) is that the dependent variable in Table 3 is the observed number of bidders whereas our structural model specifies the expected number of bidders as the dependent variable. When we predict the expected number of bidders using our estimates, we find that on average 3.3 bidders were expected to show up. Since we allowed for the fact that some bidders may be “screened out” by the minimum bid even after they decide to participate in the auction, it makes sense that the expected number of bidders is higher than the average of the observed number of bidders, which we found to be 3.05.

Using our parameter estimates, we calculated that the expected number of bidders in secret reserve price auctions is, on average, 3.9, as opposed to 3.2 bidders in ordinary auctions. Hence the intuition that secret reserve prices can be used to induce entry is justified. An interesting difference of the parameter results in Table 8 from the results in Table 3 is that entry to secret reserve price auctions is that we find that bidders do react negatively to the existence of a secret reserve price (β_7 is negative and significant, as opposed to the insignificant coefficient on *SECRET* in Table 3). This also makes sense from the standpoint of a rational bidder deciding whether to enter an auction: even though ex-ante expected profit may be higher due to the lower minimum bid in secret reserve price auction, there is still a probability that the highest bidder will not win the auction. Hence entry incentives should be curbed somewhat due to the presence of a secret reserve price.

We should point out that our structural estimates of the entry process are driven entirely by bid data; nowhere in the estimation process do we directly impose the condition that the “expected”

entry process should resemble the actual entry process. The “fit” of the structural model to the observed entry process could be improved by adding such an auxiliary condition; however, even without imposing this condition we manage to get realistic results.

6.1 How big is the winner’s curse?

Recall that the measure of the “winner’s curse” we used in section 3.1 was the amount a bidder would lower her bid when she believed that there would be an extra bidder in the auction. Plots of the bid function in Figure 2 confirm that, at least within the confines of our modelling framework, bid levels should decline with an increase in expected competition.

To assess the size of the winner’s curse in our sample, we computed the bid function in a “representative auction” whose characteristics are set at the sample averages reported in Table 1. The bid functions for the “representative auction” and the auction with one extra expected bidder are plotted in Figure 7. We find that an additional bidder decreases bids by \$1.50 in this \$47 auction. In other words, the “winner’s curse” reduces bids by 3.2% for every additional competitor in the auction.

The bid function in the representative auction is once again very close to being linear and has a slope of 0.9, with an intercept of -0.85. Therefore, with an estimate at the book value of the object (\$47), the bid should be \$41.5. This is \$5.5 less than the bid in a private values auction. Hence we see that “common values” plays an important role in the transactions in this market.

To see how reducing the uncertainty about the value of the object affects the bids, we plot the bid function for the case where the standard deviation of v is only 10% of book value (as opposed to 52%). As expected, the bids increase: for an estimate of \$47, the bid is now \$45. This is about a 7% gain! So the seller definitely has an incentive to reduce the uncertainty about the value of the coin set.

6.2 Cost of entry into an eBay auction

An interesting exercise we can perform using our parameter estimates is to try to calculate the entry or “bid-preparation” cost that would rationalize the stochastic entry process we observe. Such a cost will be borne by serious bidders who expend time and effort into formulating a good estimate of the value of the coin set and finding a profit maximizing bid. This might be a useful fact to know for sellers trying to prepare description pages for their auctions, as giving adequate information about the object for sale would help reduce this cost and induce entry.

To do this we use the framework developed by Levin and Smith (1994) to analyze auctions with endogenous entry. In this framework, n out of N bidders decide to enter an auction *before* they observe their private signal. Contingent on entry, each bidder bears a bid-preparation cost c . Levin and Smith (1994) find that there exists a symmetric equilibrium of the endogenous entry game in

which each bidder enters the auction with an identical probability p . Observe that if this probability is low enough and the number of potential bidders is high enough – conditions likely to be satisfied on eBay – then the Poisson entry process we assumed above is justified.

What determines the entry probability? Levin and Smith (1994) argue that entry will occur until each bidder’s ex-ante profit from entering the auction is zero. More specifically, let V be the ex-ante expected value of the item, and W_n be the ex-ante expected payment the winning bidder will make to the seller. Then a bidder’s expected profit, conditional on entering an auction with n bidders is $(V - W_n)/n - c$, since all bidders are ex-ante symmetric. Define $T_n(r)$ to be the probability of trade given n and minimum bid (or reserve price) r , and p_n^i be the probability that there are n bidders in the auction from the perspective of the i th potential bidder, conditional on herself being one of the bidders. Assuming all other bidders are following the symmetric entry probability, the following zero-profit condition applies to the i th bidder:

$$c = \sum_{n=1}^{\infty} \frac{e^{-\lambda} \lambda^{n-1}}{(n-1)!} T_n(r) \frac{(V - W_n)}{n} \quad (14)$$

where, assuming a Poisson approximation, we computed $p_n^i = \frac{e^{-\lambda} \lambda^{n-1}}{(n-1)!}$, $n = 1, \dots, \infty$.

We can get a feel for the entry cost to the eBay auctions by computing the above expression for the “representative auction” we considered in the previous section. One complication is that W_n and $T_n(r)$ are ex-ante expected values that are not available in closed form. Therefore we use the following simulation method:

1. Using a Poisson(λ) distribution, draw K realizations of $\{n_1, \dots, n_K\}$, the number of bidders in a series of K auctions.
2. Given n_k , draw n_k signals $\{x_1, \dots, x_{n_k}\}$ from the conditional distribution $f(x|v)$, whose parameters were estimated previously.
3. Given these signals, compute the bid corresponding to the second highest signal. This will, in the case where the reserve price r is exceeded, give the payment $W_n^{(k)}$ for the k th auction in our simulation.
4. Compute the average $\frac{1}{K} \sum \frac{(V - W_n^{(k)})}{n_k}$ to get an estimate for c , the cost of entry.

Performing this simulation procedure using 20000 draws, we get that the entry cost implied in the “representative auction” is \$3.20³⁶. This is about 6.8% of the bookvalue (\$47) in this auction and it is comparable to the shipping and handling cost.

This “estimate” of the entry cost will change with auction characteristics. However, the entry cost as a percentage of the bookvalue seems very stable. For example, for an auction with bookvalue

³⁶The number of simulation draws was enough to ensure that the standard deviation of this number was less than 0.05.

\$1000 and a minimum bid of \$600, the computed entry cost was found to be \$66, which is 6.6% of the bookvalue.

A 6-7% entry cost can be quite significant if the book value of the item is high. Therefore, it might be advisable for sellers of high value items to try and reduce this cost by giving out as much information about the item as possible, by providing pictures, certificates of authenticity, or a report from an official coin grading agency like PCGS (Professional Coin Grading Service)³⁷.

6.3 Reserve Price Policy

What is the optimal minimum bid (posted reserve price) given the seller's value for the object? Should one set a secret reserve price or not? In this section, we use our structural estimates to answer these questions, which might be of some practical interest to the sellers on eBay.

To find the optimal minimum bid for a given reservation value, we simulate the expected revenue for different choices of the minimum bid. We assume the seller will retain her reservation value as revenue if there is no sale.

Figure 8 reports the revenue simulations for an item with a bookvalue of \$47 (sample average) and a reservation value equal to the bookvalue. We observe that the seller maximizes her expected revenue by setting the minimum bid 10% higher than her reservation value. To check whether this is a robust markup rule we repeated the same experiment with reservation values equalling 50% and 75% of bookvalue, and found a similar 10% markup.

Figure 8 also plots the expected revenue for the case when the seller instead sets a secret reserve price. Once again, the 10% markup over the seller's reservation value becomes apparent.

Setting a minimum bid that is higher than the reservation value is a result that we are familiar with from the private values auction literature. Another prediction of the reservation price rule in private value auctions is that the amount of "markup" should not depend on the number of bidders. Since our parameter estimates tell us that increasing the minimum bid increases the expected number of bidders, our simulations also reveal a markup rule that is not very sensitive to the expected number of bidders in the auction.

In the data we see that the average minimum bid in auctions with no secret reserve prices is 70% of the book value. This can be explained in two ways that are consistent with seller revenue maximization: either the sellers' reservation values for the coins are significantly lower than the book value of the item, or the sellers are averse to a "no-sale" situation.

We should also remember that sellers of identical or close-substitutes in this market (for which there is limited demand) are in direct competition. If sellers have identical reservation values and one of them sets a minimum bid level that is 10% above her reservation value, then by the basic Bertrand argument, another seller might find it advantageous to lower her minimum bid. Unfortunately,

³⁷PCGS charges \$30 for its regular grading service with 21 days of turnaround time.

studying competition among auctioneers is beyond the scope of this paper. For a theoretical model that covers most private value auction formats, please see Peters and Severinov (1997) and McAfee (1993).

Should a seller make her reserve price public or should she keep it secret? Figure 8 shows that for the optimal reserve price, a secret reserve will outperform a posted reserve. We can get some intuition for this result from our entry parameter estimates in Table 8: a given reservation price level results in more bids in the auction if it is kept secret as opposed to being made public. Setting a secret reserve price (and setting the minimum bid to zero) will lower the expected number of bidders by 0.28, but setting a minimum bid equal to the bookvalue will chase away 0.78 bidders.

This intuition is not completely airtight, however. Having more bidders in an auction might not necessarily be good for revenue. As we pointed out in earlier sections, the existence of more bidders in the auction increases the magnitude of the winner's curse, and hence might reduce the bid levels. The benefit of having more bidders is to have a larger set of random numbers to draw from: the expected maximum of 3 draws is higher than the expected maximum of 2 draws.

Bulow and Klemperer (1996) provides an unambiguous resolution to the above tradeoff in the case of common value English auctions: having $N + 1$ bidders in an auction is always better (in terms of expected revenue) than having N bidders. Unfortunately, in our case, we have a second-price auction as opposed to an English auction, and the improvement in the number of bidders is not an integer. Therefore, instead of seeking an analytical result, we will use a simulation framework similar to the one used in the previous section to evaluate the use of different reserve price policies.

In the following simulation we calculate a seller's ex-ante expected revenue from using a secret vs. a posted reserve price scheme. We assume that the seller's reservation value is equal to the bookvalue. In both simulations, the seller sets the reserve price at the revenue maximizing reserve price level (10% above book); but in the first simulation, she keeps the reserve price secret (with a starting bid of zero), and in the second simulation, she announces it.

Figure 9 plots the difference between expected revenues in the two simulations. On the horizontal axis, we vary the bookvalue of the item, the vertical axis has the expected revenue difference. Other auction characteristics are set equal to the sample averages. We see that the expected revenue difference does not exceed a dollar for items with bookvalues below 50 dollars, and that the expected revenue gain in general rises with the bookvalue of the item. Hence, we can conclude from this simulation and from our data that it might not be worthwhile to run a secret reserve price auction for low bookvalue items and that the incentive to do so increases with the value of the item.³⁸

In light of these simulations, our recommendation to sellers is to use a minimum bid about 10% higher than their reservation value to maximize their revenue. Also, if the item has high enough value (above \$50) a secret reserve price auction with a zero minimum bid can induce additional entry and boost revenue. The latter strategy has indeed been discovered by eBay users and even

³⁸This is especially true now that eBay has decided to charge sellers \$1 extra to list a secret reserve auction.

Kaiser and Kaiser (1999) recommends sellers to set a low minimum bid when they are running a secret reserve price auction.

6.4 Does the optimal mechanism entail an entry fee?

Should sellers on eBay charge bidders to participate in their auctions? At first glance, this question may sound absurd, as an entry fee seems only to exacerbate the cost of entry, resulting in even fewer bidders. However, as originally pointed out by Levin and Smith (1994), having fewer bidders in an auction might actually be more desirable for sellers, since the winner's curse causes each additional bidder to reduce the bids in an eBay type common value auction.

In the presence of an entry fee e , the zero-profit condition characterizing entry in Section 6.2 becomes:

$$e + c = \sum_{n=1}^{\infty} \frac{e^{-\lambda} \lambda^{n-1}}{(n-1)!} T_n(r) \frac{(V - W_n)}{n} \quad (15)$$

Since entry occurs to the point that each bidder makes zero ex-ante expected profit,³⁹ the seller extracts all of the surplus realized from the sale of the object, minus the cost of entry borne by the bidders. Therefore, the ex-ante expected revenue of the seller becomes:

$$ER = \sum_{n=0}^{\infty} \frac{e^{-\lambda} \lambda^n}{n!} T_n(r) V - \lambda c$$

We'll analyze the case where the seller's value for the object is zero.⁴⁰ In this case, setting a positive reserve price prevents efficient trades from happening, so the reserve price will be equal to zero. Then $T_n = 1$, and the first-order condition for the optimal amount entry is given by:

$$\frac{\partial ER}{\partial \lambda} = e^{-\lambda} V - c$$

Observe that this first order condition gives a simple formula for the amount of optimal entry, λ^* :

$$\lambda^* = \log \frac{V}{c}$$

³⁹We're ignoring the fact that the number of bidders is an integer and that profits may be positive due to this restriction.

⁴⁰Unfortunately the analysis becomes much less tractable when the seller has a non-zero value for the object. Since the aim of this section is mainly to point out the intuition behind a non-zero entry fee, and do some simple calculations to gauge its magnitude, we analyze the zero seller value case and leave the detailed analysis of the general case for further work.

Substituting for c in the zero profit condition, we can also get an expression for the entry fee that achieves the optimal amount of entry:

$$e = \sum_{n=2}^{\infty} \frac{e^{-\lambda} \lambda^{n-1}}{(n-1)!} \frac{V - W_n}{n}$$

Using a similar simulation method as in the last section, we calculated the optimal entry fee for different book value coins, fixing entry and informational parameters at the sample averages, and assuming that the seller has zero value from keeping the object. Figure 10 plots the simulated values of the entry fee for different book value items. We see that the optimal entry fee is about 3% of book value; which makes sense when we consider our earlier finding that the presence of one additional, “unwanted” bidder in the auction reduces bids by 3.2% due to the winner’s curse – since the entry fee is imposed to deter the extra bidder from entering the auction.

We think imposing entry fees will probably not be a very realistic move in auctions for items with mundane values, since, with many competing auctions, the incentive for entry will greatly diminish. However, for items with very high book values, the winner’s curse maybe strong enough to justify imposing entry fees to restrict entry to the auction. Hence, we do not prescribe that eBay impose entry fees, but it should at least allow sellers to experiment with them.

7 Conclusion

In this paper, we have analyzed a data set from eBay coin auctions to investigate the determinants of bidder and seller behavior. Our main empirical findings are the following:

1. The eBay coin auctions that we analyze fit a common values framework better than a private values framework. Observed bidder behavior departs significantly from what would be expected from a private values auction. However, eBay’s mechanism design philosophy seems to rely on the private values assumption. Since bidder and seller behavior can vary significantly depending on whether private or common values are at play, we would caution eBay administrators in forming their site and mechanism design policies.
2. Our structural estimation results allowed us to quantify the “winner’s curse effect” associated with the existence of common values. We find that the expectation of one additional bidder auction decreases bids by 3.2% in a representative auction. Also, in this representative auction, bids are 10% lower than bidders’ estimate of the value of the coin set.
3. Entry into an auction seems to play a very important role in determining revenues. The entry patterns we observe seem to be in accord with a setting in which forming an intelligent bid, after observing auction listings, is a costly activity. Hence, bidders enter auctions in which

they are likely to be successful. Since a high minimum bid reduces the ex-ante expected profit of a potential bidder, it reduces the incentive to spend time and effort to enter the auction. Hence minimum bids are the most significant determinant of entry into the auction. We also find that higher book value coins attract more bidders, most likely because the dollar value of expected profits increase.

4. Since entry depends crucially on the minimum bid level, one strategy for the seller is to set a low minimum bid and a secret reserve price. Although bidders do account for the fact that there is a secret reserve price, we find that the sellers using this strategy are able to increase entry to an auction and still manage to guarantee a high selling price. We find that a secret reserve price can yield higher expected revenue to the seller. We also find that the secret reserve price strategy should be reserved for higher book value items, since eBay charges a fee to use a secret reserve price.
5. We also find that charging bidders an entry fee might be an optimal strategy if there is potential for too much entry to the auction. We found that if the seller has no value for the good, then she should charge an entry fee equal to 3% of the book value of the item.

Along the way, we have also made methodological contributions to the empirical auction literature. By devising an efficient method to compute the bid function for a common value auction, we have managed to reduce the computational requirements of our estimation procedure. Since this method is easily generalizable to other auction models, we believe future researchers can benefit from our findings when building and estimating parametric structural models of auctions.

We believe that eBay remains a fertile ground for empirical auction research. For example, we have not explicitly taken into account the “multi-market” nature of eBay, and have treated each auction in isolation. As we mentioned, in the market we studied, very often identical or very similar goods were being sold side-by-side or back-to-back. We have also realized that many sellers have more than one object to sell. Hence eBay provides an ideal testing ground for the exploration of the impact of “temporal” or “spatial” bidder and seller strategies in multi-market settings⁴¹. Since we also observe significant variation among seller’s choices of minimum bids and reserve prices, there is also room to test theories of competition between sellers (Peters and Severinov (1997), McAfee (1993)).

We also believe possible extensions to our work can be made in the setting we have examined. For example, in our structural specification, we have not modeled the sequential nature of entry and have assumed that the entry decision is made simultaneously. Relaxing this assumption could open the way to dynamic, rather than static, bidding strategies that can be taken to the data. We have also assumed away any consideration of asymmetries between bidders or errors in decision

⁴¹For example, the per unit shipping and handling costs for a bidder declines if she wins multiple auctions from the same seller. Hence her bidding strategy across different auctions might not completely separable.

making, although asymmetries and error-prone decision-making can alter the empirical implications of auction models significantly.

A Appendix

A.1 Proof of proposition 3

Without loss of generality, focus on the decision of bidder 1. Suppose the bidders $j \neq i$ adopt strategies $b_j = b^*(x_j, \lambda)$. Then the highest bid among them, conditional on there being n bidders, will be $W^{(n)} = \max_{2 \leq j \leq n} b_j(x_j)$. Since b^* is increasing in x , $W^{(n)} = b^*(y_1^{(n)})$ and the expected payoff of bidder 1 from bidding b will be

$$E_n[E[(v - b^*(y_1^{(n)}))\chi_{\{b^*(y_1^{(n)}) < b\}} | x_1 = x, n \geq 2]]Prob(n \geq 2) + E[v | x_1 = x]Prob(n = 1)$$

where χ is the indicator variable for the event that bidder 1 wins. This is equal to

$$\begin{aligned} & E_n\left[\int_{\underline{x}}^{b^{*-1}(b)} [v(x, \alpha, n) - b^*(\alpha, \lambda)]f_{y_1^{(n)}}(\alpha|x)d\alpha | n \geq 2\right]Prob(n \geq 2) + E[v | x_1 = x]Prob(n = 1) \\ &= \int_{\underline{x}}^{b^{*-1}(b)} \left\{ \sum_{n=2}^{\bar{N}} p(n, \lambda)v(x, \alpha, n)f_{y_1^{(n)}}(\alpha|x) - b^*(\alpha, \lambda) \sum_{n=2}^{\bar{N}} p(n, \lambda)f_{y_1^{(n)}}(\alpha|x) \right\} d\alpha Prob(n \geq 2) \\ & \quad + E[v | x_1 = x]Prob(n = 1) \end{aligned}$$

where

$$\begin{aligned} f_{y_i^{(n)}}(x|x) &= \Pr(y_i^{(n)} = x | x_i = x) \\ &= \int_{\underline{v}}^v f_{y_i^{(n)}}(x|v)f_v(v|x)dv \\ &= \int_{\underline{v}}^v (n-1)f_x(x|v)F_x^{n-2}(x|v)f_v(v|x)dv \end{aligned}$$

and $f_v(v|x)$ is the posterior density of the common value given the signal.

The first order condition with respect to b yields:

$$0 = \sum_{n=2}^{\bar{N}} p(n, \lambda)v(x, x, n)f_{y_1^{(n)}}(x|x) - b^*(x, \lambda) \sum_{n=2}^{\bar{N}} p(n, \lambda)f_{y_1^{(n)}}(x|x)$$

solving, we get $b^*(x, \lambda)$ as above.

A.2 Proof of proposition 4

The easiest way to see why this should hold true is to look at the von Neumann-Morgenstern (vNM) utility of a particular bidder i (the payoff given that other bidders submit bids that are monotonic in their signals, x_{-i} , and conditional on the number of competitors n):

$$\begin{aligned} u_i(b_i, b_{-i}, x_i, x_{-i}, n) &= \frac{k\mu + \sum_{j=1}^n x_j}{n+k} - \max\{b_{-i}(x_{-i})\} \quad \text{if } b_i > \max\{b_{-i}(x_{-i})\} \\ &= 0 \quad \text{otherwise} \end{aligned}$$

The term $\frac{k\mu + \sum_{j=1}^n x_j}{n+k}$ is the posterior mean of the unknown common value, given all bidders' signals. The objective function that bidder i maximizes is the expectation of the above vNM utility, and the bid he submits is the maximizer.

$$\begin{aligned} b_i &= \arg \max E_n E_{x_{-i}} u_i(b_i, b_{-i}, x_i, x_{-i}, n) \\ &= \arg \max E_n E_{x_{-i} | b_i > \max\{b_{-i}(x_{-i})\}} \left[\frac{k\mu + x_i + \sum_{-i} x_{-i}}{n+k} - \max\{b_{-i}(x_{-i})\} \right] \end{aligned} \quad (16)$$

The bid that maximizes bidder i 's expected utility in an auction where the signals are scaled linearly as $r_1 x + r_2$, and where all other bidders follow linearly scaled strategies $r_1 b_{-i}(x_{-i}) + r_2$ is:

$$\begin{aligned} b'_i &= \arg \max E_n E_{r_1 x_{-i} + r_2} u_i(b'_i, r_1 b_{-i}(x_{-i}) + r_2, r_1 x_i + r_2, r_1 x_{-i} + r_2, n) \\ &= \arg \max E_n E_{r_1 x_{-i} + r_2 | b'_i > \max\{r_1 b_{-i}(x_{-i}) + r_2\}} r_1 \left[\frac{k\mu + x_i + \sum_{-i} x_{-i}}{n+k} - \max\{b_{-i}(x_{-i})\} \right] \\ &= \arg \max E_n E_{r_1 x_{-i} + r_2 | b'_i > \max\{r_1 b_{-i}(x_{-i}) + r_2\}} \left[\frac{k\mu + x_i + \sum_{-i} x_{-i}}{n+k} - \max\{b_{-i}(x_{-i})\} \right] \\ &= \arg \max E_n E_{x_{-i} | \frac{b'_i - r_2}{r_1} > \max\{b_{-i}(x_{-i})\}} E \left[\frac{k\mu + x_i + \sum_{-i} x_{-i}}{n+k} - \max\{b_{-i}(x_{-i})\} \right] \end{aligned}$$

Comparing the last expression with the last line in (15), we see that $\frac{b'_i - r_2}{r_1} = b_i$, i.e. $b'_i = r_1 b_i + r_2$. Hence the bids follow the same linear transformation that the signals are scaled with.

B Numerical computation of the inverse bid function

Likelihood function evaluations require that we find the value of the inverse bid function ϕ for each bid in auction t . Since a closed form expression for the inverse bid function does not exist, this potentially means that we will have to evaluate the bid function $b(x|\mu_t, \sigma_t, k, \lambda_t)$ at least a few times for each bid to find the inverse of $b(\cdot)$ using a numerical algorithm like Newton-Raphson. However, as we found out while working out the closed forms for equilibrium bid functions, this requires the numerical evaluation of integrals. This is a costly step since we have to evaluate the full-likelihood about 30,000 times in our posterior simulation, and with 1500 bids in the data, the integrals would have to be evaluated at least 45 million times.

To reduce the computational complexity, we first exploit the linearity property of the bid functions. Recall that linearity implies that if $b(x|\mu^\circ, \sigma^\circ, k, \lambda)$ is a pure strategy equilibrium bid function to the auction where the common value is distributed normally with mean μ° and standard deviation σ° , then $r_1 b(r_1 x + r_2 | \mu^\circ, \sigma^\circ, k, \lambda) + r_2$ is the equilibrium bid function in the auction where the common value is distributed with mean $r_1 \mu^\circ + r_2$ and standard deviation $r_1 \sigma^\circ$. This result holds regardless of whether there is a known or random number of bidders in the auction. Similarly, if there is a posted reserve price in the auction, we can apply the linear transformation to the screening level $x^*(R, \lambda, k)$. That is, if $x^*(R; \mu^\circ, \sigma^\circ, \lambda, k)$ is the screening level in the original auction, then $r_1 x^*(r_1 R + r; |\mu^\circ, \sigma^\circ, k, \lambda) + r_2$ will be the screening level in the transformed auction.

How does linearity help us? Suppose at a particular evaluation of the likelihood function, the distributional parameters are $\{\mu_t, \sigma_t\}$ for auction t . Then we do not have to evaluate $b(x|\mu_t, \sigma_t, k, \lambda)$: setting $r_{1t} = \frac{\sigma_t}{\sigma^o}$ and $r_{2t} = \mu_t - r_{1t}\mu^o$, we interpolate the value of $b(\frac{x-r_{2t}}{r_{1t}}|\mu^o, \sigma^o, k, \lambda)$ from a set of pre-computed values, and return $r_{1t}b(\frac{x-r_{2t}}{r_{1t}}|\mu^o, \sigma^o, k, \lambda) + r_{2t}$ ⁴². Assuming the interpolation is good enough, this scheme lets us avoid computing the Nash equilibrium in the course of likelihood function evaluations. The equilibrium is solved beforehand for a normalized set of parameter values, and adapted to the parameter values generated by the posterior simulation routine.

As for inverting the bid function, one can use a simple algorithm like Newton-Raphson with relative ease, since the bid function is monotonic, and $b(x) = x$ is usually a good starting value. Alternatively, one can evaluate $b(x|\mu^o, \sigma^o, k, \lambda)$ on the grid used for interpolation, and instead of using polynomial interpolation to find b as a function of x, k, λ , one can use polynomial interpolation to evaluate x as a function of b, k, λ . This allows one to approximate the inverse bid function directly, instead of having to evaluate the bid function several times until Newton-Raphson converges. In fact, this method is much more efficient in terms of calculating the derivative of the inverse bid function (the Jacobian), since the derivative of the polynomial interpolation can be written out in closed form, and once again, be evaluated in one step.

C Estimation

C.1 Priors

Let $p(\beta)$ denote the prior distribution on the parameters. We assume that the prior is proportional to a constant if

$$\begin{aligned} 0 &< \beta_1 < 3 \\ -2 &< \beta_2 < 2 \\ 0 &< \beta_3 < 2 \\ -2 &< \beta_4 < 2 \\ 0 &< \beta_5 < 4 \\ 0 &< \beta_{6,7,8,9,10} < 5 \end{aligned}$$

and is equal to zero otherwise. We impose these flat priors to minimize the effect of prior choice

⁴²The computational simplification from linearity can be exploited fully only when we have precomputed values for all possible values of x, k and λ . Since this is not possible in a computational setting, we precompute the bid function for the "base-case" auction (where the CV is distributed $N(\mu^o, \sigma^o)$) at a large number of grid points $x \in \{x_o, x_1, \dots, x_{N_k}\}, k \in \{k_o, k_1, \dots, k_{N_k}\}, \lambda \in \{\lambda_o, \dots, \lambda_{N_\lambda}\}$ and perform a polynomial interpolation to calculate $b(\frac{x-r_{2t}}{r_{1t}}|\mu^o, \sigma^o, k, \lambda)$ at intermediate values. We took care to restrict the domain of interpolation to a region that is a rather tight cover of the region covered by simulation draws to ensure good performance.

on the shape of the posterior distribution. As for the appropriateness of the cutoff points, we made sure that these contained the "reasonable bounds" for the possible realizations of the parameters.

C.2 Posterior simulation

In this section we give an outline of the variant of the Metropolis-Hastings algorithm that we used to simulate the posterior distribution of the parameters. The goal of this algorithm is to draw a sequence of random vectors $\beta^0, \beta^1, \beta^2 \dots$ from the posterior distribution. The algorithm starts out by finding a rough approximation to the posterior mode of which we will refer to as β^o . Let H be the Hessian matrix of the posterior distribution at β^o . The algorithm starts with the first element in the sequence as β^o and is defined inductively as follows:

1. Given β^n draw a candidate value $\tilde{\beta}$ from a normal distribution with mean β^n and variance-covariance matrix $-H^{-1}$.
2. Let

$$\alpha = \min\left\{\frac{p(\tilde{\beta}|DATA)}{p(\beta^n|DATA)}, 1\right\}.$$

3. Set $\beta^{n+1} = \tilde{\beta}$ with probability α .
4. Return to 1.

Intuitively, the algorithm simulates a Markov chain (or a random walk) starting from the posterior mode, which under certain stationarity conditions, converges to an invariant distribution. Geweke (1996) provides sufficient conditions for the invariant distribution of this chain to be the posterior distribution. It is easy to check that these conditions are met in this application.

References

- Bajari, Patrick**, “The First Price Auction with Asymmetric Bidders: Theory and Applications.” PhD dissertation, University of Minnesota 1997.
- Bikhchandani, Sushil and Chifu Huang**, “Auctions with Resale Markets: An Exploratory Model of Treasury Bill Markets,” *Review of Financial Studies*, 1989, 3, 311–339.
- Bulow, Jeremy and Paul Klemperer**, “Auctions vs. Negotiations,” *American Economic Review*, 1996, 86, 180–194.
- Donald, Stephen G. and Harry J. Paarsch**, “Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions,” *International Economic Review*, 1993, 34, 121–148.
- Elyakime, Bernard, Jean Jacques Laffont, Patrice Loisel, and Quang Vuong**, “First-Price Sealed-Bid Auctions with Secret Reservation Prices,” *Annales D’Economie Et De Statistique*, 1994, 34, 115–141.
- Geweke, John**, “Posterior Simulators in Econometrics,” 1996. University of Minnesota.
- Haile, Philip**, “Auctions with Resale Markets.” PhD dissertation, Northwestern University 1997.
- , “Auctions with Private Uncertainty and Resale Opportunities,” 1999. University of Wisconsin.
- , “Partial Pooling at the Reserve Price in Auctions with Resale Opportunities,” 1999. University of Wisconsin.
- Harstad, Ronald M.**, “Alternative Common Value Auction Procedures: Revenue Comparisons with Free Entry,” *Journal of Political Economy*, 1990, 98, 421–429.
- Hendricks, Kenneth and Harry Paarsch**, “A Survey of Recent Empirical Work Concerning Auctions,” *Canadian Journal of Economics*, 1995, 28, 403–426.
- , **Joris Pinske, and Robert H. Porter**, “Empirical Implications of Bidding in First-Price, Symmetric, Common Value Auctions,” 1999. Northwestern University.
- Hong, Han and Matthew Shum**, “The Econometrics of English Auctions,” 1997.
- Judd, Kenneth**, *Numerical Methods in Economics*, Cambridge, Massachusetts: MIT Press, 1998.
- Kaiser, Laura Fisher and Michael Kaiser**, *The Official eBay Guide*, New York, New York: Fireside, 1999.
- Laffont, Jean Jacques, H. Ossard, and Q. Vuong**, “Econometrics of First-Price Auctions,” *Econometrica*, 1995.

- Levin, Dan and James L. Smith**, “Equilibrium in Auctions with Entry,” *American Economic Review*, 1994, *84*, 585–599.
- McAfee, R. Preston**, “Mechanism design by competing sellers,” *Econometrica*, 1993, *61*, 1281–1312.
- Milgrom, Paul R. and Robert J. Weber**, “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 1982, *50*, 1089–1122.
- Myerson, Roger B.**, “Optimal Auction Design,” *Mathematics of Operation Research*, 1981, *6*, 58–73.
- Paarsch, Harry J.**, “Deciding Between Common Values and Private Value Paradigms in Empirical Models of Auctions,” *Journal of Econometrics*, 1992, *51*, 191–215.
- Peters, Michael and Sergei Severinov**, “Competition among sellers who offer auctions instead of prices,” *Journal of Economic Theory*, 1997, *75*, 141–179.
- Reiley, David Lucking**, “Auctions on the Internet: What’s Being Auctioned, and How,” 1999.
- Stroud, A. H. and Don Secrest**, *Gaussian Quadrature Formulas*, Englewood Cliffs, New Jersey: Prentice Hall, 1968.
- Vickrey, William**, “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *Journal of Finance*, 1961, *16*, 8–37.
- Wilson, Robert**, “A Bidding Model of Perfect Competition,” *Review of Economic Studies*, 1977, *44*, 511–518.
- , “Sequential Equilibria of Asymmetric Ascending Auctions,” *Economic Theory*, 1995.

Figure 1:




Figure 2:




Figure 3:




Figure 4:




Figure 5:




Figure 6:




Figure 7:

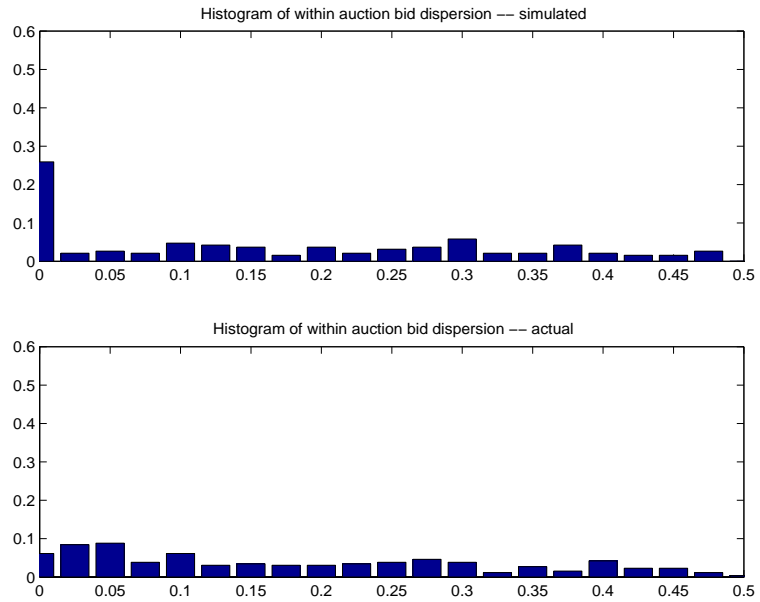


Figure 8: Expected revenue vs. reserve price




Figure 9: Expected revenue difference between secret and posted reserve schemes




Figure 10: Optimal entry fee for different bookvalue coins

