

Is Sniping A Problem For Online Auction Markets?

Matt Backus
Cornell University and eBay
Research Labs
backus@cornell.edu

Dimitriy V. Masterov
eBay Research Labs
dmasterov@ebay.com

Thomas Blake
eBay Research Labs
thblake@ebay.com

Steven Tadelis
UC Berkeley, NBER, and eBay
Research Labs
stadelis@haas.berkeley.edu

ABSTRACT

A common complaint about online auctions for consumer goods is the presence of “snipers,” who place bids in the final seconds of sequential ascending auctions with predetermined ending times. The literature conjectures that snipers are best-responding to the existence of “incremental” bidders that bid up to their valuation only as they are outbid. Snipers aim to catch these incremental bidders at a price below their reserve, with no time to respond. As a consequence, these incremental bidders may experience regret when they are outbid at the last moment at a price below their reservation value. We measure the effect of this experience on a new buyer’s propensity to participate in future auctions. We show the effect to be causal using a carefully selected subset of auctions from eBay.com and instrumental variables estimation strategy. Bidders respond to sniping quite strongly and are between 4 and 18 percent less likely to return to the platform.

General Terms

Online Auctions; Sniping

Keywords

Online Auctions; Sniping

1. INTRODUCTION

The rise of online marketplaces in general, and auction platforms such as eBay in particular, has offered researchers new sources of detailed data on the behavior of market participants. Over the past decade, there has been a growing literature focusing on auction ending rules and how these influence late-bidding behavior, often referred to as “sniping,” where some bidders submit their bid at the very last possible moment.¹ In this paper we explore the effects of

¹Sniping is a well known practice on eBay.com, and is

losing an auction to such a late bidder, in particular how such an experience affects the future participation of buyers in online auctions. We show that being sniped has a negative impact on the sniped new bidders’ likelihood of bidding in a subsequent auction.

The pioneering papers in this literature, Roth and Ockenfels [13] and Bajari and Hortaçsu [1], document the prevalence of sniping on eBay and offer some explanations that rationalize this behavior. One explanation, suggested by Roth and Ockenfels [13] and is also explored in Gray and Reiley [9], suggests that if there are some naïve bidders who do not understand the “proxy bidding” nature of online auctions, and instead choose to raise their bid every time they are outbid by others like in an English auction, then sniping by rational bidders will be a best response.² The argument rests on the fact that if a rational bidder faces an “incremental” naïve bidder, then by bidding at the very last moment he will not offer the incremental bidder enough time to raise the bid further, and hence the rational player will win the auction at a lower price than he would if he bid earlier.³

We explore a natural prediction from this framework: if a naïve bidder does not anticipate being sniped, and as a result is upset by being sniped, that bidder should be discouraged from participating in future auctions. We develop a test for this prediction and quantify the long term effects on these incremental bidders which measures the magnitude of the externality imposed by snipers using a revealed preference approach.

Using data from eBay we indeed observe that there are many incremental bidders: as many as 30 percent of bidders have bid more than once in an auction. Over 30 percent of

discussed in detail in its buying guides (<http://www.ebay.com/gds/Sniping-The-Best-kept-secret-on-Ebay/10000000002571474/g.html>, April 7, 2014).

²Proxy bidding mechanisms are akin to second-price auctions and work as follows: a bidder submits his non-proxy bid which represents his maximum willingness to pay. As subsequent bids are submitted, the auction mechanism registers the current price as one increment above the second highest proxy bid and current winner as the highest proxy bidder.

³Another popular explanation is that snipers are informed bidders who wish to hide their information from less informed bidders in a common-values setting. While the identities of the bidders are masked on eBay, the feedback score and other data which can proxy for experience, are visible. The results in Bajari and Hortaçsu [1] are suggestive of this story playing a role.

auctions have bids in the last 10 seconds of the auction, supporting the literature’s prior evidence that sniping is prevalent. Prior studies, however, have been unable to explore long-term effects.

Figure 1 suggests that there may be negative impacts to the marketplace from sniping. The horizontal axis shows the time remaining from the end of the auction at which a bidder who was winning the auction got outbid by some other bidder. The vertical axis shows the probability that the bidder who was outbid will not return in the future to participate in another auction.⁴ As the figure demonstrates, there is a rapid rise in the probability of not bidding ever again in a later auction as the time remaining from the end of the auction drops below a fraction of a minute (about 20 seconds). This means that being outbid in the last few seconds will cause the losing bidder to vote with his feet and be less likely to put in a bid on eBay again.

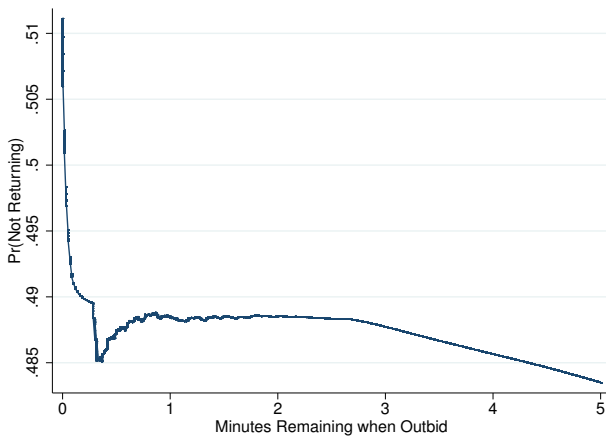


Figure 1: First-time Bidders Winning at 5 Minutes

To begin our analysis, we first define what it means to be sniped. We consider new eBay bidders who are all the high bidders with 5 minutes remaining until their auctions close. We then remove those bidders who ultimately win the auction. Hence, only new users who were winning at 5 minutes before the auction, but lost in the last five minutes, are considered. We define a bidder as sniped if he is outbid in the last 3 seconds. If he is outbid any earlier in the last 5 minutes, he will become part of the control group of losers who were not sniped. Therefore the empirical exercise compares how sniped bidders behave in the future compared to losers who were not sniped. Using a simple probit regression we first analyze the probability of leaving the site conditional on being sniped. The basic analysis shows that the likelihood that bidders who were sniped do not return to bid on eBay is about 3.5 percentage points higher than those who lose without being sniped, who fail to return 49 percent of the time.

Despite being suggestive, the simple plot in Figure 1, as well as regressions controlling for other variables such as price, category, etc., may suffer from endogeneity problems.

⁴The shape of Figure 1 is almost identical if the vertical axis shows the probability that the bidder who was outbid will not return in the future to eBay for fixed price purchases as well as auctions.

In particular, sniping is more likely to occur in markets where there are few bidders.⁵ It is these kind of markets for which a marketplace like eBay is most attractive to buyers, implying that bidders in these markets may be more likely to return to eBay. Hence, a positive correlation between sniping and auction thinness, and a positive correlation between auction thinness and the likelihood of returning to eBay, will bias downward any effect that sniping has on bidders ceasing to bid in auctions. We indeed demonstrate a positive correlation between sniping and market thinness, and develop another approach to measuring the true causal effect.

We use an instrumental variable (IV) strategy to address this concern. There are two ways that bidding occurs on eBay. First, bidders can manually insert their bid into the proxy bidding system. Second, bidders can use sniping software that does this automatically in the last seconds of the auction without their attentiveness. At night, there are fewer manual bidders active on the site, and consistent with this we observe that more auctions are won by snipers. However, 10pm in New York is only 7pm in San Francisco, while 10pm in San Francisco is 1am in New York. Therefore the 10pm San Francisco bidder is much more likely to be sniped than a 10pm New York bidder. If these bidders are otherwise comparable conditional on observables, then one can use their respective time zones as an instrument for variation in the likelihood of being sniped. This is the basis of our IV strategy.

To implement the IV approach outlined above, we use a bivariate probit regression model with local time of day and server time of day dummies as instruments to capture the exogenous variation in sniping that is due to more or less people from different time zones being awake and able to snipe. The estimates from the IV approach are much larger than the simple probit results: on average, the likelihood that bidders who were sniped do not return to bid on eBay is about 18 percentage points higher than those who lose without being sniped.

These results show that bidders respond quite strongly to being sniped.⁶ It is quite clear that rational bidders, who bid their valuation and lose, should not be upset and then vote with their feet.⁷ Hence, this behavior suggests that not only are there bidders who do not seem to understand the rules of proxy bidding, but a significant number of those are upset by the prospect of losing without being able to submit a counter-bid, and then chose not to bid in future auctions. As a result, sniping has a negative impact on the growth rate of the auction platform.

2. RELATED LITERATURE

There is an extensive literature on sniping in auction markets which seeks to answer two questions. The first is why bidders are willing to pay nonzero costs to snipe.⁸ As Zei-

⁵See Figure 4.

⁶Users often voice their distaste for snipers on eBay on forums and blogs.

⁷Losing bidders may choose not to return to auctions if they believe that prices are not a good enough deal to justify the time spent on bidding. However, there is no reason that rational bidders will choose not to return disproportionately if they lost in the very last seconds of the auction.

⁸There are a variety of third-party websites which offer sniping services and software for eBay auctions. Bidders provide their account login credentials and may pay a small fee to

thammer and Adams [14] emphasize, this fact is inconsistent with the sealed-bid abstraction which renders auction markets so empirically tractable. They show that the salience of sniping is a challenge to the theoretical underpinnings of much empirical work on auctions. The second question is the economic significance of sniping. It remains an open question whether sniping significantly affects auction revenues or the health of an auction platform, and it is on this last point we contribute to the literature.

There are a number of papers which endeavor to explain the incentive to snipe. Bajari and Hortaçsu [1] demonstrated the prevalence of sniping and suggested that if there is a common-value component to bidders' valuations, they would have an incentive to conceal their bids until the end to avoid revealing information. Roth and Ockenfels [13] and Ockenfels and Roth [11; 12] compared the soft ending times of Amazon.com auctions (now defunct) to the hard ending times at eBay.com, and proposed that sniping is a rational response to the existence of a potentially small set of "incremental bidders." Barbaro and Bracht [2] highlight the argument that bidders may use sniping as a way to avoid being squeezed by sellers who use shill bids to narrow the gap between the second-highest and the latent, highest bid.

Evidence of the returns to sniping is limited. Ely and Hossain [7] and Gray and Reiley [9] both conduct field experiments. The former find an economically insignificant payoff to sniping (about 17 cents for a new DVD). They argue that the magnitude of these gains are roughly equal to the gains of bidding early to signal one's participation, in an attempt to deter entry by other bidders. The latter find that sniped auction final prices are 2.54% lower, but the difference is not statistically significant. Glover and Raviv [8] show that auctions with soft ending times on Yahoo! sell for substantially more than those with hard ending times, and conjecture that this is attributable to less sniping. In the public finance literature on charity auctions, Elfenbein and McManus [6; 5] and Carpenter et al. [3] show that sniping has a negative effect on auction revenues. The former go on to suggest that the "warm glow" of participating in charity auctions may deter bidders from adopting more aggressive bidding strategies such as late-bidding.

3. SNIPING

We define late bidding, or sniping, as bidding that happens in the final seconds of a time-limited auction, so that other bidders have little, if any, ability to respond with an incremental bid.

Figure 2 provides a vivid illustration, where we have shown the behavior of 4 participants in a 7 day auction that was won by a sniper. The y -axis shows the dollar value of the non-proxy bids, and the x -axis shows the hours remaining when each bid was made. Bidders are color-coded and bids of each person are connected by line segments if more than one was made. There was no reserve for this auction and the starting price was \$4.99 with a Buy It Now (BIN) option of \$9.99, and an increment of 0.50. The first bid of \$7 was placed by the lavender bidder with less than 18 hours left on the clock. He was the high bidder for most of the remaining time. Around the 7 hour mark, khaki buyer made a solitary bid of 6, which was sufficient to raise the price, but not to put him in the lead. An hour later, red user became the new

have their bid placed in the last seconds of an auction.

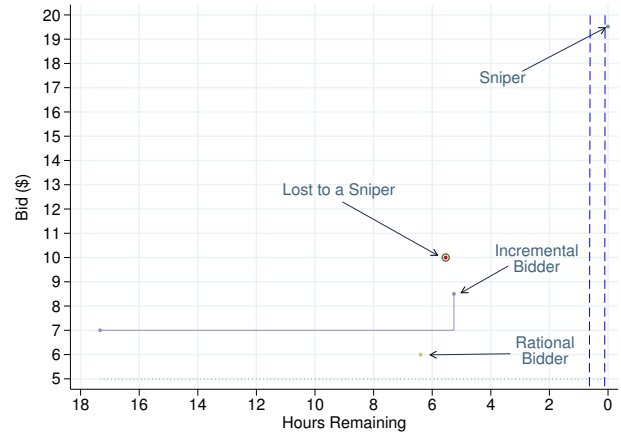


Figure 2: Bidding In An Auction

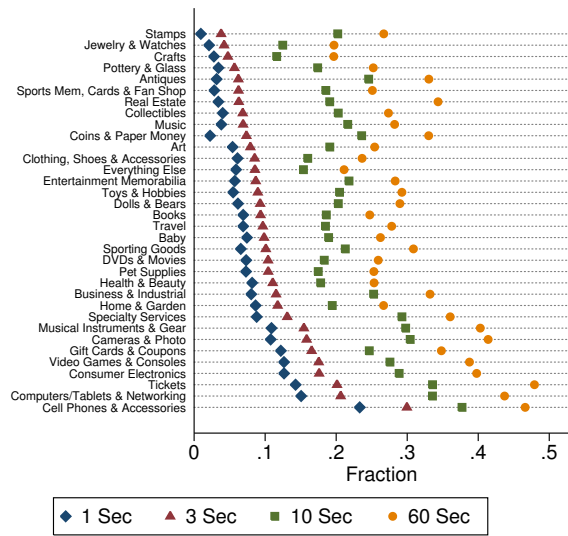


Figure 3: Sniping Rates By Meta Category

high bidder with a bid of \$10. In response, lavender bidder raised his bid to \$8.50, though that was not enough to put him back on top. Red looked to be winning for the next few hours, even with 5 minutes remaining (first dashed line).⁹ With 3 seconds to go, a sniper placed a bid of \$19.52. He won and paid 10.50 for this item. We do not know whether the red bidder would have countered, but we do know that he had no time to respond.

Incremental bidding happens fairly frequently on eBay. We find that bidders place multiple bids 29 percent of the time. Moreover, 4.4 percent of bidders place five or more bids in a single auction. The prevalence of this behavior suggests that snipers may deny a substantial fraction of bidders the opportunity to counter bid.

Figure 3 shows that there is considerable variation in how prevalent sniping is in various product categories on eBay.¹⁰

⁹Dashed lines are not drawn to scale.

¹⁰Figure 3 shows data from the universe of auctions from Oc-

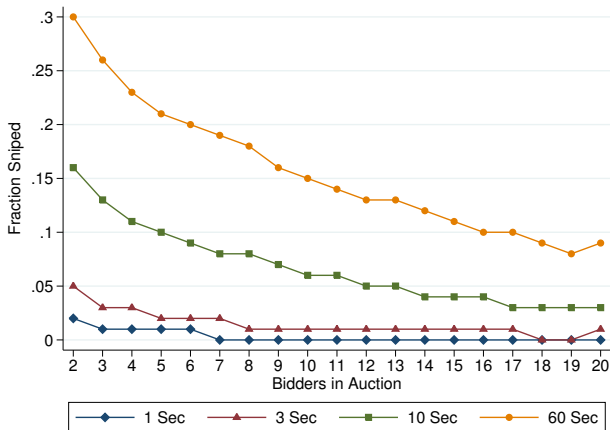


Figure 4: Sniping Rates By Number of Bidders

It plots the fraction of auctions in each meta category that were won through sniping. Mechanically, that fraction grows as we change the definition of sniping from bidding in the last one second to bidding in the last minute. For example, over a third of cell phone auctions were won in the last second, but over half were sniped in the last 60 seconds.

Using our preferred 3 second definition, sniped tickets are over twice as common as with crafts. We note that a rank order of the categories varies somewhat by definition, but not substantially. This confirms that sniping occurs at different rates for different products. In our sample, sniping is more prevalent in thinner markets, a fact that may confound analysis of its effects on follow-up behavior. Figure 4 plots the fraction of sniped auctions by the number of bidders in those auctions from our sample, which is described in Section 5.

We also see substantial variation in the sniping rate at different times of the day. Sniping, as a percentage of all bids, peaks in the middle of the night when many bidders are likely sleeping but automated sniping bots are still active.¹¹ By contrast, manual bids are more likely to happen in the evening hours when the site is busiest. We will use this variation in our IV estimation.

4. EMPIRICAL DESIGN

The ideal experiment to test our conjecture that sniping drives out buyers from the auction platform would have the following two features. First, we would want to focus on a cohort of *new* eBay users.¹² This group is still forming their expectations about how much surplus they can expect from the site, and so any effects should be readily apparent. Moreover, all users start out as new, and once buyers amass experience, it becomes considerably harder to adjust for that heterogeneity and to deal with the selection issues caused by attrition. For instance, if sniping-averse buyers

tober 2009 to October 2011. Our chosen sample, discussed in Section 5, is a subset of these auctions.

¹¹Figure 5 shows this pattern for each time zone.

¹²We consider an eBay user to be new if the observed auction is the first action, bid or purchase, made on the site. We allow users to have signed up as early as 6 months prior to our window.

leave eBay, the population of surviving seasoned users may consider sniping a feature since it allows them to purchase items at a lower price or because it has some ludic value for them, even if they are occasionally sniped themselves or suffer as sellers. On the other hand, this focus on new users somewhat circumscribes the external validity of the effects we estimate.

Once we have a population of auctions in which a new buyer becomes the current high bidder, we can flip a coin to select a random subset to be treated by being sniped in the last moment. New users in the *control* group would be outbid by us earlier on. Our winning bid for them would be set high enough that a control buyer who revises his bid upward would never regain his initial top position. For both groups, at the point of randomization, the auction would become “hidden” to restrict further entry by new bidders and “locked” to keep anyone other than the new buyer from revising his bid. This minimizes substitution bias from control units being treated, but also ensures that we have full control over the information that is revealed to the losers,¹³ and over the time that they perceive to be in the lead. This is the second essential feature.

Once this hypothetical experiment starts, we can follow both groups for some period of time to track their activity. We can compare the average number of auction purchases, total revenue, or merely the proportion who bid in an auction again. The difference between the two groups in any of the those metrics would be the treatment effect of being sniped. Note that we are not comparing winners to losers. The treatment alters the manner in which a buyer lost, not whether he did.

The second aspect of this design renders the experiment infeasible since it is impossible to control the dynamics of auctions in the field. Instead we propose the following exercise. We take a large set of auctions and restrict attention to bidders who are winning with 5 minutes remaining until the auction closes. For each auction, there is a single bidder who meets this criterion. We proceed to exclude all but first-time bidders. We also exclude bidders who go on to win the auction, out of a concern that sniped bidders, who necessarily lose, have residual demand for the product being auctioned. We identify a bidder as sniped if he is outbid in the last 3 seconds. If he is outbid any earlier in the last 5 minutes, he is part of the control group. This allows us to compare outcomes for those losers who could have potentially responded by bringing their bid closer to their true reservation value, to those for whom this was impossible. Imagine that we could find another auction that had the exact same attributes and history of bids, but the final bid came in between the 2 vertical blue lines, the 5 minute and the 3 second marks in Figure 2. This would serve as a counterfactual auction for the sniped case, since the red bidder would have had enough time to respond to a bid that came in somewhat earlier.

We note that our arbitrary 3 second threshold may misclassify some bidders and cause attenuation bias in the estimated treatment effect. The bidders’ ability to respond will depend on the quality of their internet connection and the mode of their connection, whether PC, mobile, or other.

¹³For example, if the observed standing price rises for the control group as more bids are made, the control group will learn more about by how much they lost than the treatment losers.

We have chosen 3 seconds because this both fits the visual presented in Figure 1 and because this is a logical threshold at which human response time limits the ability to counter bid. We choose a five minute counterfactual window because we want to limit the sample to auctions where the bidder was winning as the auction came to a close. We examine in the appendix how sensitive our results are to the 3 second definition of sniping and to the 5 minute window.

5. SAMPLE CONSTRUCTION

We want to measure the impact of sniping on consumer behavior, so we want to construct a sample of consumers that are similar in most respects but for the fact that some have been outbid in the last few seconds of the auction where there was no chance to respond. To do this we need to find i) a sample of users that are alike, in expectation, at the time of the auction and ii) are expected to have similar behavior after the auction but for the sniping. Therefore, we focus on i) first-time bidders and ii) bidders who lose their first auction. We identify *all* first-time bids between October 2009 and October of 2011 where the bidder did not win the auction. We chose this period so that we have a full year of data for all bidders after their first auction. This is over 2 million bidders. We then identify those bidders that were the winning bidder with exactly five minutes remaining in the auction. Some of these bidders were winning for only seconds and some were winning for days, but all were the winner at the 5 minute mark and eventually lost. Note that since each auction can have only one winner at a single moment in time, this eliminates any duplicate auctions so that each bidder-auction pair is unique. In our main specifications, we further limit the sample to bidders that were *still* winning with one minute remaining. Thus these bidders were winning for at least four minutes because they were winning at the five minute mark and the one minute mark.

Table 1: Sample Size

All first time losing bidders	2226307
Winning with 5 min remaining	705957
Outbid with less than 1 min remaining	458168
Covered by catalog	89168
Thick market (10 auctions of product)	32139

Finally, in some of our specifications we limit the analysis to the subsample of auctions that are in thicker markets. To do this, we first identify those auctions that are identified by eBay’s product catalog, which are typically items with well defined manufacturer product codes. This further limits the sample because only a portion of eBay’s inventory is covered by the catalog. We then find catalog products with more than 10 transactions over the two-year period. Table 1 shows the counts of each sample truncation.

In Table 2, we present summary statistics for our full sample of auctions with first-time bidders winning at 5 minutes. Although all the bidders in our sample were outbid, only 7 percent of these auctions were outbid in the last 3 seconds, 10 percent in the last 4 seconds, and so forth. In general, 49 percent of first-time bidders do not bid again in an eBay auction. We stress that this is specific to auction: many eBay users switch to fixed price listings. We still consider these users as not returning because we are interested in

Table 2: Sample Statistics

	Mean	SD	Min	Max	N
Sniped in 3 Seconds	0.0711	0.257	0	1	705957
Sniped in 4 Seconds	0.101	0.302	0	1	705957
Sniped in 6 Seconds	0.185	0.388	0	1	705957
Sniped in 8 Seconds	0.264	0.441	0	1	705957
Sniped in 10 Seconds	0.308	0.462	0	1	705957
Sniped (Power Seller Auction)	0.0260	0.159	0	1	705957
Bidder Does Not Return	0.488	0.500	0	1	705957
Minutes Remaining When Outbid	1.161	1.397	0	5	705957
Hours Left When Bid	10.01	23.37	0.1	399.5	705957
Own Bid as a Percentage of Highest Bid	0.749	0.219	0.0000	1.0	672049
Power seller auction	0.418	0.493	0	1	705957
Bid / Expected Highest Bid	0.739	0.343	0.0000	37	144647
Bid / Expected Category Highest	0.667	1.404	0.00000	375.9	705947
Bid / Best Expected Highest	0.684	1.357	0.00000	375.9	705947

capturing an effect associated with a dis-utility of auctions, not the overall platform. Statistics for additional controls are shown as well.

6. ESTIMATION

The principal results are shown graphically in Figure 1, which shows a marked increase in the probability of not returning in the last few seconds of the auction. It plots a locally weighted smooth of not returning against minutes remaining when outbid conditional on winning with five minutes remaining in the auction.

To make the suggested results of Figure 1 concrete, we estimate probit regressions assessing the impact of being ‘sniped’, or outbid by late bidding, on future behavior. Figure 1 demonstrates a distinct jump in the probability of leaving the platform if outbid in the last few seconds of the auction. We then regress an indicator for whether the user does not participate in an eBay auction within 12 months of being outbid on an indicator for whether or not the time of outbid was within three seconds of the close of the auction.¹⁴ We limit the regression to all auctions where the new bidder was still winning with one minute remaining.¹⁵

Table 3 presents the probit estimates of marginal effects of sniped on failure to return from several specifications. Column 1 presents the estimates without any controls and indicates a 3.3 percentage point increase in the probability of not returning. This confirms that the increase in not returning as outbid time approaches zero shown in Figure 1, which rises from 48 percent to 51 percent, is a statistically significant relationship. This estimate is unconditional, and since sniping is not a random occurrence, there are a variety of confounders we first try to address with controls.

In Column 2, we add controls for two important selection mechanisms. First, bidders who are outbid earlier in the auction are likely to have lower bids simply because lower bids are more likely to be outbid sooner as bidders arrive. Bidders that bid lower have different expectations of surplus from buying on the platform generally and may be more likely to leave the site once final prices are revealed. Lower bids could therefore be correlated with the timing of being outbid and not returning. We include the ratio of the bidder’s (losing) bid to the highest (winning) bid as a control.

¹⁴A view of the data in Figure 1 suggests a kink in the return-time remaining relationship at three seconds.

¹⁵This limits the sample to all auctions where the new bidder was winning from five minutes to one minute before the auction. This improves the comparability among bidders: we are now only comparing losers in the last few seconds to the last minute.

Table 3: Probit Estimates, Dep Var = Pr(Not Returning)

	(1)	(2)	(3)	(4)	(5)	(6)
Sniped in 3 Seconds	0.0331*** (0.00236)	0.0355*** (0.00248)	0.0138*** (0.00253)	0.0483*** (0.00312)	0.01000 (0.00705)	0.0380*** (0.00721)
Sniped (Power Seller Auction)				-0.0897*** (0.00468)		
Own Bid as a Percentage of Highest Bid		-0.159*** (0.0161)	-0.0123 (0.0164)	-0.00439 (0.0164)		
(Bid / Highest) ²		0.191*** (0.0128)	0.0470*** (0.0131)	0.0434*** (0.0131)		
Hours Left When Bid		0.00421*** (0.0000789)	0.00451*** (0.0000787)	0.00450*** (0.0000787)	0.00664*** (0.000459)	
(Hours Left When Bid) ²		-0.0000160*** (0.000000651)	-0.0000177*** (0.000000639)	-0.0000176*** (0.000000639)	-0.0000333*** (0.00000370)	
Bid / Expected Highest Bid					-0.0712** (0.0226)	-0.111*** (0.0236)
(Bid / Expected Highest) ²					0.0313* (0.0135)	0.0511*** (0.0139)
Category Dummies			Yes	Yes	Yes	
Window	1 min	1 min	1 min	1 min	1 min	1 min
Thick Markets					Yes	Yes
Product RE						Yes
N	458168	440670	440670	440670	32116	32116

Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We also include controls for the timing of the sniped bidders' bid. Since we condition on the bidder winning at 5 minutes remaining, auctions where bidders had been winning for longer are auctions with less activity and may be more likely to be sniped because potential bidders are lying in wait. Bidders that bid very early in the auction may have a different probability of returning. It is not a priori clear what direction these potential effects would bias the results, but the coefficient on sniping remains consistent at 3.5 percentage points.

From Figure 3 we see that there is some variation in sniping rates across categories so in Column 3 we include fixed effects for eBay's 35 top line product meta categories. As was shown in Figure 3, sniping rates vary across product types which naturally attract bidders with different propensities to return. The estimated effect decreases to 1.38 percentage points which hints that cross category effects may bias the results, but meta category is a very broad classification that masks important product level variation. So we explore several methods of controlling for more specific product attributes.

In Column 4, we add an interaction of the sniped indicator with an indicator for whether the item was listed by a power seller, eBay's designation for high volume, high quality sellers. These items are likely to attract bidders looking for very common items on eBay. It is perhaps not surprising that being sniped is much less burdensome for these buyers because there is a greater likelihood that they will try again.

In Columns 5 and 6, we limit the sample to auctions for items that are in 'thick' markets: the item is identified in eBay's product catalog and has more than 10 auctions in our full 2 year sample. These sets of auctions are likely to be very

similar to one another and should reduce all selection effects related to product characteristics. We replace the own bid-to-highest-bid ratio with the ratio of the bid to the expected price for that product (average final price of the product), which is a better proxy for bid's seriousness. Column 5 shows that the estimate is not significantly different from the estimates in Column 3. The estimate loses significance only because we have a lower sample size. In Column 6, we estimate a probit model with product level random effects and notice that the coefficient rises to 3.8 percentage points.

We are concerned about the endogeneity of sniping selection and return rates, but the sign of the bias is not obvious. The returns from sniping would be higher in thinner markets, where incremental bidders are less likely to be bid up by non-sniping but rational bidders. And these thinner markets are more prevalent on eBay so bidders in these markets are more likely to return. Thus, the positive shock to sniping would be a negative shock to exit rates and would bias our estimates downward. Column 6 supports this view because removing much of that bias by focusing on within product variation in sniping increases the coefficient.

In order to more fully address the selection endogeneity, we leverage an instrumental variables strategy that relies on the fact that auctions on eBay are national events but that sniping is a function of the number of bidders on the site in the last few minutes of the auction. We want an instrument that shifts the probability of being outbid in the last few seconds, as opposed to being outbid in the last few minutes. Auctions that have more rational, non-sniping bidders watching the close are less likely to be sniped because the first-time bidder will be outbid in a way that allows them to respond. We also need this instrument to be correlated

with the probability of a bidder returning to the platform only through the probability of being sniped.

Auctions end on eBay simultaneously across the country and therefore at different *local* times of day depending on the time zone of the bidder. We found that the probability of being sniped varies dramatically by the closing time of the auction because the number of users browsing the site varies dramatically by time of day. When there are fewer users on the site, there are fewer non-incremental, non-sniping bidders to drive up the auction price earlier in the auction. This creates a cyclical relationship between the proportion of auctions that are sniped and the time of day. We show

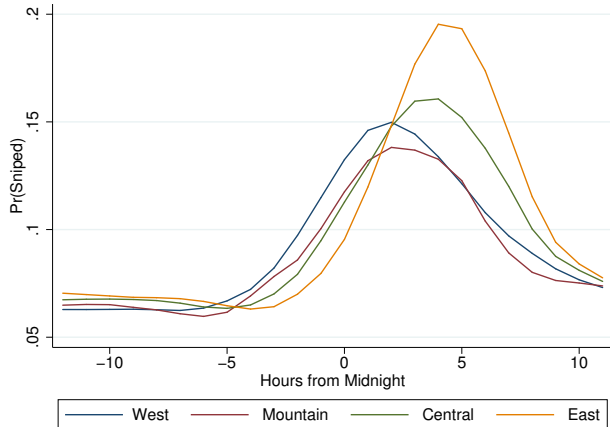


Figure 5: Probability of Sniping by Time Zone

in Figure 5 that this relationship actually varies by time zone. We plot a dummy for being sniped against the local auction end time of the first-time bidder for the auctions in our sample and then fit a local polynomial. As is evident, a bidder on the West Coast watching an auction at midnight has almost a 15 percent chance of being sniped but the same bidder on the east coast has only a roughly 8 percent chance of being sniped. This is because midnight on the east coast is still prime time on the west coast and there are more non-sniping rational bidders in the auction. Thus, two bidders bidding at midnight have a different chance of being sniped and their only difference is time zone. We assume that the probability of returning is otherwise uncorrelated with time zone, an assumption supported by eBay’s national presence.

To implement this instrument, we utilize a bivariate probit regression with local time of day and server time of day dummies.¹⁶ We use the server time dummies as the excluded instruments which are uncorrelated with churn, conditional on the local time of day.

Table 4 presents the marginal effects and average treatment effects estimates.¹⁷ Columns 2 and 3 add category dummies and the hours remaining controls used before. Column 4 restricts the sample to power seller listings and first-time bidders still winning at one minute.

¹⁶Bivariate probit can serve as an instrumental variables estimator where both the dependent and independent variables are binary Greene [10].

¹⁷We use the approach of Chiburis et al. [4] to calculate the average treatment effect (ATE) and average treatment effect on the treated (ATT).

Table 4: Bivariate Probit Estimates

	(1)	(2)	(3)	(4)
β	1.105*** (0.13670)	0.726*** (0.08794)	0.397*** (0.06510)	0.493*** (0.14566)
Marginal Effect	0.426*** (0.04999)	0.280*** (0.03301)	0.153*** (0.02483)	0.186*** (0.05413)
Average Effects				
ATE	0.384*** (0.03353)	0.269*** (0.02334)	0.152*** (0.01929)	0.189** (0.06479)
ATT	0.403*** (0.03790)	0.275*** (0.02516)	0.152*** (0.01950)	0.179** (0.05878)
Category Dummies		Yes	Yes	Yes
Hours Remaining Controls			Yes	Yes
Perc of Highest Bid Controls			Yes	Yes
Power Seller				Yes
Window				1 min
N	320909	320909	305474	82819

The estimates are much larger than the probit results but are in line with the random effects estimate. In our most restrictive sample shown in Column 4, the ATE and ATT are about 18 percentage points. This demonstrates, at the very least, that the 3 percent estimates are biased downward and that the true effect is substantial. Taken together, the evidence bolsters the claim that being sniped has large effects on users’ propensity to participate in future auctions.

7. CONCLUDING REMARKS

There is strong evidence that being sniped is an experience which discourages new bidders from returning to bid again. We find a sizeable effect in our probit specifications, which we believe to be biased downward, and a much larger effect when we use bidder time zones to instrument for the likelihood of being sniped. These findings are consistent with the story advanced in Roth and Ockenfels [13] that sniping is a rational response to the presence of incremental bidders.

The effect is economically significant from the perspective of an auction platform. Even small changes in the rate at which new participants join the platform can have a tremendous effect on long-term platform size.

The recent surge in availability of e-commerce data is a unique opportunity to re-examine hypotheses that had been formerly pursued using only scraped, field experiment, or lab experiment data. The advantage of using data directly from the e-commerce platform is that it is so extensive that we are able to narrow our empirical strategy to a subset of comparable auctions in order to identify the effects of interest, while still enjoying a wealth of data. We believe that this development offers researchers the opportunity to rigorously test theory where before we could have been little more than suggestive, and that this is an important direction for new work.

References

- [1] Patrick Bajari and Ali Hortaçsu. The winner’s curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34: 329–355, 2003.
- [2] Salvatore Barbaro and Bernd Bracht. Shilling, squeezing, sniping: Explaining late bidding in online second-price auctions. *Working article. University of Mainz, Mainz, Germany*, 2004.
- [3] Jeffrey Carpenter, Jessica Holmes, and Peter Hans Matthews. Jumping and sniping at the silents: Does it matter for charities? *Journal of Public Economics*, 95(5):395–402, 2011.
- [4] Richard Chiburis, Jishnu Das, and Michael Lokshin. A practical comparison of the bivariate probit and linear iv estimators. *World Bank Policy Research Working Paper 5601*, 2011.
- [5] Daniel W Elfenbein and Brian McManus. A greater price for a greater good? Evidence that consumers pay more for charity-linked products. *American Economic Journal: Economic Policy*, 2(2):28–60, 2010.
- [6] Daniel W Elfenbein and Brian McManus. Last-minute bidding in eBay charity auctions. *Economics Letters*, 107(1):42–45, 2010.
- [7] Jeffrey C Ely and Tanjim Hossain. Sniping and squatting in auction markets. *American Economic Journal: Microeconomics*, 1(2):68–94, 2009.
- [8] Brent Glover and Yaron Raviv. Revenue non-equivalence between auctions with soft and hard closing mechanisms: New evidence from Yahoo! *Journal of Economic Behavior & Organization*, 81(1):129–136, 2012.
- [9] Sean Gray and David Reiley. Measuring the benefits to sniping on eBay: Evidence from a field experiment. *New York University working paper*, 2004.
- [10] W.H. Greene. *Econometric Analysis*. Prentice Hall, 2003. ISBN 9780130661890. URL <http://books.google.com/books?id=JjkWAQAAMAAJ>.
- [11] Axel Ockenfels and Alvin E Roth. The timing of bids in internet auctions: Market design, bidder behavior, and artificial agents. *AI magazine*, 23(3):79, 2002.
- [12] Axel Ockenfels and Alvin E Roth. Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic Behavior*, 55(2):297–320, 2006.
- [13] Alvin E Roth and Axel Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and amazon auctions on the internet. *The American Economic Review*, 92(4):1093–1103, 2002.
- [14] Robert Zeithammer and Christopher Adams. The sealed-bid abstraction in online auctions. *Marketing Science*, 29(6):964–987, 2010.

APPENDIX

A. ROBUSTNESS

A.1 Window

The probit results in Table 3 use a sample of auctions where the first-time bidder was winning at 5 minutes and still winning with 1 minute remaining. In Table 5 we present estimates from samples with different control windows. The estimates from Column 3 are those presented in Column 3 of Table 3. In Columns 1 and 2, we narrow the sample to auctions where the first-time bidder was outbid in 6 and 10 seconds. Column 1 shows that there is still a 1.1 percentage point increase in the probability of not returning for bidders who are outbid with less than 3 seconds left compared to those outbid between 3 and 6 seconds. This is slightly smaller than the other windows, which confirms the relationship shown in Figure 1. Conversely, Columns 4 and 5 show that there is actually very little effect of adding auctions where the outbidding occurs early in the auction.

A.2 Snipe Definition

Table 6 presents probit results for increasingly expansive definitions of sniping, from 6 seconds to almost 5 minutes (300 seconds). When we vary the sniping definition even a little, the estimates are negative suggesting that the up-tick in exit is only present for those auctions where the first-time bidder actually did not have time to respond. The fact that these probit results are negative is at first alarming, but actually consistent with the instrumental variable motivation.

A.3 Linear Probability Model

We use a probit model in our primary specifications but we can easily estimate a linear probability model for robustness. Table 7 presents the results from Table 3 with a linear model. The estimates are qualitatively similar. To check the robustness of the downward bias on exit from product selection, we leverage the time zone instrument with two stage least squares estimation of the linear model. Here the estimates are substantially larger, so large that we believe the linear model for the binary outcome is not well behaved. We interpret these results as consistent with the probit results because the negative bias persists and simple OLS estimates are similar to the base probit models. The bivariate probit is a more appropriate model given that both the dependent and independent variables are binary.

Table 8: Linear Probability Instrumental Variables

	(1)	(2)	(3)	(4)
Sniped	0.525*** (0.0852)	0.411*** (0.102)	0.489*** (0.101)	0.141 (0.145)
Hours Remaining			0.00372*** (0.000146)	0.00404*** (0.000175)
(Hours Remaining) ²			-0.0000127*** (0.00000113)	-0.0000156*** (0.00000145)
Constant	0.406*** (0.0147)	0.359*** (0.0133)	0.309*** (0.0125)	0.308*** (0.0183)
Category Dummies		Yes	Yes	Yes
Non-Power Sellers				Yes
Window		1 min	1 min	1 min
N	415285	415285	415285	172660

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Varying the Control Windows

	(1)	(2)	(3)	(4)	(5)
	6 Second Window	10 Second Window	60 Second Window	150 Second Window	300 Second Window
Sniped in 3 Seconds (d)	0.0106*** (0.00312)	0.0120*** (0.00278)	0.0135*** (0.00254)	0.0135*** (0.00249)	0.0131*** (0.00246)
Own Bid as a Percentage of Highest Bid	0.00790 (0.0262)	-0.00711 (0.0218)	-0.0172 (0.0164)	-0.0366* (0.0149)	-0.0652*** (0.0136)
(Bid / Highest) ²	0.0204 (0.0217)	0.0335 (0.0178)	0.0510*** (0.0131)	0.0722*** (0.0118)	0.0991*** (0.0108)
Hours Left When Bid	0.00461*** (0.000140)	0.00447*** (0.000112)	0.00451*** (0.0000787)	0.00449*** (0.0000716)	0.00451*** (0.0000664)
(Hours Left When Bid) ²	-0.0000183*** (0.00000113)	-0.0000173*** (0.000000919)	-0.0000177*** (0.000000640)	-0.0000177*** (0.000000582)	-0.0000177*** (0.000000542)
Category Dummies	Yes	Yes	Yes	Yes	Yes
N	125967	209920	440674	549116	672049

Robust standard errors in parentheses

(d) for discrete change of dummy variable from 0 to 1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ **Table 6: Varying the Sniping Definition**

	(1)	(2)	(3)	(4)	(5)
	6 Seconds Left	30 Seconds Left	60 Seconds Left	150 Seconds Left	250 Seconds Left
Sniped (d)	0.00144 (0.00160)	-0.00170 (0.00125)	-0.00257 (0.00131)	-0.00278 (0.00160)	-0.00226 (0.00263)
Own Bid as a Percentage of Highest Bid	-0.0732*** (0.0135)	-0.0750*** (0.0135)	-0.0753*** (0.0135)	-0.0751*** (0.0135)	-0.0745*** (0.0135)
(Bid / Highest) ²	0.104*** (0.0107)	0.105*** (0.0107)	0.105*** (0.0107)	0.105*** (0.0107)	0.105*** (0.0107)
Hours Left When Bid	0.00452*** (0.0000665)	0.00453*** (0.0000666)	0.00453*** (0.0000666)	0.00453*** (0.0000665)	0.00453*** (0.0000665)
(Hours Left When Bid) ²	-0.0000178*** (0.000000542)	-0.0000178*** (0.000000543)	-0.0000179*** (0.000000543)	-0.0000179*** (0.000000542)	-0.0000178*** (0.000000542)
Category Dummies	Yes	Yes	Yes	Yes	Yes
Window	5 min	5 min	5 min	5 min	5 min
N	672049	672049	672049	672049	672049

Robust standard errors in parentheses

(d) for discrete change of dummy variable from 0 to 1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ **Table 7: Linear Probability Model**

	(1)	(2)	(3)	(4)
Sniped in 3 Seconds	0.0331*** (0.00236)	0.0350*** (0.00246)	0.0132*** (0.00245)	0.0461*** (0.00302)
Sniped (Power Seller Auction)				-0.0863*** (0.00463)
Own Bid as a Percentage of Highest Bid		-0.162*** (0.0159)	-0.0170 (0.0158)	-0.00952 (0.0158)
(Bid / Highest) ²		0.193*** (0.0126)	0.0497*** (0.0126)	0.0463*** (0.0126)
Hours Left When Bid		0.00416*** (0.0000745)	0.00439*** (0.0000739)	0.00437*** (0.0000739)
(Hours Left When Bid) ²		-0.0000162*** (0.000000603)	-0.0000175*** (0.000000593)	-0.0000174*** (0.000000592)
Category Dummies		Yes	Yes	Yes
Window	1 min	1 min	1 min	1 min
N	458168	440674	440674	440674

Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$