

Public Versus Secret Reserve Prices in eBay Auctions: Results of a Pokémon Field Experiment

Rama Katkar and David Lucking-Reiley¹

July 19, 2000

Abstract

Sellers in eBay auctions have the opportunity to choose both a public minimum bid amount and a secret reserve price. We ask, empirically, whether the seller is made better or worse off by setting a secret reserve above a low minimum bid, versus the option of making the reserve public by using it as the minimum bid level. In a field experiment, we auction 50 matched pairs of Pokémon cards on eBay, half with secret reserves and half with equivalently high public minimum bids. We find that secret reserve prices make us worse off as sellers, by reducing the probability of the auction resulting in a sale, deterring serious bidders from entering the auction, and lowering the expected transaction price of the auction.

¹ Katkar: Northwestern University, email r-katkar@nwu.edu. Lucking-Reiley: Vanderbilt University, email reiley@vanderbilt.edu. An earlier version of this paper was Katkar's senior honors thesis in Mathematical Methods in the Social Sciences; we are grateful for the financial support of an Undergraduate Research Grant from the Weinberg College of Arts and Sciences at Northwestern University, enabling us to purchase the cards auctioned in this experiment.

I. Introduction

Online-auction services such as eBay allow sellers to specify a number of different parameters when listing an item for auction. Among these are the number of days the auction will take place, the level of the opening bid, and the amount of a secret “reserve price” below which the seller will not sell the item. In this paper, we investigate the effects of setting a reserve price higher the opening bid amount. Does the use of a reserve price increase seller revenues relative to the use of a public minimum bid with no reserve?

As shown previously by Lucking-Reiley (1999a, 1999b), online commerce presents economists with exciting opportunities to conduct field experiments. Rather than waiting passively for firms and consumers to generate data that may or may not contain the exogenous variation required to test a theory, the researcher can participate actively in a market to conduct a controlled experiment specifically designed to answer the question of interest. In this paper, the question is whether the use of a secret reserve price affects seller revenues and/or the probability of selling a card . Our experiment involves auctioning 50 matched pairs of Pokémon cards on the eBay Web site. Within each pair, we auctioned one card with a nontrivial minimum bid X but no reserve, and the other card with a trivial opening bid but a reserve price of X . By carefully controlling the experiment in this manner, we can isolate the effect of the secret reserve.

Field data on online auctions are quite plentiful, with publicly available information on hundreds of thousands of auctions closing each day on eBay alone.²

² See Lucking-Reiley (2000a) for a survey of online-auction institutions as of 1998-99. One of the most important developments in online auctions since the conclusion of that survey has been the introduction of online payment services such as Paypal and Billpoint, which make it possible for individual sellers to

Indeed, a variety of authors have begun to exploit this rich source of data to investigate economic questions. For example, Lucking-Reiley *et al* (2000) and Hauser and Wooders (2000) examine the effect of eBay “feedback ratings” on the final auction price. Roth and Ockenfels (2000) show that the auction’s closing rule has significant effect on bidders’ strategic timing of bids, while Wilcox (1999) focuses on how the timing of bids varies with bidders’ experience. Easley and Tenorio (1999) quantify the amount of jump bidding present in online retail auctions . However, we note here that some questions cannot be easily answered with existing field data, no matter how vast the quantity of data may be.

In particular, the question of secret reserve prices is quite difficult to study with data on eBay auctions run by other sellers, for several reasons. First, it can be very difficult to find two auctions where everything is held constant except the use of a reserve price. Even if one tries to find two auctions for exactly the same good, the auctions will likely differ slightly in characteristics (two cards may not be in exactly the same condition, be described somewhat differently, have different shipping costs, involve very different seller feedback ratings, etc.). Second, even if one found a sample of auctions with minimal noise in these extraneous variables (and minimal possibility for omitted-variable bias), one still cannot collect perfect information on the main variable of interest. When a seller uses a secret reserve price, an outside observer can never know for sure the exact amount of the reserve. One can only observe a lower bound on the reserve price in an auction where the reserve was not met, and an upper bound in an auction where the reserve price was met; only the seller knows for sure what the reserve price was.

accept credit-card payments without having to obtain a formal merchant agreement. This speeds up payments relative to the alternative system of having the buyer mail a check to the seller.

Admittedly, heroic econometrics and painstaking data collection might conceivably allow one to tease an empirical measurement of the reserve-price effect out of eBay field data. Indeed, Bajari and Hortag̈su (2000) present an indirect measurement of the effects of reserve prices in eBay auctions, as part of a larger effort to provide a structural econometric model of bidding to answer several other questions. Their complex model involves some strong econometric assumptions, such as the assumption that the unobserved reserve prices are set as if they were bids from an independent bidder. By contrast, our approach carries out a single empirical measurement, that of the effect of the secret reserve price, and seeks to measure it in a manner that is a simple, direct, and assumption-free as possible. We not only observe the levels of reserve prices, but control them (and the other attributes of our auctions), helping us to establish conclusive empirical evidence even from a relatively small data set (100 observations). Results from such experiments may also prove useful as building blocks for researchers trying to build structural models of market behavior.

The remainder of this paper is organized as follows. Section two considers the arguments for and against the use of secret reserve prices by sellers, both in the academic literature and in discussions by actual sellers. Section three provides details of our experimental design, including the relevant eBay institutional details. Section four presents our empirical results, and section five concludes.

II. Secret Reserve Prices

Secret reserve prices have been used in auctions for many years. In a brick-and-mortar auction house such as Sotheby's, the bids continue increasing until the point when no bidder is willing to raise the current bid higher. But if this bid amount does not exceed a reserve price that may be specified in advance by the seller, the auctioneer will refuse to "hammer down" the good, and it will not be sold to the highest bidder. No bidder knows in advance the amount of the secret reserve, and in fact no one knows for sure whether there is a reserve price at all. Ebay similarly allows sellers to keep reserve-price amounts secret, but they do inform bidders whether or not a reserve price is in effect. When the seller specifies a secret reserve price, the auction begins at the opening bid amount with a public indication on its Web page that "the reserve price has not yet been met." Though eBay's computer accepts no bids at amounts less than the public minimum bid, it does accept bids less than the unknown secret reserve. As bidding proceeds, the current high bidder's identity and bid amount are updated, and if the reserve price is finally exceeded, the reserve price messages changes to "the reserve price has been met."³ Auctions without reserve prices have no such message, so that bidders know in such a case that a sale will definitely occur at the high bid price.

Why would sellers want to use secret reserve prices? A reserve can increase a seller's expected profit by raising the winner's bid (as if the reserve were a more aggressive second-highest bidder), even though it may sometimes cause the good to go

³ One other effect of the reserve price shows up in the use of "proxy bidding" on eBay (see Lucking-Reiley (2000b)). The proxy-bidding system keeps secret the highest amount actually submitted by the high bidder, and instead makes the current high bid equal to one increment over the amount of the second-highest bid. The exception is when a reserve price is involved. If the current high bid is below the reserve price and a bidder submits a new amount that happens to exceed the reserve price, the high bid becomes one increment above the reserve (as if the reserve were a bid).

unsold.⁴ However, this explanation is as true for announced reserve prices (i.e., public minimum bids) as it is for secret ones; it begs the question of why the seller might choose to make her reserve price amount secret.

The most common argument in favor of a secret reserve appears to be that a high public minimum bid tends to scare away potential bidders, which may result in the good not being sold at all. By contrast, a low opening bid (with a high secret reserve price) can grease the wheels of bidding, building up bidding “momentum” that can propel the price past the amount of the secret reserve. On eBay’s community message boards, for example, we observed one experienced user⁵ stating that reserve prices “are simply a form of marketing strategy,” which can get better results because “high minimums get fewer bids.” Kaiser and Kaiser, in *The Official eBay Guide to Buying, Selling, and Collecting Just About Anything* (2000), explain the philosophy as follows: “A high minimum bid is a turnoff even to bidders willing to pay full market price. Set your reserve and start the bidding low. Bidders are likely to bid early or track your item. Such auctions can generate a lot of curiosity, which can translate into bids.”

This reasoning appears to rely on a proposed psychological effect, that bidders can get “caught up in the bidding” at low bid amounts, and end up bidding more than they would have if the bidding had started relatively high. While this “getting caught up” reflects the stated beliefs of a number of bidders and psychologists (see, e.g., Malhotra and Murnighan (2000)), we have not yet seen this effect documented convincingly

⁴ See Lucking-Reiley (1999b) for an experiment that demonstrates the effects of announced reserve prices in online first-price sealed-bid auctions.

⁵ eBay user oscarsale@ixpnet.com, with a feedback rating of 417, placed this message on Mon, 08 May 2000 08:45:34 -0700.

through actions (observed behavior) rather than just words (typically self-reported introspection by bidders).

Vincent (1994) gives a slightly different explanation for the use of secret reserve prices, using a model of rational bidders. Vincent considers a situation where bidders uncertain about their own valuations for the good, and bidders' signals are positively correlated with each other (frequently called the "common values" or "affiliated values" model in the auction literature). He relies on results from Milgrom and Weber (1982), who show that in affiliated-values environments, the seller's expected revenues are enhanced by providing as much information as possible about values to the bidders. Vincent observes that an auction with a low minimum bid and a high secret reserve can provide more information to bidders than an auction with a high minimum bid. When the auction starts at a high minimum bid X , Joe Bidder may be unwilling to meet the minimum bid when no one has yet bid on the item, because he fears the winner's curse. But when the auction starts at a low minimum and other bidders begin to submit bids, Joe then has the opportunity to observe a lower bound on what other bidders are willing to pay. Observing this bidding protects him from the winner's curse, and therefore makes him more likely to bid above the amount X . The key observation in his model is that an auction with a substantial public minimum bid suppresses more bid information than does an auction with a secret reserve price. Although Vincent's model provides the same outcome as the psychological model of "bidding momentum" described above, it reaches the outcome in almost the opposite manner. Vincent relies on the idea that bidders bid conservatively in order to avoid the "winner's curse," while the psychological

explanation relies on the idea that bidders bid too aggressively when they get “caught up” in the heat of bidding.

We are not so ambitious as to try to distinguish between these two observationally equivalent models. Instead, we test their shared implicit assumption that secret reserve prices actually do produce higher expected revenues (after all, goes the reasoning, sellers are observed to use secret reserve prices, and they must know what they are doing!). We feel it is not obvious which selling mechanism is optimal, because we observe considerable disagreement about the subject of reserve prices on eBay’s message boards. Although some sellers appear to use reserves quite frequently, others do not. User **joeaglefeather** wrote, “as a seller, i am FORCED to use them on rare occasion to protect an item from being GIVEN away,” but also indicated that “as a buyer i LIKE reserve auctions.... they 'turn off' a lot of my competitive bidders.”⁶ User **mikejock** appeared to empathize with **joeaglefeather**’s competitors, writing that “using a reserve price and a separate beginning bid is pretty damn STUPID. It not only waste bidders time, but is also an insult.”⁷ Similarly, user **bowerbird-oz** indicated that “I usually hit the back button when I see a reserve auction, especially those which start at \$2. Can’t be bothered wasting my time, I used to bid on them and found every time that the reserve was way above what I was willing to pay.”⁸ He also indicated that when acting as a seller he never used secret reserves, for fear of deterring bidder participation. Clearly, the presence of a secret reserve is capable of the same entry-deterring effects ascribed above

⁶ eBay user joeaglefeather, no feedback profile available, placed this message on Thu, 11 May 2000 14:19:22 -0700.

⁷ eBay user mikejock, with a feedback rating of 26, placed this message on Fri, 05 May 2000 12:04:07 -0700.

⁸ eBay user bowerbird-oz, with a feedback rating 460, placed this message on Mon, 08 May 2000 05:25:16 -0700.

to the presence of high minimum bids.⁹ To the extent that secret reserve prices deter entry as described by these eBay bidders, it seems possible that, contrary to the theories advanced above, the use of a secret reserve price could actually hurt the seller.

Secret reserves may also reduce seller profits in a second way, because their use entails an extra fee. As described in Lucking-Reiley (2000a), eBay charges sellers both a listing fee based on the greater of the minimum bid or the reserve price, and a “final value fee” equal to a percentage of the final sale price. An additional fee is assessed by eBay in cases where seller chooses a secret reserve price and the auction does not result in a sale, eBay assesses an extra fee of \$0.50 for reserves less than \$25, or \$1.00 for reserves greater than \$25. This fee, added by eBay in 1999, discourages eBay sellers from using high reserve prices in their auctions.

Finally, we note that there exists at least one possible explanation for the use of secret reserve prices on eBay which does not predict them to produce higher revenues. When an eBay auction ends, the seller has access to the email addresses of all the bidders. Thus, when a reserve-price auction ends without a sale, the seller can email the high bidder to offer to sell the good to her anyway at the amount of her bid. This strategy allows the seller to avoid paying eBay its percentage commission on the final sale price. Of course, such an offer violates the terms-of-service agreement at eBay, but this would appear to be difficult to enforce. We do not have any quantitative estimates of the frequency with which sellers attempt to sell “offline” in this way, but we have observed

⁹ We should note that the current bid (equal to the minimum bid if there have not yet been any bids) is visible to bidders as they browse lists of auctions on eBay. However, the presence of a secret reserve price is not revealed until the user clicks on that auction listing to view the full auction page. This might lead toward that high opening bids having more entry-deterrent effects than reserve prices do, because noticing the reserve implies a higher level of involvement in the auction. Of course, the relative size of the effect ought to depend on the bidders’ beliefs about how high the secret reserve is likely to be; if bidders (like

several instances of sellers making such offers. It seems clear that the mere presence of reserve prices does not necessarily imply that they increase sellers' expected auction revenues, so this is still an open question. To our knowledge, there has not yet been a direct, quantitative measurement of the effects of secret reserve prices in any auction market. We aim to fill this gap.

III. Experimental Design

For our experiment, we chose to auction matched pairs of cards from the Pokémon trading-card game, which has been the focus of one of the largest collectible toy crazes of 1999 and 2000.¹⁰ Introduced in early 1999, Pokémon game cards appeal both to game players and to collectors; over 50 million cards had already been sold by November 1999.¹¹ The cards come in both Japanese-language and English-language versions, in different sets of cards (Basic Set, Jungle, Fossil, etc.), and in both limited “first editions” and “unlimited editions” (a distinction primarily of interest to collectors). Within an edition, some cards are designated rare, some uncommon, and some rare. Especially rare are the “holofoils,” printed on special foil paper, and a few special promotional cards not sold in the ordinary editions. For the experiment, we chose 50 matched pairs of cards with values high enough to attract bidder interest on eBay. Our cards were all either rare first edition, rare unlimited edition, rare holofoil, or uncommon first edition cards. Our purchase prices at the local card store ranged from \$1.50 to

mikejock above) expect reserves to be outrageously high, then secret reserves might deter more entry than public ones.

¹⁰ Incidentally, Pokémon is published by Wizards of the Coast, maker of the first collectible trading-card *Magic: the Gathering*, whose cards were featured in earlier online auction field experiments; see Lucking-Reiley (1999a, 1999b).

¹¹ “Pokemania vs. Globophobia,” The Economist, Nov. 20, 1999.

\$25.00 per card, with a mean of \$7.19 and a median of \$6.00. We made sure that each card we purchased was unplayed and in excellent condition (without scratches, tears, or nicks).

We conducted our auctions on eBay in April 2000. We created a set of HTML descriptions of the cards to post in their eBay auction listings. The descriptions used for each card in a pair were exactly identical, and all followed roughly the same scheme. We stated that we were willing to sell to bidders in the United States or Canada, and that we would accept as payment a personal check, cashier's check, or money order. In addition, consistent with a number of other card auctions taking place on eBay at the time, we stated that the winning bidder could choose one of two shipping options: USPS First Class Mail for an additional \$0.70 per shipment, or USPS Priority Mail for an additional \$3.20 per shipment. We described the card's edition, rarity, and exact condition, and posted a scanned digital photograph of the card. eBay included an automatic link on each auction page to all other auctions we were running concurrently, as well as to our personal eBay "me" page, as we ran all of these auctions under the same eBay username (**rka469**). We concluded with a notice that we intended to use data on bids for academic research, and provided contact information for questions or concerns.

We started the first fifty auctions on a Sunday between 7pm and 9pm Eastern time,¹² with each auction scheduled to last exactly seven days. Each of the fifty cards was unique. We divided the sample of fifty in half, attempting to make the two groups' distributions of book values as similar as possible. Twenty-five cards (set SP) had a

¹² Because so many bids tend to be received at the very end of an eBay auction (Bajari and Hortag̃su (2000), Roth and Ockenfels (2000)), we attempted to maximize participation by starting and ending the auctions at a time when bidders in all four U.S. time zones were likely to be at home and awake.

minimum bid of \$0.05 and a secret reserve price equal to 30% of the card's book value,¹³ while the other twenty-five (set PS) had a minimum bid of 30% of the card's book value and no secret reserve price. The reserve prices in set SR had a mean of \$2.40 and a median of \$2.10, while the minimum bids in set NR had a mean of \$2.19 and a median of \$1.80. By splitting the sample in half like this, we were able to design the experiment to control for other effects that might vary over time (such as general shifts in demand or supply for these types of cards).

We waited one week after the end of the first group of auctions before starting the second group. Again, the auctions began (and ended) on a Sunday between 7pm and 9pm Eastern time. This time, set PS had a \$0.05 minimum bid and a secret reserve equal to 30% of book value, while set SP had a minimum bid equal to 30% of book, but no reserve. The mnemonics SP and PS are intended to remind the reader which cards had the secret reserve price (S) first and the public minimum bid (P) second, and vice versa. For each card, the reserve price used in the S treatment was exactly equal to the minimum bid used in the P treatment.

At the end of an auction, eBay informed us via email of the results of our auction, along with contact information for the winning bidder. We contacted winning bidders via email and arranged for them to send payment for the cards they had won. After receiving

¹³ Our primary source for book values of the cards was the *Collector's Value Guide: Pokémon*. However, because this guide did not list every single card we purchased (in particular, it excluded special promotional cards), we also used our actual purchase prices as a guide. In general, the book values were slightly higher than the prices we paid in a card store, and card store prices were quite a bit higher than the auction prices we saw for these cards on eBay. For holofoil and promotional cards, we found that the Value Guide prices were generally either missing or were considerably higher than the actual prices we paid, so for these cards we used the purchase prices as our measure of book value. The purpose of collecting book values was merely to help us come up with a reasonable reserve price level – one that was substantial enough to affect bidding, but not so high as to completely suppress all bidding. We aimed to have the probability receiving bids above the reserve be more than 50% but less than 100%, and we were fairly successful in this goal.

a payment check in the mail, we put the cards in protective packing, and shipped the cards using the bidder's preferred shipping method.¹⁴

In an attempt to keep the environment constant between the two sets of auctions in the experiment, we asked each of the winning bidders to refrain from entering feedback information about us on eBay, at least until after the date when our final auctions would be over. Most bidders were very cooperative, but one zealous bidder must have forgotten. On the Wednesday of the second set of auctions, probably the day that this winner received his card shipment from us in the mail, we found that our feedback rating had increased from 0 to 1, with a message stating, "Praise: Prompt, friendly, very dependable... Thanks!! AAA +++." While we appreciate the sentiment of his message, we would have preferred to avoid it for purposes of the experiment. Fortunately, our experiment was designed to be able to control for differences between the two weeks of auctions, by splitting our sample in half and using the opposite time ordering between the two experimental samples. As we shall see below, the prices we received in the second week of auctions (when our feedback rating was 1) tended to be higher overall than the prices we received in the first week (when our feedback rating was 0),¹⁵ but our design still allows us to isolate the effects of the secret reserve price.

¹⁴ We were somewhat surprised how often bidders chose to pay an extra \$2.50 for priority mail, even on purchases of only a few dollars. Furthermore, several bidders urged us to sign up for Paypal (see footnote 1) in order to accept instantaneous payments via credit card, to speed up the transaction process relative to the mailing of checks. The speed of completing a transaction is highly valued by some eBay participants.

¹⁵ Lucking-Reiley et al (2000) find an effect of negative feedback ratings on auction prices for Indian-head pennies, but they find no statistically significant effect of positive ratings. However, their model uses a log-log model of the effect of positive ratings on prices. It is possible that a more nonlinear specification could capture a significant effect, in particular by allowing the difference between 0 and 1 to have a much greater effect than the difference between 10 and 11.

IV. Data and Results

We measure the effects of a secret reserve price (relative to an equivalent public reserve) on three different independent variables: the probability of the auction resulting in a sale, the number of bids received, and the price received for the card in the auction.

IV.A. Probability of Sale

Perhaps the simplest statistic to examine is probability that an auction results in a sale. If secret reserve prices, coupled with trivially low opening bids, tend to encourage bidder entry and more aggressive bidding, they should result in a higher probability of sale. By contrast, if they discourage entry relative to auctions with public minimum bids, they should result in a lower probability of sale.

In the first round of auctions, 14 of 25 cards sold in the public-reserve format, while 10 of 25 sold in the secret-reserve format. In the second week of auctions, 21 of 25 public-reserve cards sold, compared with 13 of 25 secret-reserve cards. Thus, both the secret reserve price and the time ordering of the auctions appear to have separate effects on the probability of sale. The percentage of public-reserve auctions ending in a sale is 70%, while the corresponding figure for secret-reserve auctions is only 46%. This difference is quite statistically significant, with a p-value of 0.0003. There is a difference almost as great between the auctions in the first week (48% sold) and the auctions in the second week (68% sold), a difference which is also statistically significant ($p=0.002$). Because we used matched samples of cards, and auctioned an equal number of cards in each format in each week, we can conclude that the difference we attribute to reserve prices is not merely an artifact of the order of the auctions. In particular, the probability

of sale decreases from 56% to 52% when reserve prices become secret in card set PS, and increases from 40% to 82% when reserve prices become public in card set SP. Thus, the effect of the secret reserve price goes the same direction, independent of the time ordering of experimental treatments.¹⁶

To look at the results on a card-by-card basis, we present a contingency table that displays the fraction of cards selling under both methods, under neither method, or in one method but not the other:

Secret Reserve	Public-reserve	
	SOLD	UNSOLD
SOLD	40%	6%
UNSOLD	30%	24%

As we can see, 40% of the 50 independent cards sold with both secret and public reserve prices, while 24% went unsold in both formats. The remaining 36% of cards make the effects of the secret reserve quite clear: 30% reached the public minimum bid but failed to reach the equivalent secret reserve, while only 6% reached the secret reserve but failed to reach the equivalent public minimum bid.

Overall, these results indicate that using a secret reserve price caused the probability of sale to decrease. Since our experiment used equivalent levels for both public and secret reserves, we conclude that the use of a secret reserve represents a clear loss for the seller, even without taking into account the extra fees imposed by eBay for the use of secret reserves.

¹⁶ The most honest way to treat the statistical test is to consider card set SP and card set PS as two independent tests of the difference in probability of sale. The difference for set SP is statistically significant ($z=4.31$, $p<0.0001$), while the difference for set PS is not ($z=0.40$, $p=0.344$). However, because these two tests are completely independent, we know that their two test statistics are independent standard normal random variables under the null hypothesis of no difference, and therefore we know that their sum is standard normal with mean zero and variance 2. This allows us to compute that the test statistic for a

IV.B. Participating Bidders

As noted above, the effects of secret versus public reserve prices are most likely to occur through their effects on bidders' entry decisions. In this subsection, therefore, we present some measurements of the behavior of individual bidders.

In the first week of auctions, the public-reserve cards received a total of 26 bids, while the secret-reserve cards received 85 bids.¹⁷ In the second week, the public-reserve cards received 53 bids, while the secret-reserve cards received 100 bids. These raw statistics appear to go in the opposite direction as those presented on the probability of sale, as here we observe a higher number of bidders when reserve prices are secret than when they are public. This proves, however, to be an artifact of the data-generating process. In an auction with a substantial public minimum bid, the eBay system accepts no bids less than that amount. By contrast, in an auction with a low minimum bid and a high reserve price, eBay accepts any bids that are above the current bid amount, so bids are not screened out in the same way they are in the public-reserve treatment. To provide a more meaningful measurement, we define a "serious bid" as one above the amount of the reserve price. Focusing only on serious bids, we find 26 in the public-reserve treatment versus 17 in the secret-reserve treatment during the first week, and 53 public-reserve versus 26 secret-reserve serious bids during the second week.

For hypothesis testing, we turn to a regression analysis of the number of serious bids. We run a least-squares regression of SERBIDS, the number of serious bids on a

joint test is $z=3.33$, which means that the overall difference is statistically significant, with p-value equal to 0.0004.

card, against two dummy variables: PUBLIC, which equals 1 with a public reserve versus 0 with a secret reserve, and SECOND, which equals 1 for the second week of auctions and 0 for the first. To control for differences in demand among the 50 different card types used, our specification also includes a full set of card-specific fixed effects.

OLS regression of SERBIDS, the number of serious bids on a card.

(N=100, $R^2=0.6526$, 49 card-specific fixed effects suppressed)

	Beta	Standard Error	t-statistic	p-value
PUBLIC	0.64	0.263	2.270	0.028
SECOND	0.88	0.263	3.121	0.003
Constant	0.78	0.949	0.822	0.415

Both of the independent variables have positive, statistically significant effects on the number of serious bids. The coefficient on PUBLIC indicates that on average, cards with public reserve prices attracted 0.64 more serious bidders than did cards with secret reserve prices. The coefficient on SECOND indicates that cards auctioned in the second week attracted 0.88 more serious bidders on average than did cards auctioned in the first week. Overall, then, our results indicate that although auctions with secret reserves register more bidding activity on eBay, they actually attract fewer serious bidders (i.e., bidders willing to bid at least as much as the reserve price) than do auctions with equivalently high public minimum bids.

As an interesting aside, we present some descriptive statistics on the timing of bids. As noted by Roth and Ockenfels (2000) and Bajari and Hortacsu (2000), many eBay bidders wait until the near the end of an auction to submit their final bids. We find

¹⁷ When we refer to “number of bids” in this paper, we always use a count that measures only one bid per bidder per good. Although a single bidder might raise his bid multiple times, eBay records only the final

results in keeping with theirs: 34% of bidders' final bids occurred in the last 90 minutes of the seven-day auction:

Late bids received in our 100 Pokémon auctions.

Auction Type	Bids in Last 90 minutes	Percentage of All Bids	Winning Bids in Last 90 minutes	Percentage of All Winning Bids
No Reserve	33	42%	20	56%
Secret Reserve	56	31%	15	65%
Total	89	34%	35	59%

We find that secret-reserve auctions produced less late bidding than did public-reserve auctions, with 31% of all final bids falling in the last 90 minutes for secret-reserve auctions, and only 42% of final bids falling in the last 90 minutes in for public-reserve auctions. On the other hand, as shown in the table, the proportion of *winning* bids submitted in the last 90 minutes was higher in the case of the secret-reserve (65% versus 56%). We do not have any proposed explanation for the differences, but merely find them interesting and present them for the benefit of others who may be able to explain them. In our view, these observed differences in the timing of bids are not nearly so economically significant as our main results, on the probability of sale, the number of serious bidders, and (in the next subsection) on the final auction price of the good.

IV.C. Auction Price

The variable of most interest to sellers is surely the final price obtained in the auction. As a descriptive matter, we note that of the 56% of cases where the auction

bid submitted by each individual on each card, and this is what we count in our statistics.

actually resulted in a sale, the majority of cards sold either at or only slightly above the reserve price we chose, and therefore well below our measure of book value. This is consistent with our belief that published lists of book values for collectibles tend to overstate considerably the average transaction price on eBay. (Such price lists tend to be derived from surveys of collectible dealers' in-store prices, rather than from online-auction transactions).

To measure the effect of secret reserve prices on the realized auction price, we define our variable PRICE to be the "current bid" reported by eBay at the end of the auction. Because of the way eBay's proxy-bidding system works, PRICE generally equals one increment over the amount of the second-highest bid submitted. The exceptions are: (1) if no bid exceeds the minimum bid, PRICE equals the minimum bid, and (2) if the highest bid exceeds the reserve price but the second-highest bid does not, PRICE equals the amount of the reserve (whether public or secret). In the case of exception (1), we note that what we really want to measure is the price that would have resulted if the public minimum bid had not been "in the way." That is, when eBay registers no bidders in an auction, we know that the PRICE variable is censored for that observation; the "latent price" is less than or equal to the PRICE recorded.

To examine the effects of a secret reserve on the auction price, we report maximum-likelihood linear regression results where PRICE is the dependent variable. Our Tobit-type maximum-likelihood procedure assumes the error term to be normally distributed, and takes into account that the latent value of PRICE is less than the amount

of the minimum bid when the number of bids equals zero.¹⁸ We use the same dependent variables as in the previous section, again including fixed effects to account for card-specific differences.¹⁹ Our first specification is linear in the dependent variable PRICE:

Maximum-likelihood regression of PRICE, the auction price.

(N=100, log-likelihood = -96.28, 49 card-specific fixed effects suppressed,
18 observations left-censored because they received no bids)

	Beta	Standard Error	t-statistic	Prob > t
PUBLIC	0.613	0.170	3.603	0.001
SECOND	0.620	0.170	3.650	0.001
Constant	2.097	0.572	3.664	0.001

The coefficients on PUBLIC and SECOND are again both positive and statistically significant. The results indicate that a card sold during the second week of auctions will earn an average of \$0.62 more than the same card sold during the first week. More importantly, the results show that, holding all else constant, a public-reserve auction will generate a price \$0.61 higher, on average, than will a secret-reserve auction.

¹⁸ The only difference between our censored-normal regression and a standard Tobit is that our censoring point (the public minimum bid) varies from one observation to the next. We found that the software program *Stata* can easily handle such an estimation problem, with no programming required.

¹⁹ As a side note, we also tried another specification that replaced the card-specific fixed effects with bookvalue as a single regressor (thereby restricting the functional form considerably, in exchange for increased degrees of freedom). To our surprise, this was a gross misspecification: the book value coefficient was insignificantly different from zero (though we expected it to be positive), while the coefficient on PUBLIC had the opposite sign (and statistically significant) from that in the fixed-effects specification. We found this particularly surprising because the experiment was constructed to ensure that the fixed effects and the book value were completely uncorrelated with the PUBLIC dummy variable, in order to eliminate the possibility of omitted-variable bias. Indeed, if our model had been ordinary least squares, the PUBLIC coefficient would be guaranteed to be the same in both regressions, no matter how much measurement error we might have in the book value. However, we find that in our censored-normal specification, the book-value regressor (presumably measured with error) can yield quite misleading results. We feel much more sanguine about our fixed-effects results, reported here, because the model is a much more flexible specification, robust to measurement error in the book value.

As a robustness check, we also report a regression where the dependent variable is the natural logarithm of price:

Maximum-likelihood regression of ln(PRICE).

(N=100, log-likelihood = -80.68, 49 card-specific fixed effects suppressed, 18 observations left-censored because they received no bids)

	Beta	Standard Error	t-statistic	Prob > t
PUBLIC	0.641	0.133	4.825	0.000
SECOND	0.426	0.133	3.210	0.002
constant	0.291	0.468	0.623	0.536

Again, both coefficients turn out positive and statistically significant. The economic interpretations are that on average, a second-week auction generates a 53% higher price than a first-week auction, and a public reserve generates a 90% higher price than a secret reserve.

V. Conclusion

We conclude that the use of secret reserve prices caused us to earn less money as sellers, relative to the practice of making our reserve prices publicly known. Making our reserve prices secret had negative effects on probability of selling a card, the number of serious bidders in the auction, and the price received from the winning bidder. Only 46% of secret-reserve auctions resulted in a sale, compared with 70% of public-reserve auctions for the same goods. Secret-reserve auctions resulted in 0.88 fewer serious bidders per auction, and \$0.62 less in final auction price, than did public-reserve auctions on average.

We can recommend with confidence that sellers avoid the use of secret reserve prices, at least for Pokémon cards with expected auction prices of \$25 or less, as in this experiment. Our recommendation would remain the same even if eBay did not charge an additional fee for the use of secret reserves, as the effects on the auction outcome are unambiguously negative for the seller. We do not know how far our results will generalize; it is possible that they will fail to hold for some other types of goods or other ranges of prices. Our results provide quantitative support for the following recommendation by Kaiser and Kaiser (1999) to eBay sellers in *The Official eBay Guide*: “If your minimum sale price is below \$25, think twice before using a reserve auction. Bidders frequently equate reserve with expensive.”²⁰

The quote from Kaiser and Kaiser implies a belief that more expensive goods (over \$25) might tend to benefit more from the use of secret reserves, though they present no quantitative evidence on this topic. And indeed, Bajari and Hortaçsu (2000) document the empirical fact that, for mint sets of US coins, “items with higher book value tend to be sold using a secret as opposed to posted reserve price with a low minimum bid.” Using a structural econometric model to estimate bidding parameters, they compute results consistent with observed behavior. In particular, they conclude from their simulations that for items with book values less than \$10, a public minimum bid dominates a secret reserve for the seller, but for items with book values greater than \$10, a secret reserve dominates a public minimum bid.

Our results are somewhat inconsistent with those of Bajari and Hortaçsu. Their estimated model predicts that “expected revenue from a secret reserve price exceeds the revenue from an ordinary auction where the minimum bid is set at the same level as the

²⁰ Kaiser and Kaiser (1999), pg. 106.

secret reserve price.” In their model, it is only eBay’s secret-reserve fees that cause secret reserves not to be worthwhile for low-valued goods. By contrast, we find in our experiment that secret reserve prices have negative effects on expected revenues, even without taking into account the additional fees incurred. Perhaps Bajari and Hortaçsu have made an inaccurate modeling assumption, or perhaps there is some important difference between bidding for coin sets and bidding for Pokémon cards.

At any rate, we do take the valuable point both from Kaiser and Kaiser (1999) and from Bajari and Hortaçsu (2000) that secret reserves might well become more useful to sellers when the goods being auctioned are more expensive. A new experiment, auctioning one hundred items each in the \$100 range, could shed some important light on this question. Indeed, we would very much like to see growth in the use of such field experiments, as we feel they are valuable methods for gathering data, intermediate between the use of laboratory experiments²¹ and the difficult econometric modeling of naturally occurring field data. For the time being, however, we note that our results appear to be the only available direct measurements of the effects of secret reserve prices, and therefore we plan to eschew secret reserves when trying to maximize our own expected revenues in eBay auctions.

²¹ Laboratory experiments, with induced values for fictitious goods, have been used to investigate a wide variety of hypotheses from auction theory. See Kagel (1995) for a review.

References

Ashenfelter, Orley (1989): "How Auctions Work for Wine and Art," *The Journal of Economic Perspectives*, 3(3), 23-36.

Bajari, Patrick, and Ali Hortaçsu (2000): "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions," Stanford University Working Paper, 1-49.

Collector's Value Guide: Pokémon, premiere edition, Middletown, CT: CheckerBee Press, 1999.

Easley, Robert F., and Rafael A. Tenorio, "Bidding Strategies in Internet Yankee Auctions." Working paper, University of Notre Dame, 1999.

Engelbrecht-Wiggans, Richard, John A. List, David Lucking-Reiley (1999): "Demand Reduction in Multi-unit Auctions with Varying Number of Bidders: Theory and Field Experiments," Vanderbilt University Working Paper.

Hendricks, Kenneth, Robert H. Porter, Charles A. Wilson (1994): "Auctions for Oil and Gas Leases with an Informed Bidder and a Random Reservation Price," *Econometrica*, 62(6), 1415-1444.

Kagel, John H. (1995): "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, J. Kagel and A. Roth, eds. Princeton: Princeton University Press, pp.501-585.

Kaiser, Laura Fisher and Michael Kaiser, *The Official eBay Guide to Buying, Selling, and Collecting Just About Anything*, Simon & Schuster, New York, 1999.

Lucking-Reiley, David (1999a): "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, December 1999.

Lucking-Reiley, David (1999b), "Magic on the Internet: Evidence from Field Experiments on Reserve Prices in Auctions," Vanderbilt University working paper, February 1999.

Lucking-Reiley, David (2000a), "Auctions on the Internet: What's Being Auctioned, and How?," *Journal of Industrial Economics*, forthcoming.

Lucking-Reiley, David (2000b), "Vickrey Auctions in Practice: From 19th-Century Philately to Twenty-First-Century E-Commerce," *Journal of Economic Perspectives*, forthcoming.

Lucking-Reiley, David, Doug Byran, Naghi Prasad, Daniel Reeves, (2000): “Pennies from eBay: the Determinants of Price in Online Auctions,” Vanderbilt University working paper, April 2000.

Milgrom, Paul R., Robert J. Weber (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50(5), 1089-1122.

“Pokemania v. Globophobia,” *The Economist*, November 20, 1999.

Roth, Alvin E., and Axel Ockenfels (2000): “Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet,” Harvard University Working Paper.

Vincent, Daniel R. (1995): “Bidding Off the Wall: Why Reserve Prices May be Kept Secret,” *Journal of Economic Theory*, 65, 575-584.

Wilcox, Ronald T. (1999): “Experts and Amateurs: The Role of Experience in Internet Auctions,” Carnegie Mellon University Working Paper.