

Public versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment

Rama Katkar and David H. Reiley¹

First Version: July 2000

This Version: April 2005

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. We also present evidence that some sellers choose to use secret reserve prices for reasons other than increasing their expected auction prices.

¹ Katkar: Graduate School of Business, Stanford University, email <katkar_rama@gsb.stanford.edu>. Reiley: University of Arizona, email <reiley@eller.arizona.edu>. An earlier version of this paper was Katkar's senior honors thesis in Mathematical Methods in the Social Sciences at Northwestern University; she acknowledges the financial support of an Undergraduate Research Grant from the Weinberg College of Arts and Sciences. Reiley acknowledges the National Science Foundation for support under grants SBR-9811273 and SES-0094800, and the hospitality of the Department of Management & Strategy, Kellogg School of Management, Northwestern University. We are grateful to Adam Bayer and Michael Urbancic for excellent research assistance, and to Rob Porter for extremely thorough comments and suggestions. This paper has benefited from comments received in presentations at the National Bureau of Economic Research, the Economic Science Association, the University of Chicago Graduate School of Business, the University of Illinois, the University of Oregon, the University of Hawaii, the University of Michigan School of Information, and the Stern School of Business at New York University.

1. 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?

Online commerce presents economists with exciting opportunities to conduct field experiments (see, for example, Lucking-Reiley (1999), Reiley (2005)). 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.² 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.³

² Laboratory experiments, with induced values for fictitious goods, have also been used to investigate a wide variety of hypotheses from auction theory. See Kagel (1995) for a review.

³ Lucking-Reiley (2000a) surveys 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

Indeed, a variety of authors have begun to exploit this rich source of data to investigate economic questions. For example, Lucking-Reiley, *et al* (2004), Hauser and Wooders (2005), and Melnik and Alm (2002) examine the effect of eBay “feedback ratings” on the final auction price. Roth and Ockenfels (2002) 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 (2004) 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.⁴ One can only observe a lower bound on the reserve price in

payment services such as Paypal and Billpoint, which make it possible for individual sellers to 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.

⁴ It might be possible to learn the exact amount of the secret reserve if the data were obtained directly from eBay, which has been assembling a data warehouse for statistical research. Unfortunately, eBay’s data warehouse does not currently include information on the actual item that was auctioned, because the description of the item would take up too much space in the database (personal communication with

an auction where the reserve was not met, and an upper bound in an auction where the reserve price was met. One could in principle get more information by observing exactly when the auction switched from the status of “the reserve has not yet been met” to “the reserve has been met,” but eBay’s public bid history does not record this information.

Bajari and Hortaçsu (2003) provide an indirect measurement of the effects of secret reserve prices in eBay auctions. In a sample of 407 eBay auctions of uncirculated and proof coin sets, they observed 14% of items to have had a secret reserve. The secret-reserve items featured lower minimum bids than the other auctions (28% of book value versus 69% of book value), and were less likely to result in a sale (49% versus 84%). Both the minimum bid amount and the presence of a secret reserve tend to be associated with fewer bidders in an auction, though the latter effect is not statistically significant. Though these reduced-form results are interesting, it is hard for Bajari and Hortaçsu to determine the precise effects of the secret reserve price, because the reserve-price amounts are set endogenously by the sellers but unobserved by the researchers. They write, “Since eBay does not divulge the secret reserve price at the close of the auction, we cannot compare the secret reserve prices to minimum bid levels in ordinary auctions,” and they can only “conjecture that secret-reserve price levels are, on average, higher than the minimum bids in ordinary auctions.”

With enough data and enough identifying econometric assumptions, one could conceivably tease out an empirical measurement of the reserve-price effect from eBay field data. Indeed, Bajari and Hortaçsu present the above results as part of a larger effort to provide a structural econometric model of bidding. Such structural models have to

Michael Dearing, eBay’s Director of Strategic Planning, November 2000). Thus, even though the data on reserve prices and bids might in principle be available from eBay, a study of reserve prices would be

make strong identifying assumptions in order to recover economic unobservables (such as bidders' private information about the item's value). For example, Bajari and Hortaçsu assume that the unobserved secret reserve-price amounts are set as if they were bids from an independent bidder. By contrast, our own research project is much less ambitious, as we focus only on the effect of secret reserve prices. Our experiment allows us to carry out this measurement in a manner that is as 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). While we have designed this experiment to answer a question that we find economically interesting in its own right, we also hope that such experiments may prove useful to researchers trying to build structural models of market behavior, by providing consistency checks on some of the implications of their models.

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 actual sellers' discussions. Section three provides details of our experimental design, including the relevant eBay institutional details. Section four presents our empirical results, and section five concludes.

2. 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

hindered by a lack of data on what the goods are worth.

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.”⁵ 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.

Using a secret reserve price in an eBay auction entails an extra fee for the seller. For all auctions, 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: \$0.50 for reserves less than \$25, or \$1.00 for reserves greater than \$25. This fee, added

⁵ One other effect of the reserve price shows up in the use of “proxy bidding” on eBay (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).

by eBay in 1999, appears designed to discourage eBay sellers from using high secret reserve prices in their auctions.

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 unsold.⁶ 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 user⁷ 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* (1999), 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."

⁶ Reiley (2005) discusses an experiment that demonstrates the effects of announced reserve prices in online first-price sealed-bid auctions.

⁷ eBay user **oscarsale@ixpnet.com**, with a feedback rating of 417, placed this message on Mon, 08 May 2000 08:45:34 -0700.

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., Ku, Malhotra, and Murnighan (2004)), we have not yet seen this effect documented convincingly through actions rather than just words (i.e., observed behavior rather than self-reported introspection by bidders).

Vincent (1995) 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 equivalent⁸ models. Instead, we test their shared implicit assumption that secret reserve prices actually do produce higher expected revenues. 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."⁹ 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."¹⁰ 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

⁸ If the value of the item could be observed after the auction, the two models might be empirically distinguishable. Vincent's model predicts conservative bidding, so that the auction price should be below the winner's realized value. Depending on the psychological assumptions, the "momentum" model might well predict overbidding relative to realized value.

⁹ eBay user **joeaglefeather**, no feedback profile available, placed this message on Thu, 11 May 2000 14:19:22 -0700.

¹⁰ eBay user **mikejock**, with a feedback rating of 26, placed this message on Fri, 05 May 2000 12:04:07 -0700.

them and found every time that the reserve was way above what I was willing to pay.”¹¹ He also indicated that when acting as a seller he never used secret reserves, for fear of deterring bidder participation. It seems the presence of a secret reserve may be capable of the same entry-detering effects ascribed above to the presence of high minimum bids.¹² 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.¹³

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 provide such a measurement.

3. 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 2000.¹⁴ Introduced in early 1999, Pokémon game cards appeal both to game players and to collectors. Over 50 million cards, bearing individual names

¹¹ eBay user **bowerbird-oz**, with a feedback rating 460, placed this message on Mon, 08 May 2000 05:25:16 – 0700.

¹² 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 high opening bids having more entry-deterent 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 **mikejock** above) expect reserves to be outrageously high, then secret reserves might deter more entry than public ones.

¹³ In some affiliated private-value auction models, it is possible for revenues to decrease in the number of bidders, so the effect on revenues is ambiguous. Pinske and Tan (2005) show that reducing the number of bidders can increase revenues even in situations where bidders’ values are privately known, a result previously known to have been true for common-value models (see, for example, Klemperer, 1998).

¹⁴ Pokémon happens to be 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 (1999), Reiley (2005). Wizards of the Coast is now a fully-owned subsidiary of Hasbro, Inc.

such as Charizard, Pikachu, and Magneton, were sold by November 1999.¹⁵ 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 common. Especially rare are “holofoils,” printed on special foil paper, together with 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 promotional, 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 \$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 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

¹⁵ “Pokemania vs. Globophobia,” *The Economist*, Nov. 20, 1999.

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.

Each card was auctioned twice, once with a public reserve and once with a secret reserve. We started the first fifty auctions on a Sunday between 7pm and 9pm Eastern time,¹⁶ 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 card values as similar as possible. Twenty-five cards (set SP) had a minimum bid of \$0.05 and a secret reserve price equal to 30% of the card’s book value,¹⁷ 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 mnemonics PS and SP are intended to remind the reader which cards had the public minimum bid (P) first and the secret reserve price (S) second, and vice versa. The reserve prices in set SP had a mean of \$2.40 and a median of \$2.10, while the minimum bids in set PS had a mean of \$2.19 and a median of \$1.80. (See Figure 1 for histograms of these two distributions.) 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).

¹⁶ Because many bids tend to be received at the very end of an eBay auction (Bajari and Hortag̃su (2003), Roth and Ockenfels (2000)), we attempted to maximize participation by starting and ending the auctions at a time when bidders in all four primary U.S. time zones were likely to be at home and awake.

¹⁷ Our primary source for book values of the cards was the *Collector’s Value Guide: Pok mon*. 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. (Of the fifty cards, we had twelve holofoils and four promotional cards.) 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 proportion receiving bids above the reserve be more than 50% but less than 100%, and indeed the overall proportion turned out to be 59%.

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. We use numbers to refer to the first and second auction of each card set, thereby distinguishing our four different experimental treatments. For example, “Treatment PS1” refers to the auctions with public minimum bids for card set PS, with public minimum bids, while “Treatment PS2” refers to the auctions with secret reserve prices for the same cards. Table 1 summarizes this experimental design, and provides some descriptive statistics.

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 a payment check in the mail, we put the cards in protective packing, and shipped the cards using the bidder’s preferred shipping method.¹⁸

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

¹⁸ 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.

feedback rating had increased from 0 to 1, with a message stating, “Praise: Prompt, friendly, very dependable... Thanks!! AAA +++.” While we appreciated 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),¹⁹ but our design still allows us to isolate the effects of the secret reserve price.

4. 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.

4.1. Probability of Sale

Our first question is whether secret reserve prices affect the probability of 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.

¹⁹ Lucking-Reiley *et al* (2004) find an effect of negative feedback ratings on auction prices for Indian-head pennies, but they find no statistically significant effect of positive ratings. By contrast, Houser and Wooders (2005) and Melnik and Alm (2002) find significant effects of both positive and negative feedback ratings, in auctions for Pentium III chips and U.S. gold coins, respectively.

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, 22 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 72%, while the corresponding figure for secret-reserve auctions is only 46%. This difference is quite statistically significant, with a p-value of 0.0061. There is a difference almost as great between the auctions in the first week (48% sold) and the auctions in the second week (70% sold), a difference which is also statistically significant ($p = 0.022$). 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 88% 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. The average effect is statistically significant.²⁰

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. As the Table 2 shows, 20 of the 50 individual cards sold with both secret and public reserve prices, while 11 went unsold both times. The remaining 19 of cards make the effects of the secret reserve quite clear: 16 reached the public minimum

²⁰ To conduct a hypothesis test, we regressed a SALE dummy on three other dummy variables: week, card set (SP versus PS), and reserve type (secret versus public). The coefficient on the secret-reserve dummy then constitutes a difference-in-difference estimate of the effect on the probability of sale. The point estimate is 47% with a standard error of 11%, for a p-value less than 0.0001.

bid but failed to reach the equivalent secret reserve, while only 3 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 price 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.

4.2. 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.²¹ In the second week, the public-reserve cards received 53 bids, while the secret-reserve cards received 101 bids. These raw statistics appear to go in the opposite direction from those presented on the probability of sale, as we observe a higher number of bidders when reserve prices are secret than when they are public. However, this proves 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

²¹ 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, at the time of the experiment eBay's bid history recorded only the final bid submitted (their maximum willingness to pay) by each individual on each card, and this is what we count in our statistics. In current eBay bid histories individual bidders may be listed multiple times.

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 per auction. We run a least-squares regression of SERBIDS, the number of serious bids on a 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.²²

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 0.72 more serious bidders than did cards with secret reserve prices. The identical coefficient on SECOND indicates that cards auctioned in the second week attracted 0.72 more serious bidders on average than did cards auctioned in the first week.²³ The estimated fixed effects ranged from -1.5 to $+3.5$ bidders per card.

²² Given that both PUBLIC and SECOND turn out to have significant effects, a previous reader of this paper suggested including an interaction term PUBLIC*SECOND in the regression as well. However, such an effect would not be identified in our specification. Because the experiment is designed to measure differences between identical pairs of cards, we already include a full set of card-specific dummy variables. The proposed regressor PUBLIC*SECOND turns out to be a linear combination of PUBLIC, SECRET, and the dummy variables, so it cannot be included in the regression. If we drop the card dummies, and include only PUBLIC, SECOND, and the interaction term, we find that all three coefficients are positive but none are statistically significant at the 10% level, so the PUBLIC treatment effect is not statistically significantly different in one week versus the other.

²³ Because the two independent variables produced identical coefficient and standard-error estimates, we initially feared that we had made a computational error. However, it turns out (as one can verify from the

Overall, 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.

One might wonder whether this difference in the number of observed serious bidders could be due to data censoring, rather than due to differences in the number of bidders who *wanted* to participate. In particular, because eBay will not register a bid amount lower than the current high bid in an auction, it is possible that a bidder arriving late in the auction will find his willingness to pay will already have been exceeded, so he will not show up in the data. If secret-reserve auctions reached high prices faster than did public-reserve auctions, then our results could be caused by this data-censoring effect rather than by differences in actual bidder participation. We did not save detailed information about the time path of prices in our auctions, so we cannot be certain about this possible source of bias, but we believe the data-censoring effect is likely to be small. Bajari and Hortaçsu (2003) and Roth and Ockenfels (2002) both find that eBay bidders tend to submit bids very late in auctions, which would minimize the possibility of bidders being screened out merely because they arrive on the last day of the auction.

Inspired by these other authors, we did record some statistics on the extent of late bidding. First, we recorded the number of bidders submitting their final bids in the last 90 minutes of the seven-day period: 43% of participating bidders in the public-reserve auctions (34 of 79), and 31% in the secret-reserve auctions (58 of 186). These figures

descriptive statistics in Table 1) that the conditional mean number of serious bids per card increases from 0.86 to 1.58 on average when moving from a secret to a public reserve price, and again from 0.86 to 1.58 on average when moving from the first to the second week of the experiment. The identical standard errors

include all bidders, not just serious bidders who bid above the reserve price. We also found that 58% of winning bids occurred in the last 90 minutes of public-reserve auctions (21 of 36), compared with 65% of winning bids in secret-reserve auctions (15 of 23). Of the two statistical measures we collected, the second one would seem to be a better indicator of bidding activity by serious bidders, because it restricts attention to those auctions where the reserve price was exceeded. Overall, any differences in late bidding activity between public-reserve and secret-reserve auctions appear to be insignificant, which indicates that any data-censoring bias on the number of serious bidders is also likely to be minimal.

To summarize, we find the number of serious bidders observed in our experiment to be higher in public-reserve auctions than in the equivalent secret-reserve auctions. This could well be the reason for our final finding, presented in the next section: that realized auction prices are also higher with public than with secret reserve prices.²⁴

4.3. Auction Price

The most interesting outcome to measure is the price realized in the auction. Simple descriptive statistics in Table 1 show that of the 100 auctions in the experiment, 41 auctions failed to reach the reserve price, 27 sold exactly at the amount of the reserve,

result from the experimental design, which ensures that each of the two dummy variables equals 1 exactly 50% of the time, which means that they have identical variances.

²⁴ As noted in footnote 13, increased participation need not increase the price, if bidders' values are not independent of each other. Bajari and Hortaçsu (2003) assume a common-value model of eBay auctions for proof sets of coins, and in their econometric model secret reserve prices decrease entry and thereby increase revenues. In their reduced-form regressions, they estimate that the presence of a secret reserve reduces the number of bidders per auction by .162 on average, though the effect is statistically insignificant.

and 32 sold at prices above the reserve.²⁵ We now wish to measure the effect of public versus secret reserve prices on the level of the price reached in the auction.

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 the first exception, 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 assume that the PRICE variable is censored for that observation; the “latent price” is less than or equal to the PRICE recorded.²⁶

To examine the effects of a secret reserve on the auction price, we report censored-normal maximum-likelihood linear regressions where PRICE is the dependent variable. Our Tobit-type maximum-likelihood procedure assumes the error term to be normally distributed, and considers the latent value of PRICE to be less than the amount of the minimum bid whenever the number of bids equals zero.²⁷ We use the same dependent variables as in the previous section, again including fixed effects to account

²⁵ Recall that we set our reserve prices equal to 30% of published book values, so the actual eBay transaction prices do tend to be considerably lower than published lists of book values. Such price lists tend to be derived from surveys of collectible dealers’ in-store prices, rather than from online-auction transactions

²⁶ Note that here we are assuming away any potential “bidding momentum” effect.

²⁷ 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.

for card-specific differences.²⁸ Our first specification is linear in the dependent variable PRICE, and the results are given in Table 4.

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.

In Figure 2, we plot the estimated card fixed effects versus the reserve prices we employed. Two cards (Haunter, with a reserve price of \$1.20, and Seaking, with a reserve price of \$0.60) are omitted because these two cards had estimated fixed effects of negative infinity in the censored-normal regression. These two cards failed to elicit any bids in either treatment, even when the minimum bid was \$0.05, so they provided no information about the effects of either PUBLIC or FIRST on PRICE. Only three other cards, all in set SP, received no bids in their secret-reserve auctions, but these did receive bids in their subsequent public-reserve auctions. The figure shows that our reserve prices were reasonably well correlated ($r = 0.64$) with the resulting auction prices for each card. The regression line for this plot yields an intercept of -0.13 (not significantly different from zero at the 10% level), and a slope of 0.68 (significantly less than 1 at the 1% level),

²⁸ 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 were 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

indicating that auction prices increase somewhat less than one-for-one with the reserve prices we chose as 30% of book value. Even when we omit the outlier on the upper right of the plot (Pikachu), the estimated slope of the line increases only to 0.74 (still significantly less than 1 at the 5% level).

As a specification check, we also report, in Table 5, a regression where the dependent variable is the natural logarithm of price. Again, both coefficients are positive and statistically significant. By exponentiating the point estimates, we find 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. The main difference between the linear and the logarithmic specifications is that in the logarithmic specification, the effect of the main treatment variable (PUBLIC) is greater in magnitude than that of the nuisance variable (SECOND). (This difference is not statistically significant at the 5% level, however.) We conclude that our main result is robust to our choice of specification: a secret reserve price has a negative, statistically significant effect on the realized auction price.

5. Conclusion

We have found that the use of secret reserve prices caused us to earn less revenue 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

results. We feel much more sanguine about the fixed-effects results reported here, which are robust to measurement error in the book value.

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 therefore recommend that sellers avoid the use of secret reserve prices, particularly for Pokémon cards. Our negative 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 negative even without taking this fee into account. 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. In fact, as we only experimented with one reserve-price level, they might even fail to hold for other reserve-price levels with these same goods. Our results provide some 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.”²⁹

The quote from Kaiser and Kaiser implies a belief that sellers of 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 (2003) 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. They note that this prediction is

²⁹ Kaiser and Kaiser (1999), pg. 106.

consistent with their data on seller behavior, as they observe secret reserves more often in auctions for higher-valued proof sets.

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 secret reserve price,” because of winner’s-curse effects. 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.

In any case do take the valuable point both from Kaiser and Kaiser (1999) and from Bajari and Hortaçsu (2003) 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, for example, could shed some important light on this question. Particularly useful would be to identify a category of goods for which reserve prices are very frequently used on eBay, because one might expect reserve prices to be employed most often in those categories of goods where they prove to be most useful for sellers. We know from casual experience that secret reserve prices are used in considerably less than half of Pokémon auctions, but we regrettably do not have any detailed statistics on the use of reserve prices in Pokémon cards or other categories. One exception we have is a dataset of Indian-head pennies from July and

August 1999, described in Lucking-Reiley, *et al.* (2004). In that data, with 2795 auctions of Indian cents with a mean book value of \$150 (minimum \$1, maximum \$25,000), we find that 19.3% employed secret reserve prices and 10.0% had secret reserve prices that were never met.

We note that there exists at least one possible explanation for the use of secret reserve prices on eBay that does not require 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.³⁰ Sellers might therefore choose to employ high secret reserve prices to implement this strategy even if they do not expect the reserve to produce high auction revenues. Indeed, the more valuable the item, the higher eBay's standard commission would be, and hence the more incentive to use secret reserves to try to circumvent eBay's collecting the fee.

In order to estimate the prevalence of this practice, we conducted a survey of bidders on eBay. By examining completed auctions on eBay, we identified the high bidders in 171 different auctions where the reserve price was never met. We collected this data across all categories on eBay, collecting a random sample according to eBay auction ID numbers. We then emailed a survey to those high bidders, identifying the

³⁰ Since the time of our experiment, eBay has implemented changes to make it more difficult for eBay users to contact each other outside of eBay. In particular, messages between users are now expected to go through the eBay messaging system, and users are prohibited from choosing usernames that reveal their private email addresses.

specific auction in which they had been the high bidder, and asked each of them whether the seller had contacted them to offer to sell the item anyway. We received 48 responses, a response rate of 28%. Of the 48 respondents, 13 (or 27%) indicated that the seller did email them to offer to sell the item even though the auction itself had not resulted in a transaction. Of the thirteen, six were offered the item at the amount of their bid, six were offered a higher price than the bid amount, and one was offered a lower price than the bid amount. When asked if they had ever been contacted by the seller in such a situation, 52% of the 48 respondents indicated that they had, and 31% indicated that they had completed a transaction in this way at some point. Thus, the practice of unofficially completing a “reserve not met” transaction off of eBay does not happen every time, but it is reasonably prevalent, occurring in slightly more than a quarter of “reserve not met” auctions.³¹

Overall, then, we have learned that secret reserve prices can reduce realized auction prices, at least in auctions for Pokémon trading cards. We have also learned that there are reasons why sellers might wish to use secret reserve prices, in particular the strategy of contacting the high bidder to try to conduct a transaction without actually paying eBay’s commission. Open research questions remain, of course. First, for what types of goods are secret reserve prices most often employed? Second, for those types of goods, do secret reserve prices actually increase the prices realized in the auction? For example, a seller who fears bidder collusion might wish to use a secret reserve to deter

³¹ Sample-selection bias is likely to cause us to underestimate the frequency of this practice, to the extent that bidders might fear revealing that they were involved in a transaction that violates eBay’s rules. However, we did promise anonymity to all of our respondents, and we also made sure to ask buyers about the sellers’ behavior rather than about their own behavior in that specific transaction, to minimize such concerns on the part of respondents. Therefore, we expect the extent of sample-selection bias to be relatively modest.

pooling of bids at the public minimum bid amount. For now, 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 auction revenues on eBay.

References

- Bajari, Patrick, and Ali Hortaçsu (2003): “Winner’s Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions,” *RAND Journal of Economics*, 34(2), 329-355.
- Collector’s Value Guide: Pokémon*, premiere edition (1999), Middletown, CT: CheckerBee Press.
- Easley, Robert F. and Rafael Tenorio (2004): “Jump-Bidding Strategies in Internet Auctions,” *Management Science*, 50(10), 1407-1419.
- Engelbrecht-Wiggans, Richard, John A. List and David H. Reiley: “Demand Reduction in Multi-unit Auctions with Varying Numbers of Bidders: Theory and Field Experiments,” forthcoming, *International Economic Review*.
- Houser, Daniel and John Wooders: “Reputation in auctions: Theory and Evidence from eBay,” forthcoming, *Journal of Economics and Management Strategy*.
- 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 (1999), *The Official eBay Guide to Buying, Selling, and Collecting Just About Anything*, Simon & Schuster, New York.
- Ku, Gillian, Deepak Malhotra, and J. Keith Murnighan (2004): “Milked for All Their Worth: Competitive Arousal and Escalation in the Chicago Cow Auctions,” working paper, Northwestern University Kellogg Graduate School of Management.
- Lucking-Reiley, David (1999): “Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet,” *American Economic Review*, 89(5), 1063-1080.
- Lucking-Reiley, David (2000a): “Auctions on the Internet: What’s Being Auctioned, and How?” *Journal of Industrial Economics*, 48(3), 227-252.
- Lucking-Reiley, David (2000b): “Vickrey Auctions in Practice: From 19th-Century Philately to Twenty-First-Century E-Commerce,” *Journal of Economic Perspectives*, 14(3), 183-192.
- Lucking-Reiley, David, Doug Byran, Naghi Prasad, Daniel Reeves, (2004): “Pennies from eBay: the Determinants of Price in Online Auctions,” working paper, University of Arizona.

- Melnik, Mikhail I., and James Alm (2002): "Does a Seller's eCommerce Reputation Matter?" *Journal of Industrial Economics*, 50(3), 337-350.
- Milgrom, Paul R., Robert J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50(5), 1089-1122.
- Pinske, Joris and Guofu Tan (2005): "The Affiliation Effect in First Price Auctions", *Econometrica*, 73(1), 263-267.
- "Pokemania v. Globophobia," (1999) *The Economist*, November 20.
- Reiley, David H.: "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet," forthcoming, *RAND Journal of Economics*.
- Roth, Alvin E., and Axel Ockenfels (2002): "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet," *American Economic Review*, 92(4), 1093-1103.
- 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," *Marketing Letters*, 11(4), 363-374.

Table 1. Experimental design and descriptive statistics.

<i>Treatment name</i>	<i>PS1</i>	<i>PS2</i>	<i>SP1</i>	<i>SP2</i>	Total
Type of reserve price	Public	Secret	Secret	Public	
Week	First	Second	First	Second	
Number of auctions	25	25	25	25	100
Minimum card value	\$1.50	\$1.50	\$1.50	\$1.50	\$1.50
Median card value	\$6.00	\$6.00	\$6.00	\$6.00	\$6.00
Maximum card value	\$14.00	\$14.00	\$25.00	\$25.00	\$25.00
Mean card value	\$6.88	\$6.88	\$7.50	\$7.50	\$7.19
Mean reserve price amount	\$2.19	\$2.19	\$2.40	\$2.40	\$2.30
Total number of bids*	26	101	85	53	265
Total “serious” bids**	26	26	17	53	122
Items sold	14	13	10	22	59
at the reserve price	10	2	5	10	27
above the reserve price	4	11	5	12	32
Items receiving no bids	11	3	2	3	19
Total revenue on sold items	\$31.88	\$34.85	\$29.35	\$65.61	\$161.69
Mean revenue on sold items	\$2.28	\$2.68	\$2.94	\$2.98	\$2.74
standard deviation	\$1.96	\$1.52	\$1.27	\$1.74	\$1.65

* Here “number of bids” actually refers to the number of bidders. See footnote 21.

** For a definition of “serious” bids, see section 4.2.

Table 2. Probability of sale in the two auction treatments.

Secret Reserve	Public Reserve	
	SOLD	UNSOLD
SOLD	40%	6%
UNSOLD	32%	22%

Table 3. OLS regression of SERBIDS, the number of serious bids on a card.

(N = 100, $R^2 = 0.6526$, 50 card-specific fixed effects suppressed)

	Beta	Standard Error	t-statistic	p-value
PUBLIC	0.72	0.263	2.736	0.009
SECOND	0.72	0.263	2.736	0.009

Table 4. Maximum-likelihood regression of PRICE, the auction price.

(N = 100, log-likelihood = -96.84, 50 card-specific fixed effects suppressed,
18 observations left-censored because they received no bids)

	Beta	Standard Error	t-statistic	p-value
PUBLIC	0.627	0.168	3.722	0.001
SECOND	0.660	0.168	3.919	0.000

Table 5. Maximum-likelihood regression of ln(PRICE).

(N = 100, log-likelihood = -85.44, 50 card-specific fixed effects suppressed,
18 observations left-censored because they received no bids)

	Beta	Standard Error	t-statistic	p-value
PUBLIC	0.688	0.139	4.968	0.000
SECOND	0.497	0.139	3.584	0.002

Figure 1. Distribution of card values.



Figure 2. Estimated card fixed effects versus reserve price

