

Measuring the Benefits to Sniping on eBay: Evidence from a Field Experiment

Sean Gray

Sullivan & Cromwell LLP, New York University School of Law, U.S.A.

David H. Reiley¹

Google, Inc., U.S.A.

Bidders on eBay frequently engage in sniping: that is, submitting a bid seconds before an auction closes. Using a field experiment, we attempt to measure the size of the benefit of sniping (if any). To do this we selected pairs of auctions that were as identical as possible (same item, same quality, same seller, similar closing time, *et cetera*). In one auction of each pair we submitted an early bid, and in the other we submitted the same bid exactly ten seconds before its close. Our results, from a set of 70 such pairs, indicate no benefit to sniping: we found evidence of 2.54% lower prices for the sniped auctions, but we did not find this benefit to be statistically significant. Additionally, we use our data to make a survey of the incidence of sniping and compare our findings with those of Roth and Ockenfels (2002).

Keywords: bidding, second-price auction

JEL Classification: C93, D44

1 Introduction

Online auctions have been in use for over a decade. By far the most well-known online auction site, eBay, was started in September 1995 and has since grown to

¹Gray: Sullivan & Cromwell LLP, <sean.gray@gmail.com>. Reiley: Google, Inc., <reiley@google.com>. This paper grew out of Gray's senior honors thesis at the University of Arizona. Reiley gratefully acknowledges financial support from SES- 0094800, and valuable research assistance from David Mack, Steve Manos, Mike Urbancic, and Alex Shcherbakov. Sarah Goss, Jill Sazama, and an anonymous referee provided helpful comments.

include tens of millions of registered members from around the world.² As noted by many scholars, eBay's scale and diversity encourage its use as a platform for empirical studies of auction theory.³

One of the salient features of eBay, that the site's users and observers quickly notice, is the ubiquitous phenomenon of sniping. eBay employs a "hard close" in their auctions (see Ockenfels and Roth (2002)), meaning each auction has a definite, publicly known time at which it will end. Many eBay users believe it can be advantageous to submit the very last bid in a given auction, so that no rival bidder has an opportunity to react or counter the bid. For this reason, sniping is commonly perceived to increase the chance of winning an auction, decrease the final price paid for the item in question, or both.

The default process for bidding in eBay is the proxy bidding procedure (see Lucking-Reiley (2000)). Each bidder is asked to submit the maximum amount she is willing to pay for an item. When each bid is submitted, eBay automatically increases the current price of the auction to one increment above the second highest bid submitted to that point. The exact level of the current highest bid is unknown until it can be inferred by exceeding it. Proxy bidding, thus, essentially entails a collapsed English auction between the standing high bid and each new rival bid. eBay gives the following description of proxy bidding to their users:

1. *When you place a bid, you enter the maximum amount you'd be willing to pay for the item. Your maximum amount is kept confidential from other bidders and the seller.*
2. *The eBay system compares your bid to those of the other bidders.*
3. *The system places bids on your behalf, using only as much of your bid as is necessary to maintain your high bid position (or to meet the reserve price). The system will bid up to your maximum amount.*
4. *If another bidder has a higher maximum, you'll be outbid. BUT, if no other bidder has a higher maximum, you win the item. And you could pay significantly less than your maximum price!*

²For a more detailed history of online auctions, see Lucking-Reiley (2000).

³Bajari and Hortacsu (2004) and Ockenfels, Reiley, and Sadrieh (2007) provide useful reviews of the academic literature on online auctions.

Proxy bidding allows bidders maximum flexibility in placing their bid (they may submit their bids at any time—day or night—during the auction period) and requires minimum effort since each bidder needs make only one bid.

If bidders submit their true maximum willingness to pay, there is no need for them to submit a higher bid later. In particular, if all bidders in an auction were to submit their true maximum willingness to pay (assuming this remains constant through the duration of the auction), there would be no benefit at all to sniping; in this situation no rival would want or need to react to a late bid. eBay encourages this use of the proxy bidding process and admonishes each bidder to submit her maximum willingness to pay, as reflected in the following notice regarding sniping:

One way to help avoid disappointment is to ensure that the maximum bid you enter on the item page is the highest price that you're willing to pay. The eBay system automatically increases your bid up to the maximum price you specify, so entering a higher maximum may help prevent you from being outbid in the closing seconds of a listing.

Yet non-proxy bidding appears to have remained quite popular. We find a large incidence of sniping below, as did Roth and Ockenfels (2002) in two categories of items, computers and antiques:

20 percent of all last bids on eBay ... were submitted within the last hour ... Furthermore ... a considerable share of bidders submit their bid in the last five minutes (9 percent in Computers and 16 percent in Antiques) ... At the auction level: 40 percent of all eBay-Computers auctions and 59 percent of all eBay-Antiques auctions ... have last bids in the last 5 minutes ... In the 240 eBay-auctions, 89 have bids in the last minute and 29 in the last ten seconds.

Perhaps the prevalence of sniping is due to a misunderstanding by users about the nature and purpose of proxy bidding. Despite the excellent description of the bidding mechanism, some bidders may still treat eBay auctions as live English auctions. In particular, they may believe that the full value of their bid will be

posted, as if it were a first-price rather than a second-price auction rule. This would discourage such bidders from bidding their maximum willingness to pay for the item. It is quite common to observe a bidder rapidly submitting a sequence of incrementally higher bids, implying that the bidder wishes to bid slightly more than a rival's proxy bid. This indicates that some bidders are hesitant to reveal their true maximum willingness to pay, and that they misunderstand the nature of proxy bidding. In the presence of such naïve bidders who do not fully understand the mechanism, a sophisticated bidder might reason: why should I submit my full bid early in the auction, giving the naïve bidder a chance to respond and raise her bid, rather than waiting until the end to snipe and not give her a chance to respond? Ariely, Ockenfels and Roth (2006) state this reasoning as follows:

Sniping may also be a best response to incremental bidding that is observed both in the field... and in our experimental setting. An incremental bidder starts with a bid below his value and is then prepared to raise his bid when he is outbid. Bidding late on eBay may be a best reply to incremental bidding, because this strategy would not give the incremental bidder any opportunity to respond to being outbid.

Ely and Hossain (2009) point out a second theoretical reason for sniping, which is that bidders may be choosing among multiple auctions for very similar goods. Suppose I begin shopping for a Luke Skywalker action figure on Wednesday at 1pm. My search turns up ten different auction listings for the item I want, with scheduled auction end times ranging from Wednesday at 2pm to the following Monday at 6pm. I have some incentive to bid on the auction with the soonest end time: in this example, the auction ending only an hour from now, because if I win I get the item sooner, while if I lose I still have the option to bid in the subsequent auctions. If everyone tends to bid in auctions ending soon, this automatically generates late bids in auctions without bidders engaging in particularly strategic behavior.

Sniping has become so popular that a few businesses have been created for the sole purpose of facilitating the submission of last-second bids on eBay. eSnipe is

one example of an online service that allows a bidder to automatically submit a sniping bid for a given online auction at the amount and time of her choosing (say, within a few seconds of the auction's close).⁴ The ease of use of this service as well as the policy to only charge the service fee to customers that actually win the item on which they bid has facilitated the growth of sniping in the late 90s to the levels noted above.

Ariely, Ockenfels, and Roth (2006) conducted sniping experiments in a laboratory context. They found that as bidders gained more experience in the controlled online bidding marketplace, they were more likely to place late bids; this correlation implies sniping may offer a tangible benefit to bidders. This is also consistent with the finding of Roth and Ockenfels (2002) that more experienced (higher-reputation) bidders were more likely to snipe on eBay. In this paper, we go outside the laboratory in order to measure this effect on eBay using a field experiment.

2 Experimental Method

We designed our experiment to measure the benefit of sniping by identifying matched pairs of identical items listed on eBay. We placed an early bid for one of the items, and placed a snipe bid for the other item in the pair. In identifying potential pairs of items, we controlled for a number of variables that could add noise to our data, such as seller, shipping cost, day and time of the auction end, and starting bid level. As much as possible, we tried to choose each pair of auctions so that our bidding strategies were the only factor that varied between them. Thus, any change in the prices paid in the paired auctions would be primarily due to the respective bidding strategies used.

In choosing the types of items to purchase, we looked for categories high availability of identical items. The types chosen were: PlayStation 2 video games, US coin proof sets, DVDs, Xbox video games, Hot Wheels die-cast cars, and Game

⁴As of 2004, this service had 50,000 registered users—as compared to eBay's 65,000,000 users—and placed more than 10,000 bids a day, and those bids totalled about \$16 million per month. For reference, the value of goods sold on eBay in 2003 was \$20 billion. As the \$16 million figure from eSnipe represents the total value of all eSnipe bids (not just winning bids), the proportion of eBay community that used eSnipe was fairly small. While most users apparently submitted snipes manually, we chose to employ eSnipe to guarantee automated, consistent timing of our snipe bids.

Boy Advance video games. These items were more likely to be identical because most of the goods were new, and therefore in the same condition. For the goods that were not new, they were of a measurable condition to which a value could be attached (e.g., the coin proof sets). The categories used were also useful in that they ranged from relatively small bidder pools (Xbox video games) to relatively large bidder pools (DVDs). This also allowed us to look at the effects when different amounts of bidders are taking part in the auction.

Once we identified the types of goods we were interested in, we selected 20 pairs of PlayStation 2 games and 10 pairs of items for each of the remaining five types, for a total of 70 pairs of items. To find these pairs of items, we browsed through each specific category of newly listed goods until we found two items with identical names listed near each other. We started with newly listed goods so that if a suitable pair were found we would be able to place an early bid in one of the two auctions. Identical item listings, including the same use of punctuation and capitalization, usually indicated that the items were being sold by the same seller. Searching for multiple identical items offered by the same seller was the fastest way to locate potential pairs of nearly identical auctions, since such auctions are very often listed at approximately the same time, last for the same number of days, have the same starting bid, require the same shipping costs, and are advertised by identical item display pages. Upon finding a potential pair of auctions, we examined the details of each to confirm that they had the same seller, shipping costs, and auction length. The two auctions in each pair closed typically within hours of each other, as described in Table 1. Our data show that the average number of bidders is lower for the proxy-bid auctions than the sniped auctions, with the averages being 3.34 and 5.56, respectively. We believe this is due to the fact that a late (sniped) bid keeps the price lower for a longer period of time, and thus encourages more entry.

Table 1: Differences in Auction Pair Close Times by Item Type

Item	<i>Time Difference</i>			
	Mean	Max	Min	Median
Playstation 2 Games	14:04:45	23:59:34	3:03:35	23:59:26
Game Boy Advance Games	0:00:48	0:02:07	0:00:00	0:00:37
Xbox Games	3:05:13	3:12:20	2:58:24	3:01:48
Coin Proof Sets	0:21:09	0:25:02	0:04:50	0:23:35
Hot Wheel Cars	0:00:57	0:02:04	0:00:26	0:00:34
DVDs	0:09:16	1:00:00	0:00:25	0:01:49
Average Total	2:57:01	4:46:51	1:01:17	4:34:38

Once we selected a particular auction pair, we needed to determine an appropriate bid level. We wanted the bid to be high enough to win the auction, because in the case that we lose the auction, we do not learn how high the other bidders were willing to go.⁵ However, we did not want it to be an unreasonably high bid, to avoid a large monetary risk. We decided to use the Wal-Mart retail price—conveniently found on their website—for our bid levels for specific Playstation 2 video games, DVDs, Xbox video games, and Game Boy Advance video games. The items in question were in new condition, and Wal-Mart prices are usually the lowest price at which a consumer can buy a product at an easily accessible “brick-and-mortar” store. Based on the widespread availability of Wal-Mart stores and the additional cost of shipping items won on eBay, we felt that most consumers were unlikely to bid above this threshold. The US coin proof sets and Hot Wheels die cast cars were all in mint condition and could be priced at their mint condition price guide value. The coin proof sets were valued by the Professional Coin Grading Service: <http://www.pcgs.com/prices/>, while the Hot Wheels cars were priced by

⁵Ely and Hossain (2009) experimented with several lower bid amounts. We believe that our strategy of very high bids, designed to win the auction, maximizes statistical power per observation. Ely and Hossain measure both the impact on the final price and the impact on the probability of winning the auction, but with our strategy, the total effect of the snipe can be summarized in the final auction price, given that we submitted very high bids. Submitting a low experimental bid and losing the auction translates into submitting a high experimental bid, winning, and paying a higher final price in the auction. The lower-bid case cannot tell us exactly how much the experimenter would have had to pay to win, so the high-bid case less data censoring and more information.

Hallsguide.com online price guide: <<http://www.hallsguide.com/hotwheels.html>>. The authority and widespread availability of these pricing guides gave us confidence that we could win most of these auctions at book price without risking excessively high payments.

In half of the auctions (one item in each pair) we submitted our bid several days before the auction was to end. In the other half of the auctions, we used automated sniping software to place an identical bid exactly 10 seconds before the auction concluded. We hoped that 10 seconds before close would be as near as possible to the end of the auction while ensuring enough time to make sure that our bid would get through and be processed. Since our bids were selected to be high relative to going market prices on eBay, we tended to win both items in each pair of auctions. We randomized whether the snipe treatment was on the first or the second auction in each pair, in order to eliminate any correlation between our treatment and any order effects that might be present on eBay for multiple auctions from the same seller. All auctions took place between November 2003 and April 2004.

3 Description of the Data Sample

For each of the 70 pairs of items on which we bid, we copied both the auction page and the bidding page. Clicking for a description of the item and viewing the auction page allowed us to obtain the most basic information about an auction, including:

- Item name
- Item number
- Starting bid
- Winning bid
- Auction length
- Number of bids
- Number of bidders
- Time ended
- Seller

Clicking the link on the number of bids on the auction page allowed us to view the bidding page where there was a record of when each bid was made and by whom. These two pages provided all of the necessary data for our analysis. Descriptive statistics of bids that won their auctions and the bids we placed in each category are shown in Tables 2 and 3 respectively.

Table 2: Descriptive Statistics of Winning Prices by Item Type

Item	<i>Winning Price</i>			
	Mean	Max	Min	Median
Playstation 2 Games	\$13.70	\$33.07	\$4.95	\$11.50
Game Boy Advance Games	\$20.19	\$26.91	\$11.53	\$21.04
Xbox Games	\$19.99	\$36.00	\$3.95	\$20.01
Coin Proof Sets	\$17.91	\$32.03	\$10.50	\$15.50
Hot Wheel Cars	\$8.83	\$11.57	\$6.99	\$8.55
DVDs	\$13.25	\$18.50	\$8.10	\$13.05
Overall	\$15.64	\$36.00	\$3.95	\$12.80

Table 3: Descriptive Statistics of Our Bids by Item Type

Item	<i>Our Bid</i>			
	Mean	Max	Min	Median
Playstation 2 Games	\$27.12	\$49.82	\$14.85	\$19.82
Game Boy Advance Games	\$30.81	\$34.82	\$29.79	\$29.82
Xbox Games	\$35.88	\$47.72	\$19.72	\$44.82
Coin Proof Sets	\$32.62	\$45.00	\$16.00	\$25.00
Hot Wheel Cars	\$17.19	\$30.00	\$10.00	\$15.00
DVDs	\$18.02	\$19.98	\$15.87	\$17.98
Overall	\$26.94	\$49.82	\$10.00	\$19.98

Although our raw data contained 70 pairs, we were forced to remove some of the pairs of items. Three pairs of US coin proof sets, one pair of Playstation 2 games, and two pairs of DVD movies were eliminated for the purposes of our

analysis because we were outbid in one or both of the auctions, with eight total auctions affected.⁶ Of these eight auctions, we were outbid on three sniped bids and five early bids. We needed to win both auctions in each pair to measure the benefit of sniping, so the affected auction pairs did not provide us with useful data. Also, eBay administrators removed five of the ten pairs of Game Boy Advance games due to suspicion of piracy after we had placed our bids. The total number of pairs remaining for statistical analysis was thus 59.

It is important to note that there were some auctions in which we were the sole bidder. Moreover, five complete pairs of items had no other bidders, which may potentially cause us to underestimate the benefit of sniping. In addition, there were seven individual auctions for which we were the only bidder, while the other auction of the pair had more than one bidder. We retain these auctions in our analysis because they give us the opportunity to include in our analysis the possibility that early bids deter entry by competing bidders (relative to sniping).

4 Data Analysis

We first analyze our claim that the price of a sniped good should be relatively lower than its unsniped counterpart. To get a feel for the data, we present a histogram displaying the difference between the sniped price and proxy price for an auction pair. The bins are sorted by the difference in prices, where lower bins correspond to a larger benefit to sniping. Visual inspection shows that price difference is unimodal but skewed to the left. This indicates a benefit to sniping, which we now quantify numerically.

⁶The auctions in which we were outbid were surprising to us in that the winners of those auctions paid more than retail price for the Playstation 2 games and DVDs, and more than the “book value” for the coin proof sets. Malmendier and Lee (2011) document systematically that overbidding behavior, relative to available posted prices, is fairly common on eBay.

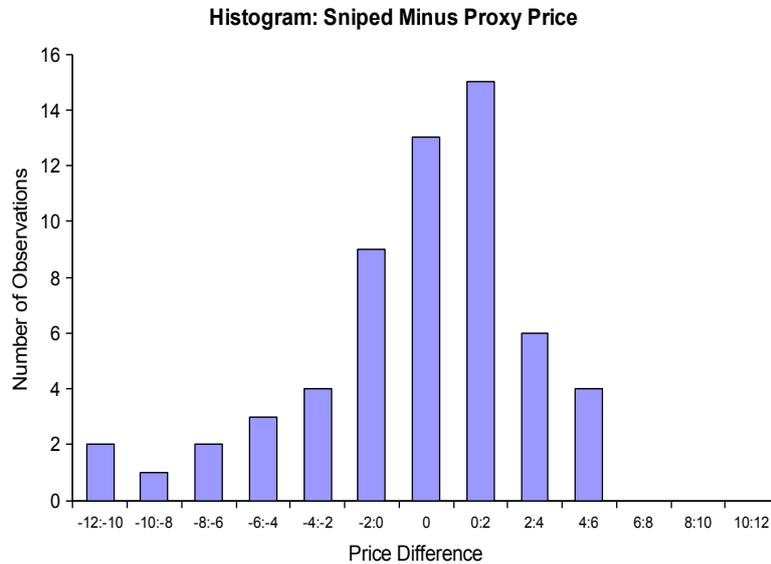


Figure: The Differences between the Sniped Price and Proxy Price for an Auction Pair

Although we find a small benefit to sniping, it is not statistically significant. We perform paired-sample tests of means between the closing price of a sniped auction and that of the paired auction in which we placed an early bid. The mean percentage difference between the sniping price and the early-bid price is -2.54% , with a 95% confidence interval of $[-9.39\%, 4.13\%]$. Therefore, on average, sniping resulted in a slightly lower closing price relative to bidding early. However, the test statistic of -0.74 indicates that this result is not significant at any reasonable level. Testing instead on the absolute differences in price between the auctions in each pair (sniping minus early bid) produces similar results. The mean absolute difference was $-\$0.50$, with a 95% confidence interval of $[-\$1.43, \$0.43]$. This suggests a possible benefit to sniping, but is not significant at the 5% significance level ($t=-1.09$).

We also perform tests of proportions. Under the null hypothesis of no effect, the proportion of pairs where sniped auctions closed at a higher price than early bid auctions should be equal to the proportion where the reverse occurred. A test of this null yields a t-statistic of -1.16 , so we cannot reject this null hypothesis at a 5% significance level.

Overall, we obtain small point estimates of the benefits to sniping, and find that

these estimates are not statistically significant. For items with closing auction prices on the order of \$15, we estimate the benefit to sniping to be approximately $\$0.50 \pm \1.00 . While we clearly have limited statistical power, we do see positive point estimates of the benefits to sniping, and we find that these benefits are relatively small. The upper bound to our confidence interval is on less than \$1.50 per item, consistent with the result of \$1.00 obtained by Ely and Hossain (2009). It is also worth comparing this benefit to the \$0.25 fee charged by eSnipe for bids on items with prices up to \$25.00 on eBay.

We next consider the possibility that the benefit to sniping may vary according to the number of competing bidders. The following are two competing stories that can be told to motivate intuition either way. It is plausible that sniping may be more effective when there are more bidders participating in an auction, since a snipe is less likely to affect the outcome of an auction that has one or zero rival bidders (thin markets will have fewer naïve bidders to best-respond to). Alternatively, sniping may be less effective when there are many bidders, since the probability that the high bidder has significantly overbid (see Malmendier and Lee (2011)) using eBay's proxy bidding may increase in the number of bidders. In this case, if the sniping bid is lower than the current high bidder's overbid, it will only have the effect of raising the closing price for that winning bidder.

Sorting our dataset into two subsets, we attempt to answer whether the benefits to sniping vary with the number of competing bidders. We sort the 59 pairs of auctions by the average number of bidders in each pair, grouping by pair in order to avoid introducing endogenous variation in the number of bidders. We then split our dataset above and below the median pair, leaving ourselves with two subsets of 29 pairs.⁷ One subset consists of pairs with relatively low numbers of bidders; the average number of bidders for this subset is 2.62, and the average number of bidders for the high subset is 6.28.

Focusing first on the low subset, we find that sniped auctions closed an average of \$0.72 or 4.19% less than the early-bid auctions. Running tests of means, the respective *t*-statistics (absolute and relative differences, respectively, for sniping minus early bid) are -0.73 and -1.22, so the small point estimate of a benefit to

⁷We dropped the median pair, which happened to be unfavorable to the sniping treatment: we paid \$3.29 more in the snipe treatment for an item with \$20 book value.

sniping is not statistically significant. Turning our attention to the high subset, we find a similar result. Among those pairs, the sniped auctions also yielded cheaper closing prices than their proxy counterparts. On average an auction in the high subset yielded a sniped price lower by \$0.50, or 1.91%, relative to the the early-bidding price. However, a test of the equivalence of means in the two groups fails to reject the null of no effect, with t -statistics of -0.69 and -0.48 for absolute and percentage differences, respectively. While we do not have clear statistical significance, our point estimates for the benefits of sniping are higher when there are fewer competing bidders, suggesting that thicker markets may reduce the benefits of sniping.

5 Incidence of Sniping

We were surprised to find small, statistically insignificant benefits to sniping, given past results by Roth and Ockenfels (2002) and Ariley, Ockenfels, and Roth (2006). There are three possibilities that could explain these results. First could be low statistical power, since our confidence intervals are twice as wide as our point estimates. Second, it is possible there was less sniping on eBay as a whole during the period when we were bidding on our auction pairs than during the period when Roth and Ockenfels (2002) collected their data. Third, it may be that sniping is less prevalent within the categories of goods that we chose to bid on than in the categories studied by Roth and Ockenfels (2002).

To investigate these latter two possibilities, we decided to collect additional data in spring 2004 to determine whether sniping still occurs at frequencies similar to those reported by Roth and Ockenfels (2002). To accomplish this, we mimicked their data gathering by choosing 120 laptop computer auctions and 120 antiques auctions, for a total of 240 auctions. We chose the first auctions we saw where the reserve price was met and the Buy-It-Now option was not exercised.⁸

Like Roth and Ockenfels, we then sorted late bids into four overlapping bins: bids made in the last hour, bids made in the last ten minutes, bids made in the last minute, and bids made in the last ten seconds. Our results were that 78.3% of the

⁸In some auctions a buyer can elect to use the Buy it Now option to negate the auction format and simply directly purchase the item at a known, stated price.

bids were made in the last hour, 63.8% in the last ten minutes, 50% in the last minute, and 25.4% in the last ten seconds. Roth and Ockenfels (2002) found 67.9% of bids occurred in the last hour, 55% in the last ten minutes, 37.1% in the last minute, and 12.1% in the last 10 seconds. Our results actually show a noticeable increase in the prevalence of sniping over time, with the percentage of bids occurring in the last hour, ten minutes, one minute, and ten seconds all higher in our sample than in Roth and Ockenfels (2002).

We also added data from two other categories in which we had attempted to measure the benefit to sniping: Playstation 2 games and US coin proof sets. The results are summarized in Table 4. We do find less sniping in our categories than in the categories chosen by Roth and Ockenfels, particularly in coin proof sets, where we see no sniping at all (other than our own) in the last ten seconds of the auction. Sniping appears to have increased in prevalence in the categories originally chosen by Roth and Ockenfels, but it turns out that at least some of our new categories have less sniping activity than those original categories. The fact that we measure small, insignificant benefits of sniping in our experiments may be due in part to our having chosen categories where sniping is relatively less beneficial.

Table 4: Incidence of Sniping on eBay in Spring 2004

Item Type	<i>Percentage of</i>				# of Items
	<i>Bids Made</i>				
	<i>Within:</i>				
	1 hour	5 min	1 min	10 sec	
Laptops + Antiques	78.3%	63.8%	50%	25.4%	240
Playstation 2 Games	90.0%	50.0%	33.3%	7.0%	30
US Coin Proof Sets	70.0%	36.7%	6.7%	0.0%	30

6 Conclusion

Previous research has implied benefits to sniping in online auctions. In a controlled experiment, we directly measure these benefits to be small and statistically insignificant, with sniping yielding a price reduction of $\$0.50 \pm \1.00 relative to placing our bid early. We hypothesized that the benefits to sniping

might be small due to bidder competition, but we found no evidence of the benefits to sniping being greater in thinner versus thicker markets. We also hypothesized that the benefits of sniping may have eroded over time as bidders have become more sophisticated, but we found that the frequency of sniping remained as high during the period of our study as it did during Roth and Ockenfels' data collection five years earlier. We did find that the categories studied by Roth and Ockenfels experience somewhat more sniping than the categories we chose to study in our experiment.

Our lack of statistical significance may be due in large part to lack of statistical power. Ely and Hossain (2009) have independently performed a similar experiment with four times as many auctions, and managed to find a statistically significant benefit of slightly more than \$1.00 per auction. Several differences between our experimental designs are worth noting. While we always submitted a very high bid, in order to have a high chance of winning and therefore measuring a quantitative effect, Ely and Hossain varied their treatment between high and low bids. While we chose a number of different types of goods, from Hot Wheels cars to collectible coin sets to Playstation games, Ely and Hossain focused exclusively on DVDs. A further difference between our studies is that we focused on identical matched pairs of items, while Ely and Hossain allowed for other differences between auctions. Our studies complement each other in showing that the small measured benefits of sniping are robust over a relatively wide range of eBay auctions.

References

- Ariely, D., A. Ockenfels, and A. E. Roth, (2005), "An Experimental Analysis of Ending Rules in Internet Auctions," *RAND Journal of Economics*, 36, 890-907.
- Bajari, P. and A. Hortacısu, (2004), "Economic Insights from Internet Auctions," *The Journal of Economic Literature*, 42, 457-486.
- eBay, Inc., (2004), "Proxy Bidding (for Buyers)," from <http://pages.ebay.com/help/buy/proxy-bidding.html>.
- Ely, J. and T. Hossain, (2009), "Sniping and Squatting in Auction Markets," *American Economic Journal: Microeconomics*, 1, 68-94.

eSnipe, Inc. ,(2004), “eSnipe Information,” from <http://www.esnipe.com/AbouteSnipe/>.

Lucking-Reiley, D., (2000), “Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce,” *Journal of Economic Perspectives*, 14, 183-192.

Lucking-Reiley, D., (2000), “Auctions on the Internet: What's Being Auctioned, and How?,” *Journal of Industrial Economics*, 48,227-252.

Malmendier, U. and Y. H. Lee, (2011), “The Bidder’s Curse,” *American Economic Review*, 101,749-787.

Ockenfels, A., D. H. Reiley, and A. Sadrieh, (2007), “Online Auctions,” In: Terrence, H. (ed.), *Handbook of Information Systems and Economics*, Elsevier Science, 571-628.

Roth, A. E. and A. Ockenfels, (2002), “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review*, 92, 1093-1103.