

Consumer Surplus in Online Auctions

Ravi Bapna⁺, Wolfgang Jank, Galit Shmueli*

March 22, 2005

ABSTRACT

Despite the growing research interest in Internet auctions, particularly those on eBay, little is known about the quantifiable consumer welfare accrued from such mechanisms. Using an ongoing novel field experiment, we collect and examine a unique dataset to empirically quantify and understand determinants of consumer surplus in eBay auctions. Our analysis, based on a sample of 5157 eBay auctions, indicates that the median surplus level per auction on eBay auctions is \$3.53, which roughly translates to \$1.47 billion in accrued consumer surplus for the year 2003 alone. We find that consumer surplus is significantly different across currencies and item categories, negatively influenced by seller experience, auction duration and competition, and positively influenced by bidder experience, bidder aggressiveness and item price.

Keywords: Consumer Surplus, Online Auctions

⁺ Contact author.

* Bapna: Dept. of Operations and Information Management, UConn School of Business, Storrs, CT 06269, rbapna@sba.uconn.edu phone: (860) 486-8398; Jank: Decision & Information Technologies Dept., Robert H. Smith School of Business, University of Maryland, College Park, MD 20742, wjank@rhsmith.umd.edu, (301) 405-1118. Shmueli: Decision & Information Technologies Dept., Robert H. Smith School of Business, University of Maryland, College Park, MD 20742, gshmueli@rhsmith.umd.edu (301) 405 9679.

Bapna gratefully acknowledges the support of the CIDRIS and Treibeck Electronic Commerce Initiative (TECI), UConn. Jank and Shmueli's research was partially funded by the NSF grant DMI-0205489. They gratefully acknowledge support by the Center of Electronic Markets and Enterprises, U of Maryland.

The authors thank Joe Bailey, Dipak Dey, Robert Garfinkel, Alok Gupta, Ali Hortaçsu, Curt Hunter, Jim Marsden and Alex Tung for their feedback on earlier versions of this paper. The paper has benefited from feedback received at the Workshop on Information System Economics (WISE) 2004, and at research seminars at the University of Connecticut's OPIM department, University of Connecticut's Statistics department and at the University of Minnesota.

Author names are in alphabetical order.

INTRODUCTION

Classical microeconomic theory uses the notion of consumer surplus as the welfare measure that quantifies benefits to a consumer from an exchange. Alfred Marshall (1936) defined consumer surplus as “the excess of the price which he (a consumer) would be willing to pay rather than go without the thing, over that which he actually does pay...” It is also traditional to visualize consumer surplus as the roughly triangular area lying under a downward sloping demand curve and above the rectangle that represents actual money expenditure. Yet, despite its established theoretical standing, little is known empirically about consumer surplus levels in real markets. In this paper we report on consumer surplus levels in eBay auctions, a vastly popular Internet based electronic market. eBay’s popularity is evident in the reported \$23.8 billion in gross merchandise sales for the year 2003, up from \$14.9 billion in 2002.

Internet based markets are now mainstream artifacts of today’s economy. Yet, very little has been generally said about the quantifiable benefits such markets provide to consumers. There are two notable exceptions. Erik Brynjolfsson, Yu Hu and Michael D. Smith. (2003) in an interesting analysis, demonstrate how new product introduction in electronic markets can lead to significant consumer welfare gains. There is also an emerging stream of work, exemplified in Anindya Ghose, Michael D. Smith, and Rahul Telang (2004), which is looking closely at the welfare implications (including consumer welfare) of secondary markets for used books. Both studies creatively devise econometric estimation procedures to measure consumer welfare levels where they are not *directly observable* in posted price markets. Typically, in such markets, willingness to pay has to be inferred indirectly through surveys, contingent valuation techniques and price changing experiments such as promotions and discounts. Surveys of willingness to pay have credibility issues and have lead to a stream of research dealing with contingent valuation precision and bias reduction [see Peter A. Diamond and Jerry A. Hausman (1994)]. There has also been some controversy about which consumer welfare measure is most appropriate. Hausman (1981) explains the difference between primal Marshallian consumer surplus and the dual Hicksian compensating variation measures. He then goes on to provide an approach to derive the unobserved compensated Hicksian demand curve from the observed demand curve, which in turn leads to an exact

measure of consumer surplus. These developments have led to initial empirical research in measuring consumer surplus levels in specialized markets [Brynjolfsson, Hu and Smith (2003) for obscure book titles, Ghose, Smith and Telang (2004) for secondary markets]. However, it is our belief that more needs to be said and done in a wider context about consumer surplus.

In this paper we demonstrate the suitability of using direct mechanisms [Roger Myerson (1981)], such as auctions in which it is a dominant strategy for market players to reveal their true valuations, to quantify and understand determinants of consumer surplus. Auction theory is built upon the fact that a consumer with a valuation v_i for an item, strategizes and formulates a bid b_i so as to maximize her surplus $(v_i - b_i)$ [R. Preston McAfee and John McMillan (1987)]. Note the similarity here with consumer theory, which evaluates consumer surplus as the value or willingness to pay for a change in price of a good from say p^0 to p^1 . Thus, the consumer theory notion of the amount the consumer would pay (or would need to be paid) to be just as well off after the price change as she was before the price change, is no different from the auction theory notion that measures the winning bidder's surplus as how much could the price have gone up, without changing the current allocation¹. Not surprisingly, induced by the vast amount of bidding data on the Internet, connections between auction and consumer theory are beginning to be formally explored by researchers and practitioners². Christopher P. Adams and William Vogt (2004) show the highly elastic nature of the market for digital cameras using three different methods for estimating demand for differentiated products on eBay. Bidding data from eBay also has value to antitrust regulators who, while evaluating a merger, seek to understand the competitive structure of a particular product market [Christopher P. Adams and Laura L. Bivins (2003)].

While auctions have been around for centuries, only in recent times, in conjunction with advancements in Internet technologies, have they reached the scope and scale to be considered mainstream for consumers. Lucking-Reiley (2000) provides an overview of what is being auctioned by whom and under what mechanism rules on the Internet. The growing popularity of Internet auctions has

¹ We hereafter refer to surplus, bidder's surplus and consumer surplus interchangeably in this paper.

² eBay sells its price data to Andale.com which attempts to provide Blue Book like services for consumer products.

been accompanied with equal fervor amongst researchers revisiting auction theory, and finding new and creative uses for the vast amount of bidding data available. These range from classifying bidding strategies in multi-unit Yankee auctions [Ravi Bapna et al. (2004)], to determining optimal bid increments [Bapna et al. (2003)], to visualizing online auction data [Galit Shmueli and Wolfgang Jank (2004)] to studying the dynamics of online auctions [Wolfgang Jank and Galit Shmueli (2003)]. For an excellent review of research on eBay in particular see Patrick Bajari and Ali Hortascu (2004). We expect to contribute to this stream of research by shedding light on the important metric of consumer surplus on eBay auctions.

There have been several studies that have described, in depth, eBay's second-price ascending proxy-bid auction mechanism. A key feature is that at the termination of the auction, the highest bidder wins and pays a price equal to the second highest bid plus one bid increment. The only exception is the case that the two highest bids are equal, wherein the earlier bidder is awarded the item at a price equal to her bid. Another established feature, resulting primarily from eBay's hard closing time, is that last minute bidding or sniping is widely prevalent. Al Roth and Axel Ockenfels (2002) provide interesting theoretical and empirical insights into sniping behavior on eBay. They observe that in 240 antique auctions on eBay, 89 had bids in the last minute and 29 in the last 10 seconds.³ Similar findings have been reported by Bajari and Hortaçsu (2003), Shmueli, Russo and Jank (2003) and Schindler (2003). Explanations for late bidding range from tacit collusion against sellers to the presence of naïve bidders who don't understand proxy bidding, to common value components in the items being auctioned. For the purpose of this study, we make use of the fact that sniping is prevalent.

eBay posts almost the complete bid history after the auction closes, *with the exception being the value of the highest bid*. For instance consider the bid history of the following auction for a Nokia 6610 GSM cell phone. Conspicuous in its absence is the exact amount bid by the winner 'kanchenjunga'.⁴ Since the winner's bid is not disclosed by eBay, there is no direct measure of the revealed willingness to

³ Our own conversations with a senior eBay executive reveal that 80 percent of all bids arrive in the last hour of the auction.

⁴ The first author's eBay id.

pay of the winning bidder from the auction that is publicly available. To overcome this limitation, we design an ongoing field experiment that allows real-world bidders to use Cniper.com, our web based bidding agent, to snipe eBay auctions.

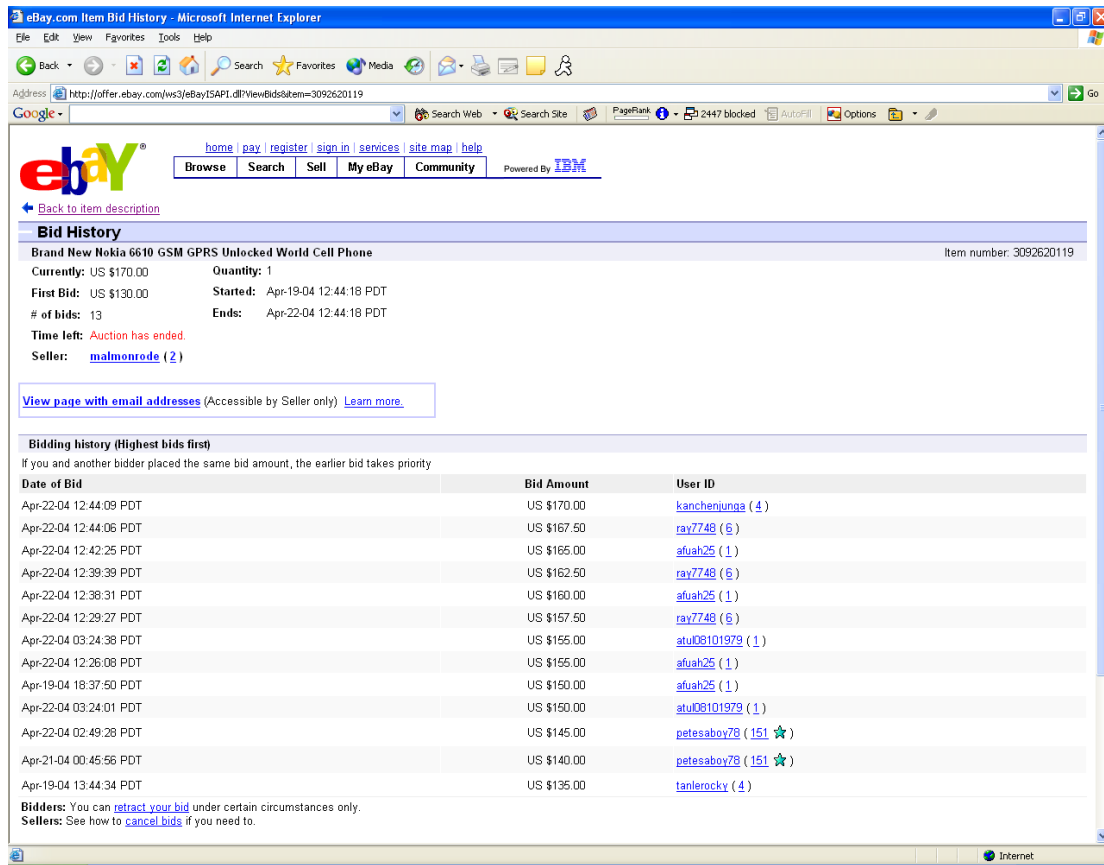


Figure 1 – Almost complete bid history, exception being highest bid

Internet-based field experiments that deal with real bidders in real markets provide a contrast to the controlled environment of laboratory experiments with student subjects. This is evident in the work of David Lucking-Reiley⁵ (1999) and John A. List and Lucking-Reiley (2002). They show how age old questions such as revenue equivalence and the importance of decisions costs respectively, can be practically examined using field experiments with real bidders and without any theoretical assumptions that would be enforced in the laboratory. In the context of the current study, given the lack of published

⁵ Now David Reiley.

high bids, researchers seeking to model consumer surplus on eBay would have to make distributional assumptions about bidders' valuation and subsequently use order statistics theory to estimate surplus. Hasker et al. (2001) demonstrate this on a data set of PC monitors. Our field experiment, by virtue of extracting exact revealed measures of consumer surplus for 28 eBay item categories, can serve to validate theoretical assumptions made by researchers, as well as to inform future analytical work aimed at understanding the dynamics and equilibria of eBay auctions.

Bidders using Cniper.com to bid on their behalf reveal their willingness to pay to the agent. Subsequently, to measure surplus, we extract from eBay the actual price paid, for auctions where our agent wins the auction. This yields a unique dataset, based on which we set out to quantify the level and characterize the distribution of consumer surplus in eBay auctions. In addition, with an objective of shedding new light on online bidding strategies, we also identify and examine the determinants of consumer surplus. Do bidder characteristics such as experience, market characteristics such as number of competing bidders, and mechanism design choices made by sellers such as opening bid, have any influence on consumer surplus? From an econometric perspective, particularly challenging is the fact that our data of 5157 auctions consist of 383 auctions with zero-surplus. The "zero-inflated" dependent variable cannot be directly modeled within an ordinary regression model. We therefore use a preliminary step that allows us to integrate the zero and non-zero surplus auctions into a single regression model. Interestingly, our data is diverse, with significant bidding activity in three important currencies (US Dollar, Great Britain Pound (GBP) and the Euro), in all but two eBay item categories⁶. Thus, we are able to shed light on whether there are cross-continental differences in surplus levels and whether certain groups of categories yield different surplus than others.

Our analysis rests on a key assumption. Even though our dataset is diverse, comprising of items from almost all of eBay's major categories, we *initially* assume a private value setting for our analysis. This permits us to claim, under the independent private values assumption, that the difference between the

⁶ eBay call these major categories "metas" internally. Our data did not contain entries from categories "Travel" and "Tickets."

actual bid made on eBay by the winning bidder and the price paid represents surplus. We are assuming that the highest bid represents the winner's valuation for the item. Put simply, we are assuming that eBay's second-price mechanism induces truth-telling. Using the lens of William Vickrey's (1961) stylized model, eBay's mechanism is a hybrid between an open ascending English auction and a sealed bid second price auction. For such hybrid mechanisms multiple equilibria are likely to exist and are being currently explored [Bajari and Hortascu (2003), Hasker et al. (2001)]. In the case of our data, since by design we limit ourselves to sniping winners, it is straightforward to prove that the absence of any response time to other bidders makes truth-telling sniping a weakly dominant strategy, and the eBay auction resembles a second-price sealed bid auction.

Subsequently, in Section III, we relax the private values assumption and fit a nested model to our data to examine whether, what would normally be considered, common value items have a significantly different surplus level as compared to private value items.

Because our dataset arises from a sniping website, it raises questions about the generalizability of our results to the overall population of eBay auctions, where sniping may or may not occur. We address this empirically by testing whether a randomly drawn validation sample of 1000 eBay auctions has similar distributions of key auction parameters as do our field experiment data. We find that in all auction parameters, including item price, bidder and seller reputations, item categories, number of bidders and opening bid, there is no significant difference between our validation and field data⁷. This, coupled with the academic and practitioner support of the notion that sniping is widespread, leads us to believe that our findings are generalizable across eBay auctions.

A simple analysis of the surplus data yields a sample median of surplus equal to \$3.53. Our results indicate that consumer surplus in online auctions can be well-approximated by a 3-parameter Weibull distribution. This distribution is one of the most popular in reliability engineering, due to its flexibility, and is related to the lifetime of the weakest link in competing risk models. We use the Weibull distribution to evaluate the sampling error in the median surplus and show that a 95% confidence interval

⁷ Appendix A provides the comparative box-plots, indicating no significant difference in any of the variables.

for the median is [\$3.30, \$3.71]. Using the median surplus level of \$3.53, and based on an estimate of 417.5 million⁸ auctions which resulted in a sale on eBay in 2003, we estimate \$1.47 billion accrued consumer surplus for the year 2003. This high and growing consumer surplus level is one reason why online auctions are an attractive retail channel for consumers.

We find that US currency auctions carry higher surpluses relative to Euro and GBP auctions, by a factor of approximately 22%. Surplus levels in Euro and GBP auctions are similar to each other. There appear to be three main groups of surplus categories. The highest surplus is accrued to the group of eBay categories that are antique or collectible in nature. This is followed by a moderate surplus group of items comprised of computers, electronics and books, among others. The lowest surplus is in the group of household items such as toys, health and beauty items and games. We find that sellers with higher feedback ratings, a proxy for experience and trust, tend to yield lower bidder surplus, and that experienced bidders tend to realize higher surplus. We find that the main effects of price, opening bid, and number of bidders have a significant influence on surplus, but so do the interactions of price with opening bid and price with number of bidders. These main effects must therefore be interpreted cautiously. Interestingly, we find that surplus is positively associated with price in auctions with many bidders, but this relationship is moderated by the opening price. We find that surplus is generally positive in auction duration and negative in sniping time, but only for “mainstream auctions” with five to seven day duration and sniping time equal to eight⁹ or nine seconds. We elaborate on these findings, and on how we reach them, in the remainder of the paper.

The next section of this paper describes the working of the bidding agent and the data sample. Section II characterizes the distribution of consumer surplus, provides estimates of levels of consumer

⁸ eBay’s 10k statement available at <http://sec.gov/Archives/edgar/data/1065088/000089161804000676/0000891618-04-000676-index.htm> has a section on operational parameters. They report that a total of 971 million items were listed in 2003. However, a significant percentage of auctions don’t get any bids. Our conversation with a senior eBay executive revealed that eBay treats the overall success rate as confidential information. Thus, to estimate our multiplier we relied on aggregate level secondary eBay data from Hammertap.com . Based on a sample of 14,000 auctions from Hammertap we conservatively find an overall success rate of about 43 percent (lower 95% CI). Thus, we use 417.5 million as our multiplier. Further research is needed on analyzing the level and determinants of auction success rates on eBay.

⁹ The default on cniiper.com

surplus on eBay auctions and explores the determinants of consumer surplus. Section III concludes by pointing out limitations of the study and directions for future research.

I. DESCRIPTION OF THE BIDDING AGENT AND THE DATA SAMPLE

The prevalence of sniping on eBay has led to several independent third party sniping agents that help bidders place last second bids on eBay. The interested reader is referred to Bapna (2003) for a review of sniping agents and their technical details. This study utilizes data from one such agent, Cniper.com, deployed by us. Cniper's logo "Snipe bids in your sleep, for free" is all explaining. Cniper is deployed just two hops away from eBay's server, making bid submissions lightning fast¹⁰. While most competing eBay sniping agents are fee based¹¹, Cniper has always been a free service and has a loyal and steadily growing user base of 2,035 bidders. It relies solely on word of mouth for advertising. In the period beginning July 23, 2003¹² and ending June 24, 2004, Cniper placed 69,571 bids on eBay on behalf of its users. Cniper is developed using PHP¹³ and MySQL¹⁴ and is deployed on an Apache webserver sitting on a Unix box. Thus, Cniper leverages the latest advances in open source software technology to keep its costs low¹⁵. This helps us provide it as a free service, and provides no incentive for any bid shading to account for bidding agent commissions. The lack of commissions also attracts entry for the tool, which in turn provides us with continuously richer observations of real economic agents acting in real markets. We believe that our approach, a first in the research community, will serve as a model for researchers doing Internet based field experiments and so called action research.

While the full technical details of Cniper.com's working are beyond the scope of this paper, a brief overview of its usage is necessary to motivate its usefulness in measuring consumer surplus. Agents

¹⁰ Roth and Ockenfels (2002) emphasize the probability of bids not getting through in the last seconds.

¹¹ BidSlammer.com, AuctionSniper.com etc.

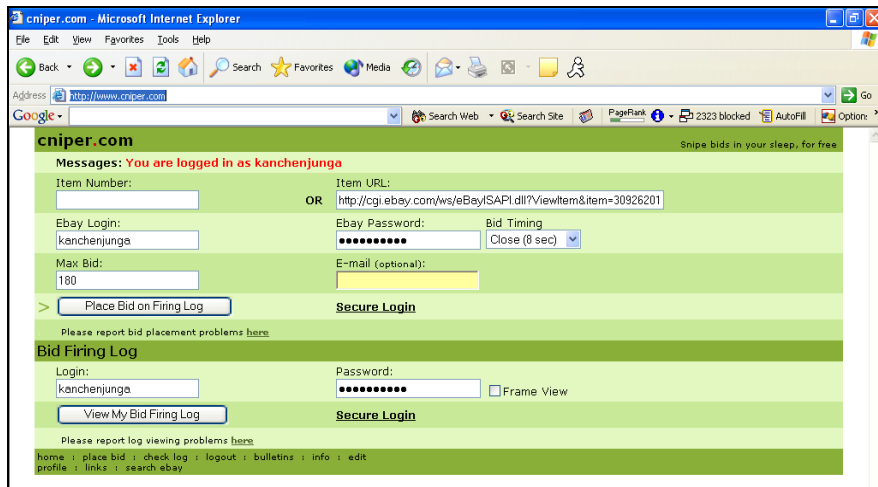
¹² This was when the site was significantly redesigned.

¹³ A fast growing server-side scripting tool.

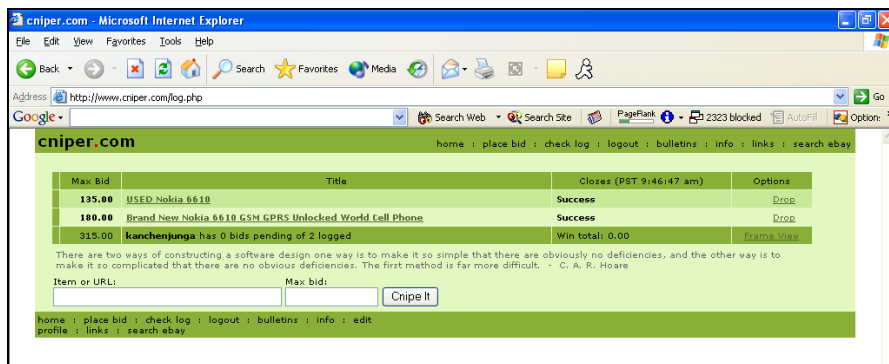
¹⁴ The standard open source relational database.

¹⁵ Cniper has zero licensing fees costs. Its only costs are those of hosting and the first author's time. The latter is bursty and can be significant at times when eBay changes its bid acceptance technology and Cniper has to respond by reprogramming its bid submission protocol.

such as Cniper, allow bidders to reveal a) their willingness-to-pay for a specific item being auctioned on eBay, and b) the number of seconds before the close of the auction they want their bid to be placed on eBay. For illustrative purposes, we continue with our earlier example of the Nokia 6610 GSM phone auction which was sniped and won by eBay user ‘kanchenjunga’ using Cniper and whose bid history is displayed in Figure 1. We show the process of the bidder sniping and the actual winning bid placed in Figure 2. Figures 2a and 2b reveal that bidder ‘kanchenjunga’ fired \$180 for the item and won the auction by outsniping bidder ‘ray7748’ by 3 seconds. Recall, from Figure 1 that the winning bid or price is \$170. Thus, it is evident that bidder ‘kanchenjunga’ derived a surplus of \$10 from this auction.



(2a)



(2b)

Figure 2a – Bidder *kanchenjunga* requests Cniper to bid \$180 for eBay Item 092620119 eight seconds before auction closes

Figure 2b – Row two of *kanchenjunga*’s Cniper log shows the \$180 bid as fired and successful

I. A. Description of the Data

The data used in our analysis consist of 5157 eBay auctions that took place between January 9, 2004 and April 21, 2004¹⁶. In all these auctions the winner was a Cniper.com user. These auctions were carried out in one of three major currencies: US Dollar (USD), Great Britain Pound (GBP), and the Euro. The items auctioned were in a wide variety of categories, spanning most¹⁷ of eBay's 30 high level categories¹⁸. To maintain a minimal cardinality level in each category, we grouped the items, using eBay's categories, into 18 major categories as displayed in Table 1. An additional 19th category was created for items in which the category description was missing¹⁹.

Category	Proportion of Auctions	Median Surplus
Antique/Art/Crafts	2.35%	9.35
Automotive	5.43%	6.08
Books	6.36%	1.60
Business/Industrial	2.13%	4.57
Clothing/Accessories	3.70%	2.20
Collectibles	12.22%	4.00
Computer/Networking	5.20%	4.84
Coins/Stamps	3.61%	2.25
EverythingElse	3.04%	5.80
Consumer Electronics	4.25%	4.26
Health/Beauty	1.84%	1.58
Home/Garden	6.75%	2.16
Jewelry/Watches	4.36%	3.84
Music/Movie/Video Games	10.92%	1.22
Missing	10.59%	3.45
Pottery/Glass	1.07%	8.05
Photography/Camera	3.24%	5.07
SportingGoods	3.66%	4.07
Toys/Hobbies	9.29%	3.12

Table 1 - Break-down of eBay Categories in Data

¹⁶ This corresponds to little more than a three month period, the duration for which eBay posts bid histories of completed auctions.

¹⁷ Only the two categories "Travel" and "Tickets" were not populated in our dataset.

¹⁸ See <http://pages.ebay.com/categorychanges/> for a list of high-level eBay categories. Notice that it includes the category "Everything Else" that contains auctions that do not fit any other classification. Also notice that the category "Automotive" is not contained in this list.

¹⁹ While auctions with missing category descriptions could have been assigned to the "Everything Else" category, we decided to keep these auctions separate in order to maintain objectivity.

To the best of our knowledge, currency and category have not featured in the extant analysis of eBay data. In addition, we recorded from eBay the following information on each auction: Opening and closing prices in their original currency and their USD equivalent²⁰, whether hidden reserve was used, the starting and ending time²¹ and date, the number of bids placed in the auction, the number of unique bidders participating in the auction, and seller and winner rating. From Cniper.com we obtained the number of seconds before the auction close that the winning bid was placed, and the winning bid itself. We then calculated the surplus by subtracting the closing price from the winning bid.

Table 2a describes summary statistics for each of the non-monetary continuous variables. Table 2b splits the summary statistics of opening bid, closing price and surplus according to currency.

	<i>Mean</i>	<i>Std Dev</i>	<i>Minimum</i>	<i>Maximum</i>	<i>Median</i>
Auction Duration (NUM_DAYS²²)	6.66	2.08	1	10	7
Number Bids	6.56	5.41	2	50	5
Number Bidders (NUM_BIDDERS)	3.83	2.63	1	29	3
Snipe Time in seconds (SNIPE_TIME)	8.71	0.89	1	13	9
Winner Rating (W_RATING)	230.24	380.74	-3	11350	111
Seller Rating (S_RATING)	2514.42	8775.54	-1	170889	350

Table 2a – Summary Statistics for Non Monetary Continuous Variables

²⁰ eBay provides approximate conversions on the web page.

²¹ Since eBay auctions last either 1,3,5,7 or 10 days, the starting and ending times are always equal.

²² Upper case corresponds to variable names used later in modeling

	SURPLUS	PRICE	OPENING BID
USD (n=3117)			
<i>Mean</i>	19.19	69.16	24.48
<i>Median</i>	4.49	16.90	6.99
<i>Standard deviation</i>	64.77	236.60	88.79
<i>Minimum</i>	0.00	0.01	0.01
<i>Maximum</i>	1117.48	7600.00	2100.00
EURO (n=1571)			
<i>Mean</i>	6.63	28.81	6.21
<i>Median</i>	2.16	8.30	1.20
<i>Standard deviation</i>	14.36	99.20	51.56
<i>Minimum</i>	0.00	1.18	1.08
<i>Maximum</i>	168.36	2249.98	1953.93
GBP (n=469)			
<i>Mean</i>	43.31	61.14	33.97
<i>Median</i>	3.60	11.75	5.34
<i>Standard deviation</i>	77.12	356.06	253.81
<i>Minimum</i>	0.00	0.9	0.45
<i>Maximum</i>	1355.55	5627.18	3570.50
Total (n=5157)			
<i>Mean</i>	14.92	56.14	19.78
<i>Median</i>	3.53	13.16	3.95
<i>Standard deviation</i>	56.30	220.61	107.28
<i>Minimum</i>	0.00	0.01	0.01
<i>Maximum</i>	1355.55	7600	3570.50

Table 2b – Currency-wise²³ Summary Statistics for Monetary Variables

We next analyze the distribution of surplus levels in online auctions, quantify its level and explore its determinants.

II. ANALYSIS OF CONSUMER SURPLUS

Figure 3 shows a bar chart of surplus of the USD data; the other two currencies are similar. Two interesting observations are worth considering. Firstly, we see that the values \$0, \$0.50, \$1.00, \$1.50 and so on are especially prevalent. Only a very small proportion of surpluses assume values between these values. For instance, while 75 of the surplus values equal exactly \$1, only 10 equal \$0.99, only 6 equal \$0.98 and only 4 equal \$0.97. On the other hand, only 32 values are \$1.01. This is surprising, since we would expect surplus to be distributed uniformly in such small intervals. Thus, our first insight into the

²³ All values are currency adjusted and reported in USD for easy comparison.

distribution of consumer surplus is that the data are apparently *semi*-continuous. By semi-continuous we mean that while surplus can theoretically assume any value in a given interval, some underlying (and unobserved) data-generating mechanism introduces discretization, causing the data to be concentrated on certain values. We explore the possible causes of this later in this section.

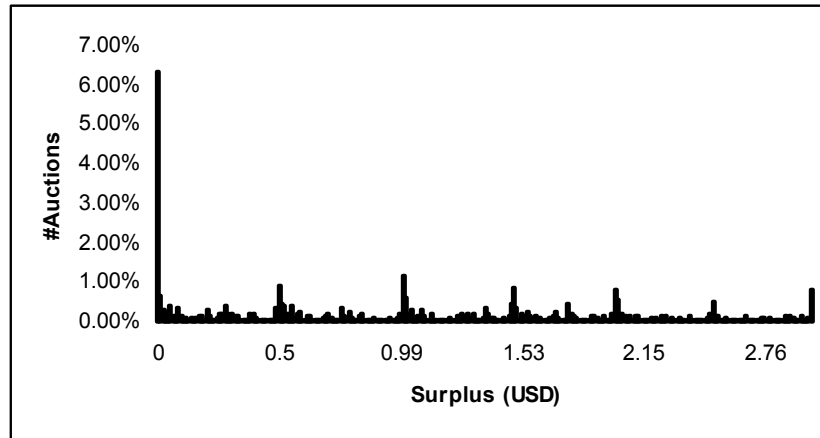


Figure 3 – Bar chart for Surplus in the Range \$0 to \$3

In order to transform surplus from auctions in various currencies into a single scale, we converted all currencies into USD using the conversion rate listed on the auction’s eBay page. Figure 4 displays a histogram of $\log(\text{USD surplus}+1)$. The shift of 1 allows us to apply the log transform to the zero surplus data. The data are clearly bi-modal, with a large “lump” of zero values ($0 = \log(1)$).




Figure 4 – Surplus Distribution is Bi-Modal

Our second observation is that 383 (out of 5157) auctions ended at a closing price exactly equal to the revealed willingness to pay indicated by the winner to Cniper. These can be categorized as highly competitive zero surplus auctions where the two highest valuations are identical. Given that our research objective is to quantify and understand the determinant of consumer surplus, the above “zero-inflated” data present a challenge, in that we cannot directly model surplus in a regression model (which assumes normal residuals).

II.A Finding a Suitable Transformation

Although there are multiple values with high frequencies, the zero-inflation is the most prominent. In order to model such data we can take two approaches: We can transform the data to achieve a more regular distribution, or develop a model that can account for such a data structure. An example would be a mixture model that results from combining a process that generates only special values such as zeros with another process that generates non-negative or positive surpluses. The first approach, of transforming the data, is more popular because it is simpler to use and understand. It requires fewer assumptions and fewer parameters, and enables to make use of a large variety of established statistical methods. The only challenge is to find a good transformation. It turns out that ordinary transformations do not yield satisfactory results. For instance, the Box-Cox transformation indicates that the log transformation brings the data closest to normality. However this does not alleviate the heavy zero-inflation and the resulting data are clearly not normal. Furthermore, we would like the transformation to deal with the other (non-zero) high-frequency values as well. Our solution is to use a novel transformation which moves the data into a slightly coarser scale. In particular, we round surplus to the next integer value. The reasoning behind this transformation is data driven: Recall the initial observation about the semi-continuous nature of the data, with large counts of surpluses at values of 0, 0.01, 0.5, 1, etc. We believe that the reason is most likely the bid increments that eBay imposes²⁴. It could also be the case that most bidders enter integer values as their bids, and this further contributes to semi-continuous nature of the data. Reasoning that surpluses within a \$1 unit range carry the same (or at least similar) information,

²⁴ See Bapna et al (2003b) for the impact of discrete bid increments on online auctions

we transform the original surplus data by applying the ceiling function (which gives the next highest integer). We then take $\log(\text{integer surplus} + 1)$ to accommodate for the zero surplus data. Figure 6 shows the histogram of the transformed integer surplus values. The new values follow a unimodal, right-skewed distribution.




Figure 5 – (log) Ceiling transformed Surplus is right-skewed

Further investigation of the probabilistic structure of the data reveals that a 3-parameter Weibull distribution approximates the transformed data fairly well. This can be seen in Figure 6 which displays a weibull probability plot of the data (left), and in comparison a lognormal probability plot of the same data (right).




Figure 6 – Three parameter Weibull distribution fits better than a Lognormal distribution

This approximation is of special interest, since the Weibull distribution often describes the time to an occurrence of the “weakest link” in competing failure processes; for example, the time to death of the first component in a system with multiple components and competing failures. In our case we have a set of bids in an auction. Among the differences between each of the bids and the winning price, the surplus is the only non-positive value, and thus the minimum of these differences. In this sense the winning/highest bid is the “weakest link” because it has the smallest distance from price compared to all other bids that were placed in the auction. Taking logs maintains the order, and therefore the result holds.

The estimated parameters of the fitted Weibull distribution are shape = 1.69, scale = 2.20, and location/threshold = -0.18. Since the surplus distribution is very skewed, the median is a better measure than the mean for the center of the distribution. Although we can compute the median surplus of \$3.53 directly from the original values (converted to USD), it is advantageous to have a measure of sampling error. We use the asymptotic normality [Herbert Aron David and Haikady Navada Nagaraja, (2003), page 241] of the sample median and the Weibull distribution to obtain the formula for the median’s standard deviation. After transforming the data back into its original units we obtain a 95% confidence interval of [\$3.30, \$3.71]²⁵. We also obtain a non-parametric confidence interval by taking bootstrap samples from the raw surplus data. A 95% bootstrap confidence intervals (based on 1000 replications) yields [\$3.36, \$3.74]. The closeness of these intervals indicates that the transformation maintains the main features of the data. Using the median \$3.53 value, and multiplying by the estimated 417.5 million transactions on eBay for 2003, we estimate an overall surplus of \$1.47 billion accrued to eBay bidders last year.

²⁵ The median is asymptotically normal, centered around the population median and with variance $[4nf^2(Med)]^{-1}$, where $f(Med)$ is the Weibull density at the median [David and Nagaraja (2003), page 241]. We estimate this quantity using the estimated Weibull parameters. This yields a standard deviation of 0.021 for the median of $\log(\text{integer-surplus})$. A 95% CI for the population median of $\log(\text{integer-surplus})$ is [1.57, 1.65]. Taking an exponent and subtracting 0.5 (the average rounding) yields the 95% CI for median surplus of [\$3.30, \$3.71].

II.B. Determinants of Consumer Surplus

Prior empirical research dealing with consumer surplus in online auctions has made relative comparisons of groups of bidders in discriminatory multi-unit Yankee auctions [Bapna et al 2003a, 2003b, Bapna et al. (2004)]. Their main finding is that *evaluators*, bidders who make singleton early high bids, leave significantly more money on the table in comparison to other groups of bidders. To the best of our knowledge, there has been no other study that has looked at determinants of absolute consumer surplus levels. In this paper, we work with the ceiling transformed surplus data and identify candidate explanatory variables that could influence the level of bidder surplus. It should be noted that much of the prior research has focused on understanding the determinants of auction price. Bidder surplus, of course, is intrinsically associated with price, albeit negatively. We believe that the following five categories of independent variables potentially influence the rents that bidders can extract.

1) *Seller's mechanism design choices*: Sellers, who strive to maximize their revenues, can be expected to strategize on eBay by choosing the appropriate combination of opening bid level, auction duration and the usage of a hidden reserve price. Opening bid can be interpreted as an open reserve price and prior research has contrasted the comparative effectiveness of open versus hidden reserve prices on sellers' expected revenue. Myerson (1981) formulates the optimal (seller revenue maximizing) auction design problem as being equivalent to deriving the optimal open reserve price. Rama Katkar and Lucking-Reiley (2000) in a field experiment selling Pokemon cards, find that hidden (secret) reserve prices make sellers worse off, 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. In contrast, Bajari and Hortaçsu (2003) based on an econometric estimation, suggest that optimally chosen hidden reserve prices can yield the seller 1 percent higher revenues. Thus, the evidence seems mixed with respect to how the seller's usage of hidden reserve prices impacts the auction price, and consequently the bidders' surplus.

It is well established that, on eBay, lowering opening bids attracts more bidders [Bajari and Hortaçsu (2003), Lucking-Reiley et al. (2000)]. In addition, prior research also suggests that when a seller chooses to have her auction last for a longer period of days, this significantly increases the auction price

on average [Lucking-Reiley, Bryan, Prasad, and Reeves (2000)]. We describe the number of competing bidders under market characteristics and examine its interaction with opening bid. Taken together, the prior research²⁶ allows us to hypothesize:

H1: Opening bid has a negative influence on bidder surplus

H2: Auction duration has a negative influence on bidder surplus

2) *Seller characteristics*: eBay's feedback reputation system has been widely studied and studies indicate that sellers with higher reputations engender trust and extract premiums [Sulin Ba and Paul A. Pavlou (2002)]. It can also be argued that sellers with more experience²⁷, also proxied by feedback ratings, make better mechanism design choices to maximize expected auction price. Thus, we hypothesize:

H3: Seller reputation has a negative influence on bidder surplus

3) *Product characteristics*: For homogenous products, by definition surplus increases with a decrease in price. However, in contrast to the above-mentioned empirical studies on eBay [Bajari and Hortaçsu (2003), Katkar and Lucking-Reiley (2000), Lucking-Reiley et al. (2000)] that controlled for product heterogeneity, our dataset is diverse, covering all but two of eBay's 30 major item categories, with prices ranging from 1 cent to \$7600. This allows us to test the implications of stakes and product attributes on surplus levels in a far more generalizable setting. Smith and Walker (1993) have predicted that individuals' behavior will more closely match the predictions of rational behavior as the stakes of the decision increase²⁸. Marketing theories suggest that as stakes get higher, consumers get more involved in finding the best price for their product [Marcel Cohen (2000)]. Based on the mixed evidence regarding the influence of price and in light of our diverse dataset we do not specify a directional hypothesis. Instead, we hypothesize:

H4: Item price has a significant influence on bidder surplus

²⁶ Prior research typically considers only main effects and our hypotheses reflect this. As our detailed analysis will show though, some of the main effects are moderated by interaction terms.

²⁷ eBay's feedback rating, which indicate the difference between positive and negative ratings, have been viewed as indicators of experience as they are generally reflective of the number of transactions conducted by the seller. Paul Resnick and Richard Zeckhauser (2001) report that only 0.6% of feedback comments left on eBay by buyers about sellers was negative or neutral.

²⁸ In the context of auctions this maps to surplus maximizing behavior.

With respect to item categories, while we assume the independent private values setting to establish truth-telling as a dominant strategy for second-price auctions, it is well-known that most items are more completely viewed under the general auction model [Paul R. Milgrom and Robert J. Weber (1982)]. This model accommodates private value components in what would traditionally be categorized as common value auctions and vice versa. This belief is based on practical limitations of observing data from online auctions and subsequently categorizing it in the context of the assumptions of the private and common value models. Issues such as condition of the good, bidder's confidence in the expressed condition of the good, the degree of expertise required in assessing the market value of the good, possibilities of easy resale in electronic markets, the value of private consumption of the good, as well as hedonic aspects of outbidding one's rivals confound any priors with respect to item category wise surplus levels. We will thus let the data speak.

Later in the Results section (III), we relax the private values assumption and fit a nested model to our data to examine whether, what would normally be considered, common value items have a significantly different surplus level as compared to private value items.

4) *Winner characteristics*: We collect information on the winning bidder's rating as a proxy for their experience on eBay. We expect more experienced bidders to have more confidence in their valuations and bids and hence we expect them to derive higher surplus. In addition, we also measure aggressiveness of the bidder by looking at how many seconds prior to the close of the auction they snipe. Given that there is a small likelihood of bids not getting through due to congestion on eBay, even when fired by agents such as Cniper²⁹, we expect risk seeking behavior with a greater desire to win and consequently higher bid levels and surplus. Thus, we hypothesize:

H5: Winner experience has a positive influence on bidder surplus

H6: Winner aggressiveness has a positive influence on bidder surplus

5) *Market characteristics*: We are fortunate to have significant data in three prominent currencies, namely USD, GBP and the EURO. This allows us to test, for the first time, whether bidders differ in their bidding

²⁹ Cniper offers a choice of as low as 2 seconds before closing, but calls it "insane."

characteristics across countries. While there has been prior research in looking at the efficiency of auction formats across different countries [Klemperer (2002)], this study represents a first in comparing bidder surplus levels across countries. Given that eBay was founded in the US and subsequently expanded to UK and Europe, we hypothesize that:

H7: The longer experience of the US market has a positive influence on bidder surplus

Lastly, as can be expected, we propose that the level of competition in an auction, reflected in the number of bidders, negatively influences the surplus accrued to the winning bidder. We also have data on the total number of bids in an auction. However, total number of bids is highly correlated with the total number of bidders and thus does not add any additional explanatory power to our regression models. Thus, we hypothesize that:

H8: Increasing competition has a negative influence on bidder surplus

II.C Modeling Approach

Our goal is to find a model that explains the effect of the above-mentioned set of explanatory variables on consumer surplus. The main challenge with surplus data such as ours is the extremely large number of auctions resulting in “special values” and especially zero-surplus. That is, auctions which ended exactly at the value placed by the winner. Recall that in our dataset, 383 auctions (from a total of 5157) carried zero surplus. Furthermore, these auctions are diverse in their other characteristics (across categories, currencies, etc.). Such data cannot be modeled directly by ordinary regression without violating the normality assumption. As mentioned before, two possible solutions are to transform the data and then use ordinary statistical models or to develop a specialized model that captures this behavior. Although “zero-inflated” models exist for discrete data (e.g., “zero-inflated Poisson”, Lambert, 1982), we are not aware of such models for continuous data, especially when the inflation is not only for zeros but also for other values that cannot be easily predetermined. In order to fit such a model, we need to specify the special values and their probabilities, the mixing distribution, and the general surplus distribution.

This results in many parameters to be estimated. In addition, the mechanism that generates the special values is not observable, and thus specifying “special values” based on the sample values alone can be misleading. If we make a simplifying assumption that only zero-surplus auctions should be separated from positive-surplus auctions, then we could use a two-stage model for price as a function of our hypothesized determinants: The first stage finds variables that are useful for classifying an auction as a zero- or positive-surplus auction. This can be done by using a classification method, such as logistic regression. The second stage fits an ordinary regression model to positive-surplus auctions, in order to find variables that determine surplus. While this two-stage approach is valid, one of its disadvantages is the loss of statistical power. Also, it is less elegant than a single integrated model, and requires the estimation of multiple parameters. We therefore strived for an alternative approach that combines zero-surplus and positive-surplus (including the other high-frequency special values) under one roof.

As described in Section II, the rounding of surplus values and taking a log transform yields a distribution that is suitable for regression modeling. We fit a regression model to the transformed data based on our hypotheses above and using model selection procedures. The best model is given in Table 3³⁰. Note that among the candidate predictors identified in Section II.A the usage of the secret reserve price, the number of bids, and the auction’s time did not add additional explanatory power to the model. The model’s R^2 and Adjusted R^2 values are 0.3188 and 0.3158, respectively. In comparison, the best model *without interaction terms* had corresponding values of 0.3029 and 0.3004. Price has the highest explanatory power, resulting in an R^2 value of 0.2500 for the simple, single-variable regression model. It is also interesting to note that when using only the “standard” eBay explanatory variables,³¹ opening Bid, winner and seller rating, number of bidders, duration of the auction and price, the best model fit amounts to only 0.2678 and 0.2668 for R^2 and Adjusted R^2 .

Notice that the model above uses $\log(\text{surplus})$ as the response variable and $\log(\text{price})$ as a predictor. This is mathematically equivalent to a model where the response is $\log(\text{surplus}/\text{price})$, or \log

³⁰ This model is very similar to the results obtained from a two-stage model, as described above.

³¹ We say standard since many other empirical studies do not consider currency or ebay categories.

relative surplus, with the predictor $\log(\text{price})$, where the coefficient for $\log(\text{price})$ is equal to $\beta_{\text{price}} - 1$ (we are subtracting $\log(\text{price})$ on both sides of the equation). The coefficients for other predictors are the same as in the first model. The practical difference is in the interpretation that now looks at the relative surplus as a function of the item's price. We describe our results on the relationship between surplus and price in both terms below.

Variable	Coefficient	SE	Pvalue
Intercept	2.51	0.52	<.0001
Categories*			
Antique/Art	0.41	0.10	<.0001
Pottery/Glass	0.28	0.07	0.00
Collectibles	0.41	0.05	<.0001
EverythingElse	0.38	0.09	<.0001
Toys/Hobbies	0.33	0.08	<.0001
Music/Movie/Games	0.39	0.15	0.01
Jewelry	-0.30	0.12	0.01
Automotive	-0.24	0.06	0.00
Home/Garden	-0.26	0.05	<.0001
Health/Beauty	-0.16	0.06	0.00
US Dollars**	0.20	0.04	<.0001
NUM_DAYS	-0.15	0.07	0.03
SNIFE_TIME	-0.23	0.06	0.00
NUM_BIDDERS***	-0.52	0.05	<.0001
PRICE***	0.36	0.03	<.0001
S_RATING***	-0.03	0.01	0.00
W_RATING***	0.03	0.01	0.02
OPENING_BID***	-0.17	0.02	<.0001
OPENING_BID x PRICE	0.04	0.00	<.0001
PRICE x NUM_BIDDERS	0.09	0.02	<.0001
NUM_DAYS x SNIFE_TIME	0.02	0.01	0.01

* Base Category: Books, Business/Industry, Clothing/Accessories, Computer, Coins/Stamps,

Electronics, Photography, Sporting Goods

** Base category: Euros and GBP

*** The variables surplus, price, opening bid, winner rating, seller rating and number of bidders were transformed to the log-scale

Table 3 - Parameter Estimates for Log-Surplus

III. ANALYSIS OF RESULTS

We organize this section by reporting first on the main effects and subsequently on those variables that have significant moderating interactions. We begin with the effect of the categorical variables of currency and product categories.

Currency: Consumer surplus is not the same across different currencies. We see that US currency auctions carry a significantly higher surplus relative to that of EURO and GBP auctions. In fact the coefficient equals .20 which implies that this difference is $\exp(.20) = 1.22$. The surplus accrued to US bidders is about 22% higher than in the European market! We believe that one source of these difference is the longer experience the US has had with electronic markets. eBay in the U.S. is, arguably, more mature than it is in Europe. Therefore, the higher consumer surplus may be explained by the fact that there is greater seller participation shifting the supply curve to the right and effectively lowering supplier surplus and increasing consumer surplus and social welfare. Greater experience with electronic markets may also translate to greater confidence in bidding. There are potentially other cultural factors that need further exploration.

Product Categories: Three main groups of categories determine the surplus level:

- *Highest Surplus:* Antiques/Art, Automotive, Collectibles, Everything Else, Jewelry and Pottery/Glass have the highest surplus by an estimated factor of 1.3-1.5 relative to the reference group.
- *Moderate Surplus:* Computers, Electronics, Books, Clothing/Accessories, Coins/Stamps, Photography, Sporting Goods, Business & Industry and missing category descriptions. These categories comprise the reference group.
- *Lowest Surplus:* Toys/Hobbies (TH), Music/Movies/Games (MG), Home/Garden (HG), and Health/Beauty (HB) have the lowest surplus, by an estimated factor of 0.75-0.85 relative to the reference group.

The highest surplus group appears to be made up of items, such as antiques and jewelry, which require expertise in assessing valuations and could potentially also have a common-value component. Thus, there is a potential winner's curse effect revealed here. The moderate and low surplus groups are made up of items that have several alternative channels, such as eBay's half.com and Amazon.com, and are therefore very competitive. In order to compare our results with other common/private value categorizations, we fit a nested model that uses the common/private value categorization as the nesting variable. Based on guidelines on what constitutes a common/private values setting suggested in Klemperer (1999)³², we tag each of the categories in our dataset as either being common value (antiques, art etc) or private value (goods deemed for private consumption; for instance computer and electronic equipment). Although the nesting variable is statistically significant (indicating a difference in surplus between these two groups), the category-level coefficients remain unchanged. Furthermore, their signs indicate that the 3-class groupings found in the original model is still prominent (see Appendix A for comparison of results from the two models). A key finding from the nested model is that in general common value items have significantly higher surplus levels as compared to private value items. To that extent, our analysis, made under the assumption of private values, paints a conservative view of consumer surplus levels. The nested model also strongly suggests that the categorization of what is private and what is common value is a non-trivial task. Our empirical analysis (see Appendix A for a detailed reasoning as to why this is evident) indicates most items are more completely viewed under the general auction model [Milgrom and Weber (1982)] that has components of both private and common value.

Next, we report on the effect of seller and winner experience, as reflected in their total feedback ratings on eBay.

Winner rating: More experienced winners are associated with a higher surplus. This indicates that with increasing experience bidders act more strategically, identifying winning bid levels and maximizing their utility from participating in eBay.

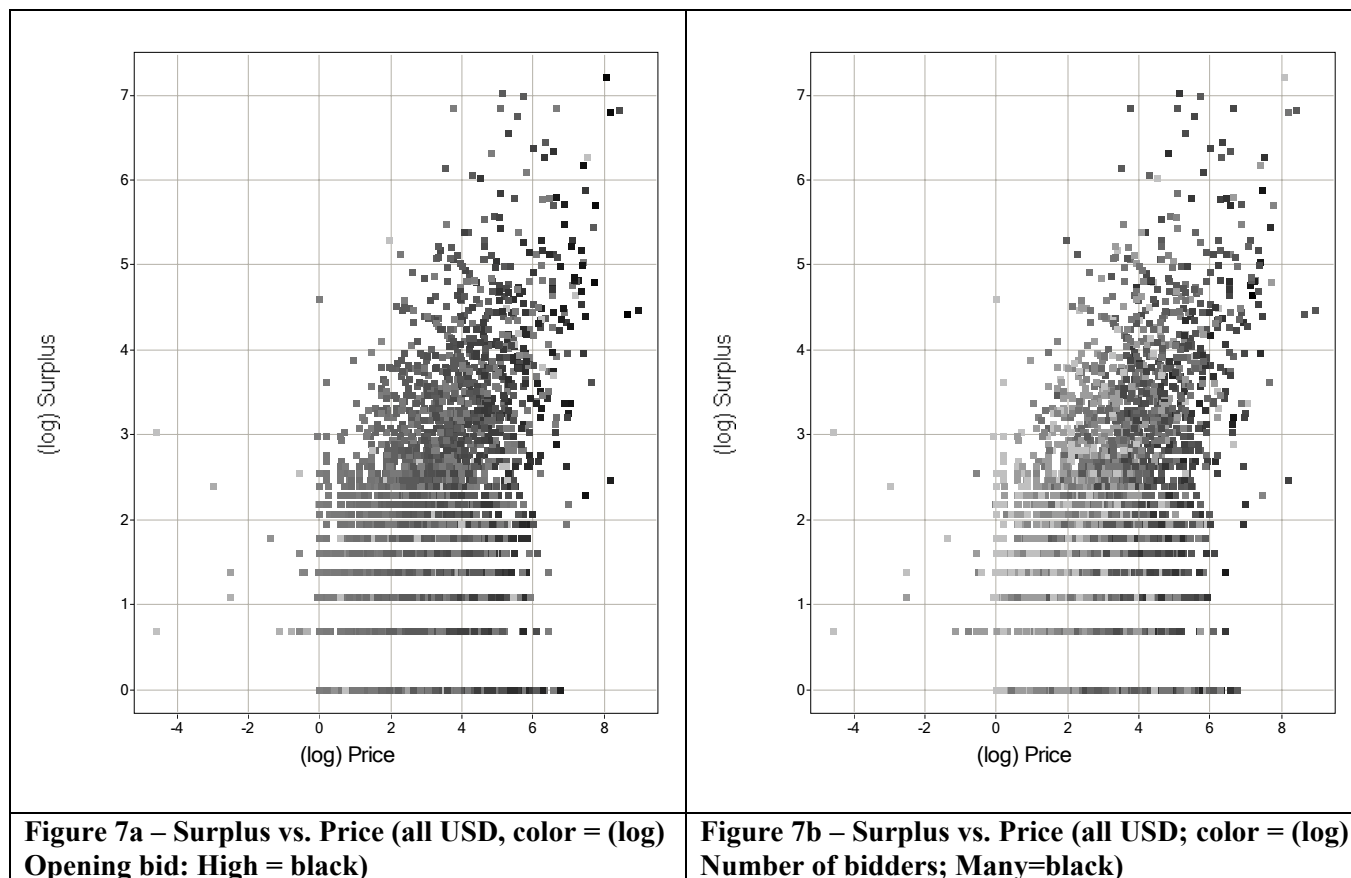
³² The key test (one that has not yet been empirically tested or even operationalized) to detect a common value component is whether a bidder would be inclined to revise her valuation based on some information about her competitors' valuations. If not, it's safe to assume a private values setting.

Seller rating: In contrast, and as expected, a seller's experience has a negative relationship with surplus. It can be expected that more experienced sellers make better mechanisms design choices. In addition, they can be expected to be more sophisticated in describing the product, by choosing a better layout and pictures to attract more bidders. This enables them to extract rents away from bidders. It is also likely that bidders simply trust more in sellers with a better reputation and this translates to more confidence in bidding. Overall, this effect is consistent with previous research [Ba and Pavlou (2002)] that indicates that higher seller reputation results in a higher price!

We now report the main effects that are moderated by interaction terms amongst them.

Price, Number of Bidders, and Opening Bid: The main effect of price is positive, while the main effect of number of bidders and opening bid is negative. It is interesting to note that all of these three main effects are statistically significant, but so are the interactions of price with opening bid and price with number of bidders. These main effects must therefore be interpreted cautiously. The relationship between price and opening bid and its effect on surplus is visible in Figure 7.

This is a scatterplot of surplus vs. price, with color denoting opening bid (continuous grayscale; gray = low, black = high). Overall, surplus is positively correlated with price, but the extent of this relationship depends on the opening bid amount: It is stronger in auctions with higher opening bids and weaker in auctions with low opening bids.



The relationship between price and number of bidders and its effect on surplus can be seen in Figure 7b, which is a scatterplot of surplus vs. price, with depth of grayscale denoting number of bidders (gray = few bidders, black = many bidders). The positive relationship between surplus and price appears to be stronger in auctions with many bidders, as can be seen by the slope of the cloud of grayish-black to black auctions. In contrast, in auctions with few bidders, the relationship is weaker (gray color has lesser slope). The reason for this could be increased competition, which in turn causes the winning bidder to be less cautious, thereby generating higher surplus.

If we look at $\log(\text{surplus}/\text{price})$, the relative surplus as the response, then the coefficient for price is negative ($0.36 - 1 = -0.64$). Considering the number of bidders (between 1-29 in our data) and the range of opening bids (\$0.01-\$3570.50) we get a negative coefficient tying relative surplus with price. This means that higher value items accrue lower relative surpluses. Table 4 shows this effect by giving the

average relative surplus for different price ranges. Notice that the interval lengths increase exponentially (reflecting the log scale).

Price range (in USD)	Average relative surplus (surplus/price)
0.01-0.03	1037.5
0.03-0.07	199
0.07-0.20	24.31
0.20-0.55	7.98
0.55-1.48	3.61
1.48-4.00	1.04
4-11	0.69
11-30	0.46
30-81	0.43
81-220	0.29
220-600	0.197
600-1,628	0.141
1,628 – 4,424	0.146
4,424 - 12,026	0.013

Table 4: Average relative surplus for items in different price ranges

To summarize, the regression model reported in Table 3 indicates a significant positive coefficient for the log-price (even when considering the interaction terms). Since the response is log-surplus it implies that a percentage increase in item price is associated with a percentage increase in surplus (ranging between 17.5-99% depending on the number of bidders and opening price). In terms of relative surplus, a percentage increase in price is associated with a percentage decrease in relative surplus. Interestingly eBay’s gross merchandise sales grew by 160%, from \$14.9 billion in 2002 to 23.8 billion in 2003. This indicates surplus should have gone up by approximately 30%-160% in the same period, indicating growing benefits to consumers participating in online auctions.

Our examination of the interaction effects reveals that the main effect of price has to be qualified by saying that overall surplus is positively associated with price, and this relation is stronger in auctions with higher opening bids. In addition, the rate of this surplus increase changes for auctions with different competition levels. Auctions with higher competition levels, in which there are larger number of bidders, see a faster surplus increase than auctions with a low level of competition.

The opening bid plays an interesting role in influencing surplus. Its main effect is negative, that is, higher opening bids are associated with lower surpluses. But also notice again the positive interaction effect with price: A higher opening bid results in a faster rate of surplus increase due to price. A closer inspection of the data reveals the explanation for this result. For low valuation items, setting the opening bid to something different from eBay's default (= \$0.01) results in a negative correlation with surplus. On the other hand, high valuation items, which are likely to have uncertain valuations, have a positive correlation with the opening bid. In this case bidders extract information from the opening bid about their own valuation for the item; which appears not to be the case for low-valuation items.

Duration and Sniping Time: Auction duration and sniping time (the number of seconds before the close of the auction that the bid fired), both have negative main effects. Overall, longer auctions and early firing times generate lower surplus. These results are as expected: Longer auctions result in higher prices and thus in a lower surplus. Very late firing times appear to be indicative of confident and perhaps risk-seeking bidding behavior.

However, the significant interaction between the two effects, which might be a result of the discrete nature of auction duration (=1,3,5,7,10 with high concentration at 7,10) and fire time (=2,4,8,9,10,..24 with high concentration at 8-9) tell us a more complicated story. The scatterplot in Figure 8 plots sniping time vs. duration with size representing surplus. It is apparent that surplus increases in auction duration for sniping time equal to eight and nine seconds but not for other values. It also appears that in five and seven day auctions surplus decreases for earlier firing times, but no such relation appears for other durations. In summary, surplus is generally positive in duration and negative in sniping time, but only for "mainstream auctions" with five/seven day duration and sniping time equal to eight³³ or nine seconds.

³³ The default on cniper.com




Figure 8 – Sniping time v. Auction Duration (size proportional to Surplus)

IV. CONCLUDING REMARKS and FUTURE RESEARCH

eBay auctions are at the forefront of e-commerce, demonstrating how the Internet can remove spatial and temporal constraints to make economic exchange mechanisms, such as auctions, mainstream. In 2003, sellers through eBay sold \$23.8 billion worth of merchandise. While gains from such trades to sellers and eBay, the market-maker, are obvious, little is known about consumer welfare levels accrued to buyers. Surprisingly, this anomaly extends beyond electronic markets to traditional markets, with the only other studies reporting consumer surplus levels being the analysis of new product offerings in electronic markets by Brynjolfsson, Hu and Smith. (2003), and the study showing significant consumer welfare gains from secondary markets for books [Ghose, Smith and Telang (2004)]. We attempt to fill this gap by studying data obtained from an ongoing field experiment that allows eBay bidders to use a web based sniping agent called Cniper.com for free. In return, roughly 2000 bidders who regularly use Cniper.com

to place last second bids on eBay, vote with their dollars and provide us with a wealth of information about actual willingness to pay for a wide variety of items sold on eBay. Keeping the service free attracts entry for the bidding agent, and provides no incentive for any bid shading to account for bidding agent commissions. Coupled with the fact that these bids are last second by definition, they are not expected to be retaliated against, which further encourages truth-telling valuation revelation behavior on part of the bidders. Thus, for those bids placed by Cniper that win the auction, the difference between the amount bid by the bidder and the auction's closing price represents an exact measure of consumer surplus accrued to the winner of that auction.

Our analysis of winning bids from 5157 auctions spanning all but two of eBay's categories reveals interesting insights into the distribution of consumer surplus. We find that eBay's discrete bid increments tend to make surplus data multi-modal and semi-continuous, rather than continuous. In addition, we observe zero-inflation in 383 auctions, indicating a high degree of competitiveness on eBay. For an auction to have zero surplus, the two highest bids have to be exactly equal. These issues, coupled with data in three currencies and multiple categories catalyze us to innovatively transform our data using a log-ceiling transformation to make it suitable for regression modeling. We expect that our finding that a 3 parameter Weibull distribution best typifies consumer surplus to be useful to future researchers making distributional assumptions for analytical models examining the dynamics of online auctions.

We find that the price of the auction has the biggest positive influence on the bidder's surplus. This reinforces the belief that stakes matter, and as stakes get higher rationality (surplus maximizing behavior) becomes more prevalent. In addition we find that surplus levels are highest for items such as collectibles and antiques, and US bidders accrue significantly higher surplus than their European counterparts. The 22% difference between the more mature U.S. and non-U.S. auctions raises interesting propositions regarding the long-term trends we can expect with consumer surplus. This result suggests that as electronic markets, such as eBay, mature consumer surplus can be expected to be increasing over time. In this regard, it is interesting to consider the possibility that supplier surplus is shrinking as there is increased participation. How these longitudinal patterns will shape out and how they will impact buyer

and seller entry promise to remain interesting questions for future research. Such a study could verify (or dispute) what we suspect; that the increases in consumer surplus are large enough to ensure an overall increase in social welfare.

We find that sellers' mechanism design choices, namely the choice of opening bids, auction duration as well as the sellers experience all have a significant influence on bidder's surplus. So do bidder characteristics such as experience and aggressiveness, as indicated by how close to the auction's end they select to fire their bid.

While we attempt to check the generalizability of our results by examining potential distributional differences in any of the independent variables between our field data and a randomly selected test data of 1000 auctions, we realize that this test data set is only a small fraction of the auctions conducted by eBay on a given date. We are limited in this aspect by having only the public access to eBay's search engine, which is clearly not designed for such a purpose.

We believe that our findings about consumer surplus levels are conservative. This belief stems from our conservative estimate of the number of transactions completed on eBay and also from our assumption of a private values setting. When we relax the private values assumption and fit a nested model to our data that tags categories as being either private or common value, we observe that common value items have significantly higher surplus levels. The nested model also strongly suggests that the categorization of what is private and what is common value is a non-trivial task. This presents a promising area for future research.

As in traditional consumer theory, having benchmark consumer welfare levels is useful in examining the benefits to society of future policy changes that may be introduced by eBay. Our study estimates that in the year 2003 alone, eBay online auctions contributed \$1.47 billion in consumer surplus. We believe that these high and growing consumer welfare levels reveal why online auctions are increasingly popular.

REFERENCES

1. Adams, Christopher P. and Vogt, William (2004). "Digital Demand: Demand for digital cameras on eBay," available at www.ios.neu.edu/papers/s4g3.pdf
2. Adams, Christopher P. and Bivins, Laura L. (2003). "Focusing on Demand: Using eBay data to analyze the demand for telescopes," available at <http://www.ftc.gov/be/workpapers/wp257.pdf>
3. David, Herbert Aron; Nagaraja, Haikady Navada (2003), *Order Statistics*, 3rd edition, John Wiley.
4. Ba, Sulin and Pavlou, Paul A. (2002). "Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior," *MIS Quarterly*, 26(3): 269-289.
5. Bajari, Patrick and Hortaçsu, Ali (2003). "The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions", *Rand Journal of Economics*, 3:2, pp. 329-355.
6. Bajari, Patrick and Hortaçsu, Ali (2004). "Economic Insights from Internet Auctions," forthcoming in *Journal of Economic Literature*.
7. Bapna, Ravi; Goes, Paulo; Gupta, Alok and Jin, Yiwei (2004). "User Heterogeneity and its Impact on Electronic Auction Market Design: An Empirical Exploration," *MIS Quarterly*, 28:1, pp. 21-43.
8. Bapna, Ravi (2003). "When snipers become predators: can mechanism design save online auctions?" *Communications of the ACM*, 46:12, pp. 152-158
9. Bapna, Ravi; Goes, Paulo; Gupta, Alok (2003a). "Replicating Online Yankee Auctions to Analyze Auctioneers' and Bidders' Strategies," *Information Systems Research*, 14:(3), 244-268.
10. Bapna, Ravi; Goes, Paulo; Gupta, Alok (2003b). "Analysis and Design of Business-to-Consumer Online Auctions" *Management Science*, 49:(1), 85-101
11. Brynjolfsson, Erik; Hu, Y. and Smith, Michael D. (2003). "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 49:11, pp. 1580-1596.
12. Cohen, Marcel (2000), "Consumer Involvement – Driving up the Cost," *Consumer Policy Review*, 10(4): 122– 125.

13. Diamond, Peter A. and Hausman, Jerry A. (1994). "Contingent Valuation: Is Some Number better than No Number?" *The Journal of Economic Perspectives*, Vol. 8, No. 4. (Autumn, 1994), pp. 45-64.
14. Ghose, A., M. Smith, R. Telang. 2004, "Internet Exchanges for Used Books: An Empirical Analysis of Welfare Implications," NYU Working paper available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=584401 .
15. Hasker, Kevin; Gonzalez, Raul and Robin C. Sickles (2001), "An Analysis of Strategic Behavior and Consumer Surplus in eBay Auctions," Rice University working paper, available at <http://www.ruf.rice.edu/~rsickles/paper/auction.pdf> .
16. Hausman, Jerry A. (1981). "Exact Consumer's Surplus and Deadweight Loss," *The American Economic Review*, Vol. 71, No. 4. (Sep., 1981), pp. 662-676.
17. Jank, Wolfgang and Shmueli, Galit (2003). "Dynamic Profiling of Online Auctions using Curve Clustering", University of Maryland working paper, available at <http://www.smith.umd.edu/ceme/statistics/AuctionProfiling.pdf>
18. Katkar, Rama and Lucking-Reiley, David. (2000). "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment," working paper, University of Arizona.
19. Klemperer, Paul (2002). "How (Not) to Run Auctions: the European 3G Telecom Auctions," forthcoming in *European Economic Review*.
20. Klemperer, Paul (1999), "Auction Theory: A Guide to Literature", *Journal of Economic Surveys*, 13(3), pp. 227-86.
21. Lambert, Diane (1992). "Zero-Inflated Poisson Regression, With an Application to Defects in Manufacturing", *Technometrics*, vol. 34, pp. 1-14.
22. List, John A. and Lucking-Reiley, David (2002) "Bidding Behavior and Decision Costs in Field Experiments" *Economic Inquiry*, vol. 40, no. 44, pp. 611-619
23. Lucking-Reiley, David (1999). "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet" *American Economic Review*, vol. 89, no. 5, pp.1063-1080.

24. Lucking-Reiley, David (2000), "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics*, vol. 48, no. 3, pp. 227-252.
25. Lucking-Reiley, David; Bryan, Doug; Prasad, Naghi, and Reeves, Daniel (2000), "'Pennies from eBay: the Determinants of Price in Online Auctions,'" working paper, University of Arizona.
26. Marshall, A. (1936), *Principles of Economics*, Eight edition, London, pp. 124-133.
27. McAfee, R. Preston and McMillan, John (1987), "Auctions and Bidding," *Journal of Economic Literature*, Vol. 25, No. 2. pp. 699-738.
28. Milgrom, Paul R. and Weber, Robert J. (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, Vol. 50, No. 5, 1089-1122.
29. Myerson, Roger (1981) "Optimal Auction Design," *Mathematics of Operation Research*, 6, pp 58-73.
30. Resnick, Paul and Zeckhauser, Richard (2001). "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System," working paper, Harvard Kennedy School.
31. Roth, Al, and Ockenfels, A. (2002), "Last Minute Bidding and Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet," *American Economic Review*, 92 (4), pp 1093-1103.
32. Schindler, Julia. 2003. "Late Bidding on the Internet," mimeo, Vienna University of Economics and Business Administration.
33. Shmueli, Galit and Jank, Wolfgang (2004), "Visualizing Online Auctions," *Journal of Computational and Graphical Statistics*, to appear. Available at <http://www.smith.umd.edu/ceme/statistics/AuctionDisplays.pdf>
34. Shmueli, Galit, Russo Ralph P, and Jank, Wolfgang (2003), "Modeling Bid Arrivals in Online Auctions," University of Maryland Working paper. Available at http://www.smith.umd.edu/ceme/statistics/BidArrivals_Jan04.pdf
35. Smith, Vernon L., and Walker, James M. (1993), "Monetary Rewards and Decision Cost in Experimental Economics," *Economic Inquiry*, 31, 245-261.

36. Vickrey, William. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance*, 16:8-37.

APPENDIX A

Comparing surplus across item categories with respect to the private/common-value nature of the category.

The table below shows the estimated regression coefficients with and without nesting the category variables within a common/private value variable³⁴. The nesting variable was constructed based on Klemperer (1999), with **bold** signifying common-valued categories. Coefficients for all predictors are the same in both models (including the other predictors that are omitted from the table below), except for the “common-valued” categories. This is only a technicality because these categories are now compared to a different reference category (Pottery/Glass) within the “common-value” group.

Variable	With nesting			Without nesting					
	Coefficient	SE	Pvalue	Coefficient	SE	Pvalue			
Intercept	2.91	0.54	<.0001	2.36	0.52	<.0001			
Private value dummy	-0.55	0.16	.0005	---	---	-----			
Automotive	Same as no-nesting model			0.43	0.08	<.0001			
Books				0.13	0.08	0.13			
Clothing & accessories				0.07	0.10	0.48			
Computers				0.19	0.08	0.02			
EverythingElse				0.54	0.10	<.0001			
Electronics				0.19	0.09	0.04			
Health/Beauty				-0.14	0.12	0.28			
Home/Garden				-0.08	0.08	0.29			
Music/Movie/Games				-0.11	0.07	0.13			
Missing				0.22	0.07	0.002			
Photography				0.25	0.10	0.013			
Sporting Goods				0.10	0.09	0.28			
Toys/Hobbies				<i>Reference category</i>					
Antiques/Art				0.03	0.18	0.89	0.57	0.11	<.0001
Business & Industrial	-0.45	0.18	0.01	0.09	0.12	0.41			
Collectibles	0.01	0.16	0.94	0.56	0.07	<.0001			
Coins/Stamps	-0.49	0.17	0.004	0.06	0.10	0.55			
Jewelry	-0.05	0.17	0.74	0.49	0.09	<.0001			
Pottery/Glass	<i>Reference category</i>			0.55	0.16	0.0005			

³⁴ This table is comparable to Table 3. But there is only one category variable in the reference category.

The nesting variable is statistically significant ($p\text{-value}=0.0005$), indicating that auctions in the private-value group have lower average surplus.

If common vs private value were to explain all the differences in surplus, then all the nested category-specific coefficients should be (statistically) zero. But observe that some of the coefficients are significantly different from zero. For instance, within the private value items, Automotive, Computer, Everything Else, Missing and Photography all have positive coefficients. Private value items have (in general) lower surplus than common value item, but these five categories have positive coefficients, implying perhaps that they have a common value component in them.

A similar picture is evident for the nested coefficients within common values. Business and Industrial and Coins and Stamps are negative and significant, implying perhaps that they have a significant private value component.

In summary, the nested model measures how close to private(common) value a particular category is. It also measures in a sense whether a category is "purely" private(common) value.

APPENDIX B

Validation Analysis

Our objective here is to test whether the sample of auctions obtained from Cniper.com is any different from a randomly selected sample of eBay auctions. To obtain the latter, we undertook a “title and description” advanced search of eBay auctions, in the three currencies, using a neutral phrase “May-13-04.” The string representing this date returned the maximum number of listings in period of plus or minus 15 days, approximately 10,000³⁵. Subsequently, after the last of these auctions closed, we obtained all the independent variable information by parsing the HTML pages of those auctions that had at least one bid submitted (1077 auctions³⁶) to form the “validation data.” The only data missing was the dependent variable (surplus), as eBay does not provide access to the winning bidder’s bid. Subsequently, we compared the distribution of each of the independent variables in the validation data and the Cniper data to test for any significant difference.


These comparative studies are presented as a series of box plots and QQ plots below. For all variables we found no significant difference between the Cniper data and the validation data, supporting the assumption that the Cniper data is no different than any other randomly drawn set of eBay data.

³⁵ It could have been the begin date, end date, modified by seller date or any other occurrence of the string “May-13-04” in the title or description of the auction listing.


³⁶ A large percentage of eBay auctions get no bids at all. Secondary eBay data sites such as www.andale.com and Hammertap.com’s DeepAnalysis tool, give a range of so called “success rate” according to eBay category.

Box plots of Cniper variable vs. validation (eBay) variable.


(black arrow = mean, blue arrow = median; left plot corresponds to Cniper and the right to validation data; all variables are log transformed; ratings are also shifted by four to the right)




Price




Opening Bid




Number of Bidders




Number of Bids




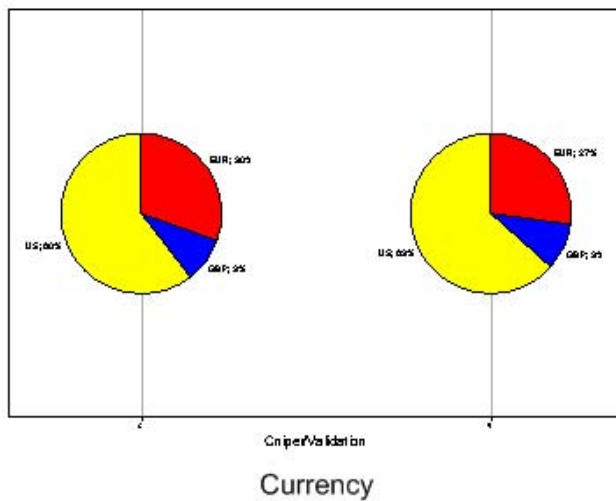
Winner rating (log +4)



Seller rating (log +4)




Categories (% auctions in each category). Blue is eBay validation data.




QQ plots of Cniper variable vs. validation (eBay) variable.


(Plots the Cniper percentile vs. its matching validation percentile (median vs. median, etc.) for each variable separately. Validation is on the horizontal axis. All variables are log transformed; ratings are also shifted by four to the right. A straight line of 45 degrees indicates that the distributions match. Note that most of the distributions closely match, and that discrepancies occur at the very top percentiles due to very right-skewed distributions.)




Number of Bids




Number of Bidders




Seller rating



Winner Rating



Price



Opening Bid