Squatting, Sniping, and Online Strategy: Analyzing Early and Late Bidding in eBay Auctions

Philip Groenwegen
Advised by Professor Philip A. Haile

Yale University
Department of Economics

April 3, 2017

Abstract: This essay builds on previous studies of bidding strategies in online auctions. Many studies have observed an increase in bidding activity close to the ending time of eBay auctions, despite the fact that such behavior is not predicted by economic theory and is explicitly discouraged by eBay. There is also some evidence that bidding very close to the beginning of an eBay auction is another widely used strategy, although this too is unexpected. I use bid data from eBay to examine the conditions under which early and late bidding is most prevalent. I then study the effect of these strategies on the likelihood of a bidder winning an auction and the closing price of an auction.

I am sincerely grateful to Professor Philip A. Haile for his thoughtful advising. I also thank Joshua Dull at the Center for Science and Social Science Information and Sachith Gullapalli for their technical advice.

1. Introduction and Literature Review

The explosive growth of the Internet over the past two decades has afforded researchers new opportunities to test theories of economic behavior on a large scale. Online bidding sites like eBay have made research on auctions particularly fruitful. In 2015, the British newspaper *The Daily Telegraph* estimated that eBay handled nearly $83 billion in sales per year and had about 800 million items for sale at any given time. The scale of eBay and the diversity of items sold on the platform make it a useful tool for testing economic theory. Studying eBay and other online auctions has also revealed many cases of unexpected bidding behavior. One of the most important of these observations is the fact that bidding tends to occur in clusters, with some bids at the beginning of the auction, relatively few in the middle, and a large spike in activity in the last few minutes.

The exact rules of an auction on eBay can vary, but one of the most prevalent methods of selling is a second-price auction.[1] When a bidder enters an auction, eBay prompts her to bid at least a minimum increment—set by the seller—over the standing price. There is no maximum on the amount a bidder can enter. If the bidder expresses a willingness to pay above the minimum price, eBay automatically raises the bid of the bidder who submitted a higher willingness to pay, such that the bidder with the highest valuation wins by paying marginally more than the second-highest bid. Thus if the minimum bid is $10, the bid increment is $0.50, and the first bidder enters a maximum willingness to pay of $20, eBay will show that the first bidder has bid $10. If a second bidder enters a willingness to pay of $15, the standing bid will be raised to bidder two's $15 plus the $0.50 increment to $15.50, the minimum amount required for bidder one, the bidder with the highest willingness to pay, to win. Therefore, at any given point, participants know the

[1] Many auctions on eBay have a "But It Now" option that allows a bidder to end the auction by paying a high, posted price. These auctions are excluded from my analysis.

value of all bids except the highest one. Bidding continues in this manner until the auction ends at a fixed time. Bidders may increase their maximum willingness to pay at any time, although eBay discourages this behavior, claiming that the automatic bidding function allows a bidder to win at the minimum price necessary.

The bidding rules on eBay mean that eBay auctions are similar to second-price sealed-bid auctions. As Vickery (1961) lays out, in auctions like these, players have a weakly dominant strategy to bid their valuations. In other words, bidders should bid their maximum willingness to pay for an item once and ignore whatever else happens in the auction. The time a bid is placed should have no effect on a bidder's strategy. Despite these strong predictions, eBay auctions display many characteristics that are not consistent with a sealed-bid second-price model. Many "incremental" bidders submit multiple bids, slowly ratcheting up their willingness to pay. This indicates that they do not initially bid their valuations, contrary to Vickery's prediction.

The feature of eBay bidding that is most strikingly different from the theoretical prediction is that many bids are submitted close to the end of an auction—a behavior known as sniping. As eBay advises its users and as the traditional theory would predict, there is not a clear benefit to sniping if bidders bid their valuations. Because of eBay's proxy bidding function, a bidder could bid a high valuation but pay a relatively low price—the price necessary to beat the second-highest bidder.[2] Moreover, by bidding extremely close to the end of an auction, bidders risk that their bids may not be transmitted before the auction ends, thus losing the opportunity to compete altogether. Yet despite all these theoretically costly consequences, sniping is a relatively common practice. In fact, Ockenfels and Roth (2006) find that sniping is especially popular

[2] As many theorists have noted, however, this logic is complicated by the fact that bidders' decisions to enter are informed by the actions of other bidders. This is discussed in greater detail in Section 2.

among more experienced bidders on eBay, suggesting that bidders learn this behavior. Moreover, although eBay discourages sniping, many online forums encourage bidders to engage in the practice, and there even exist services that will automatically submit sniping bids. None of this is to say that bidders ignore the implications of the rules of eBay auctions. In fact, Ariely, Ockenfels, and Roth (2003) compare the prevalence of late bidding on eBay, which has a fixed closing time, to that on Amazon, which extends the closing time by a few minutes after a bid is placed. As expected, there is less of a flurry of activity close to the expected closing time of an Amazon auction.

Previous studies have produced evidence of a modest return to sniping. Ely and Hossain (2009) find that sniping leads to a roughly 5% increase in the probability of winning an auction, as well as a small yet statistically significant decrease in the price paid for the item. Gray and Reiley (2007) perform a similar study and find an even smaller and statistically insignificant benefit to sniping. The statistical insignificance is likely due to the small sample sized used in their study. Gray and Reiley speculate that the observed decrease in the magnitude of the benefit of sniping could be due to the fact that markets have become more competitive or bidders more sophisticated over time. However, they find no evidence of different returns to sniping in markets of different sizes and find that sniping is just as frequent in their study as in previous ones. Both of these studies use field experiments to evaluate the returns to sniping. The authors place both early and late bids on items and compare the prices they pay under the two strategies. Ely and Hossain compare the price they pay to a hypothetical valuation corresponding to their bid in order to find the surplus of winning a given auction. By looking at data from previously completed auctions instead of actually bidding in auctions, I do not estimate the surplus of the winner, but rather the difference in closing price associated with sniping. On average, this should reflect the change in surplus a winner would receive if she were to snipe.

One popular explanation for the prevalence of sniping is the presence of “naïve” bidders. eBay has thousands of active bidders, not all of whom may be well-versed in theoretically optimal bidding strategies. Ku, Malhotra, and Murnighan (2005) demonstrate that competition can cause bidders to become more aggressive, in what they call “competitive arousal.” Similarly, Ariely and Simonson (2003) argue that bidders may enter an auction when the price is relatively low and then become attached not necessarily to the item but to the prospect of winning, leading them to increase their bids many times, perhaps at the last minute, after being outbid. Indeed, Ariely and Simonson find that bidders often buy common retail items like DVDs for more than a brick and mortar store would charge. This may suggest that bidders value winning per se, even if they must overpay to win. Moreover, as Bajari and Hortaçsu (2004) suggest, bidders may participate in auctions to have fun rather than to buy important items. Another potential explanation for sniping is that bidders likely derive more pleasure from bidding in real time than from learning that their bid has been raised automatically.

No matter the reason why these “naïve” bidders behave as they do, the presence of these bidders could influence the strategies of better-informed bidders. Ockenfels and Roth (2006) show that late bidding by experienced bidders is a best response to the existence of at least one naïve bidder, who does not initially bid his valuation but instead raises it incrementally. By sniping, sophisticated bidders avoid starting a bidding war with incremental bidders. Indeed, the authors find that bidders who bid only once tend to submit their bids later than the last bids placed by incremental bidders, which is further evidence to suggest that sniping is favored by experienced bidders as a response to their naïve competitors.

Similarly, Ely and Hossain (2009) find evidence that naïve or incremental bidding is prevalent, but they argue that this behavior does not have a significant effect due to the large size of eBay's markets. Instead, they propose a model with two counteracting effects of placing a bid. On the one hand, placing a bid early in an auction risks provoking incremental bidders and starting a price war, raising the closing price in what the authors call the "escalation effect."[3] On the other hand, placing an early bid could deter potential bidders and encourage them to look for the same good in another auction in what the authors call the "competition effect."

Ely and Hossain consider the magnitude of these effects for snipers and "squatters," bidders who place a high bid early in an auction. The authors find that both strategies allow bidders to pay a statistically significant amount less than their maximum bid. Though they predict that sniping and squatting should lead to the same closing price, in fact they find that sniping leads to lower closing prices than squatting does, in what they claim demonstrates the dominance of the escalation effect over the competition effect. This model is discussed in greater detail in Section 2.

Ely and Hossain also explain the incentive to snipe by considering the effect of concurrent auctions for similar goods. In this model, a sophisticated bidder will incrementally raise her bid in two auctions until she becomes the highest bidder in one of the auctions by a small margin. If all bidders are sophisticated, this creates an incentive for squatting, while if there is at least one naïve bidder, sniping can be rational. Anwar, McMillan, and Zheng (2006) find empirical evidence of this: in their study, bidders who bid in multiple auctions for the same good paid on average 9% less for items they won. Moreover, Ely and Hossain show that squatting will not lead to an inefficient allocation of the auctioned items. Inefficient allocation is,

---

[3] The presence of an incremental bidder is the crucial feature of the escalation effect. The authors suppose that a price war begins only if the incremental bidder has not bid his valuation.

however, a possibility once some bidders snipe.[4] Ely and Hossain's and Roth and Ockenfels' models are satisfying because they empirically show benefits to sniping and offer explanations for why rational bidders would engage in such a practice. Nevertheless, their models rely on the existence of at least some naïve, incremental bidders in an auction.

Other literature has attempted to explain sniping through the lens of imperfect information, such as common values auctions or cases where bidders are uncertain of their own private values. Rasmussen (2006) and Hossain (2008) propose models where bidders have independent private values, but some bidders do not know their exact valuation and must incur a cost to learn it. These bidders submit a relatively low bid and reconsider whether to bid again each time they are outbid, effectively using other bids to "learn" their own private valuations. In response, experienced bidders may prefer to snipe, so they avoid giving away any information that could raise the bid of a competitor. Similar results have been predicted in auctions with common values, where a knowledgeable bidder's bid may convey information to less well-informed bidders. As outlined in Section 2, Bajari and Hortaçsu (2003) propose that sniping is an equilibrium behavior in second-price sealed-bid auctions with common values. Being outbid early in an auction could inform the lower bidder that she has underestimated the true value and cause her to raise her bid to something she previously believed was too high. In general, theories of information asymmetry are attractive because they do not rely on the presence of naïve bidders to give a rational basis for sniping.

While the presence of common values may theoretically provide an incentive to snipe, testing this hypothesis rigorously is possible thanks to relatively recent work formalizing the nature private and common values auctions. Milgrom and Weber (1982) propose models of

[4] This is shown in more detail in Section 2. In short, a sniper may cause bidders other than those with the highest valuations to win a group of simultaneous auctions for similar goods.

private and common values bidding in addition to offering qualitative descriptions of factors that might influence the presence of private or common values. They also propose what they call an "intermediate" model of affiliated values. While the authors link the common values framework with auctions for things like mineral rights and private values with non-durable consumer goods, they claim that an affiliated values model is suitable for goods like paintings. With affiliated values, bidders' valuations are correlated, but they still reflect the individual preferences of the bidder. This differs from a pure common values auction, where bidders' valuations are estimates of one true, unknown valuation.

Later work attempts to distinguish between private and common values empirically. Paasrsch (1992) derives models of private and common values and tests their applicability to tree planting auctions, yet as Haile, Hong, and Shum (2003) note, this method requires modeling bidding behavior based on one's assumption about the suitability of a private or common values framework. Instead, Haile, Hong, and Shum show that in first-price sealed-bid auctions with equilibrium bidding, the presence of common values is testable against the null hypothesis of private values. This test relies on the fact that the "winner's curse" is present under common values but not under private values. As more bidders arrive in a common values auction, there is an increased probability that the winner's estimate of the actual, unknown value of the item is too high. To compensate for this, bidders should decrease their bids as more bidders arrive. This is not the case under private values, since a bidder's valuation depends only on her own preferences and information. Under private values, then, there should be no relationship between the number of bidders and an individual's expectation of winning, conditional on the number of bidders, and bidders should not adjust their bids depending on the number of competitors of they face.

The authors acknowledge that tests using bidder participation require participation to be exogenous, even though this is often not the case. However, they demonstrate that the use of instrumental variables for bidder participation can help solve the problems created by endogenous participation. The authors also rely on the assumption of equilibrium bidding, complicating the application of their method to cases such as sniping, where bidding behavior may not always correspond to what auction theory would predict. On the other hand, their method is simplified in second-price auctions, like those on eBay.[5]

Bajari and Hortaçsu (2003) apply the tests developed above to eBay auctions for coins. Because they test for common values in second-price auctions, many of the problems associated with first-price auctions mentioned above do not apply. However, second price auctions do introduce what Haile, Hong, and Shum (2003) call a "missing data problem," since the winning bid in a second-price auction is unknown. Bajari and Hortaçsu speculate that the sale price should decrease with the number of bids in the same fashion that Haile, Hong, and Shum theorize that the mean bid should decrease with the number of bids. Following Athey and Haile (2002), who demonstrate that one can test for common values in a second price sealed-bid

[5] Pinske and Tan (2005) show that first-price auctions with affiliated (private or common) values can display a negative relationship between the number of bidders and the equilibrium bid, making it difficult to distinguish between common values and affiliated private values. With common values, additional bids suggest that the bidder's estimate of the true value of the object is too high. With affiliated values, potential winners see other bids as indications that their own private values are likely on the high end of the distribution of values, causing them to decrease their bids in order to win at a lower price. However, a negative relationship is more telling in second-price auctions like those on eBay, since in these auctions, private values imply a weakly dominant strategy to bid one's valuation.

auction by observing only some of the bids, Bajari and Hortaçsu examine the relationship between the levels of the observed bids and the number of bidders.[6]

Bajari and Hortaçsu also adopt the instrumental variables approach to endogenous bidder participation explained by Haile, Hong, and Shum (2003). They use the minimum bid as an instrument for participation. The authors find that the minimum bid is negatively correlated with the number of bids, and they assume that it is uncorrelated with any other factors that could affect the value of the bids. The latter assumption may not be particularly strong, given that lower minimum bids could simply be associated with less valuable items. Finally, although Bajari and Hortaçsu provide qualitative evidence that auctions for coins have a common values element and find statistical backing for their claim, they note that they merely reject a pure private values model. Indeed, as Goeree and Offerman (2002) claim, the extreme cases of pure private and common values are likely quite rare. Rather, most auctions probably have some elements of both cases. In the present study, I attempt to improve on Bajari and Hortaçsu's analysis by using a new instrument. I also test the relationship between the median bid—rather than the average bid, the closing price, or all the observed bids—and the number of bidders in order to limit the effect of the unknown winning bid. Recognizing that a mixture of private and common values is likely, I attempt to examine not just whether a private or common values model prevails, but also the magnitude of these effects.

Many previous studies have investigated sniping and squatting by placing bids in auctions or by tracking individual bidders, in addition to studying the characteristics of completed auctions. As Ariely, Ockenfels, and Roth (2003) note, one would obtain stronger

[6] Bajari and Hortaçsu (2003) also claim that Athey and Haile (2002) imply that under eBay's rules, the number of bidders should have a *positive* relationship with the levels of the bids, rather than no relationship, under the null hypothesis of private values.

results in an experimental setting where researchers can control for factors like bidders being unaware of their true valuations. Similarly, Einav, Kuchler, Levin, and Sundaresan (2015) note the difficulty of using real auction data: While heterogeneity among items, bidders, and sellers can obscure important bidding patterns, considering bidding for only certain items under certain conditions can make it difficult to generalize one's results. For reasons of practicality, I rely solely on data from completed eBay auctions in which I did not bid. Despite the challenges this presents, the size of the dataset should help solve the power problems seen in work like Gray and Reiley's (2007).

Instead of following bidders or conducting an experiment where I place bids myself, I rely on information about the good being sold and about each bidder's identity, experience, and bids in completed eBay auctions. Like in previous work, I create indicator variables to show sniping and squatting. I use them to study the prevalence of these behaviors and whether they have an effect on the probability of winning and on the price the winner pays. I compare the behavior I observe to the null hypothesis that there is no benefit to bidding at any particular time, and that bidders should simply bid their valuation when they see an item in which they are interested.

The present study will examine the prevalence of sniping and squatting in auctions for a number of different goods. The purpose of this part of the study will be to determine if either type of unexpected bidding is more common in auctions for goods where a common values framework may be appropriate. For each type of good that I study, I test whether it might be appropriate to assume common values and compare my finding to the prevalence of early or late bidding in auctions for that type of good. For each type of good, I will examine the effect of bidding behavior on the probability that a given bidder wins an auction. I then develop a new model of a two-bidder auction to examine the probability that a sniper wins because of the strategic benefit of sniping, and not because she was selected to have a high valuation.

Additionally, I see if auctions that were won by a sniper or squatter close at a price that is significantly different from auctions that are won by ordinary bidders. This will help evaluate the magnitude of Ely and Hossain's competition and escalation effects. I consider the effect of a bidder's experience on how likely she is to snipe or squat and how likely her strategy is to be effective. Even if sniping or squatting does not have an effect on the probability of winning an auction or the price the winner pays, a disproportionate amount of sniping or squatting among experienced bidders would suggest that individuals do derive some benefit from theoretically unexpected behavior.

This paper proceeds as follows. Section 2 gives greater detail to the theoretical basis for sniping and squatting. Section 3 explains my data collection process and summarizes my data. Section 4 discusses the results of my efforts to predict sniping and squatting, with particular attention to a private versus common values explanation. Section 5 examines the potential benefits of sniping and squatting by considering their relationships to the probability that a bidder wins an auction and to the closing price of the auction. Section 6 concludes.

2. Theoretical Background

When a bidder sees an auction on eBay, she observes information about the item for sale, the progression of the price of the item throughout the auction, and, except in rare cases where they are hidden, the identities of her opponents. Crucially, the bidder sees the price that the current high bidder would pay if no one were to outbid him, but she does not see that bidder's valuation. Rather, the price shown for the high bidder is the minimum increment above the second highest bidder's valuation. Since the highest bid at any given point is unknown to the other bidders, eBay auctions are very similar second-price sealed-bid auctions. In second-price sealed-bid auctions, it is a weakly dominant strategy to bid one's valuation. Assuming private values, the time that a sealed bid is placed also should not affect its magnitude, since bidders do not know the valuations of their competitors.

One possible equilibrium in eBay auctions is that all bidders will bid their valuations, as is traditionally predicted for second-price sealed-bid auctions. However, this is complicated by the fact that only the highest bid in an eBay auction is sealed. To be sure, this is the most important bid to know, since it is the one that must be beaten in order to win, but knowledge of lower bids could reveal information about how one's competitors behave, help bidders learn their own private valuation, or, in the case of common values, help improve their estimate of an item's true value. This additional information allows for incremental bidding in during the auction. As a result, bidding one's value at the beginning of an auction is just one equilibrium strategy. Bajari and Hortaçsu (2003) and Roth and Ockenfels (2000) show that other equilibria, including one where bidders place multiple bids and bid their valuations only at the last moment, exist.

The key to these analyses is dividing an eBay auction into two stages: The first comprises almost the entire auction, from the beginning to just before the end. The second is the last moment of the auction—the moment in which sniping takes place. During the first period, bidders have time to react to other bids that come in, while in the second period they do not. The second period of the auction is exactly like a second price sealed bid auction, since even if one were to see the last-minute bids, it would be impossible to use the information they might bring. Bajari and Hortaçsu (2003) also note that this two-stage model also implies that when a bidder stops bidding or "drops out" of the first stage, it does not necessarily indicate that the price has reached the bidder's valuation, as it would in a typical ascending auction.

Roth and Ockenfels' (2000) model incorporates another element of uncertainty in the brief final period of the auction. In their model, while bids are submitted with certainty in the first period, there is a possibility in the second period that a bid will not be processed in time to be counted. This is an apt modification; bidders frequently complain that their snipes are not processed in time, and online services compete by advertising their success in transmitting snipes. The authors show that it is an equilibrium strategy for bidders to bid the minimum bid at the beginning of the auction and then to bid their valuation in the last second. This strategy allows a bidder to increase the probability that, should she win, she must pay only the starting price or the minimum increment over the starting price. If two bidders submit the minimum bid at the beginning of the auction, eBay will count whichever bidder submitted the bid first as the current high bidder. When the players bid their valuations at the end of the auction, the bidder with the higher valuation could lose if her snipe fails and if the low bidder had submitted the minimum bid first. In this equilibrium, players risk that their final bid equal to their valuation will not go through. This creates the possibility that the bidder will lose against a competitor simply due to chance, rather than because she has the lower valuation than her competitor.

Alternatively, bidders can mitigate the risk that their high bids will lose simply due to chance by bidding their valuation in the first stage of the auction, eliminating the possibility that the bid will not be transmitted in time. One's opponent may know that one is following this strategy. For example, when an opponent arrives, he sees that there is one bidder and the standing price is the minimum bid. He must bid at least an increment above the minimum. If he does this, he will find out that the first bidder has not bid the minimum, since eBay's proxy bidding will raise the price to more than one increment above the minimum. In this case, the opponent may raise his bid immediately; since the price escalation has already begun, there is no reason for him to wait until the last second to bid and risk losing by chance. These are two of many possible equilibria, and while this two-bidder model is obviously much simpler than most auctions, it demonstrates how rational bidders could believe that sniping is beneficial.

Roth and Ockenfels show that it is ambiguous which of the above equilibria is more profitable for the winner. They derive a profitability condition relating *p*, the probability of a successful snipe, *H* and *L*, the bidders' valuations, *m*, the minimum bid, and *s*, the bid increment. In particular, if $p > 1/2\ [H\text{-}L]/[L\text{-}m\text{-}s]$, the sniping strategy should be an equilibrium, while the non-sniping strategy should not. This leads to a number of interesting observations. Auctions with a higher starting price and a higher minimum increment should see more sniping. Moreover, if *p* changes with time—perhaps increasing with changes in sniping technology or faster Internet connections, or decreasing with eBay's new use of CAPTHCA to prevent robotic bidding—one might see a corresponding change in the prevalence of sniping. Finally, for items where the difference between bidders' valuations is very high, one would expect it to be less likely for non-sniping equilibrium to be profitable, so sniping should be more common. This might be the case in a common values auction, where only "informed" bidders have a good estimate of the true value of the item.

Indeed, the incentive to snipe is particularly strong in auctions with common values. Bajari and Hortaçsu (2003) show that in a common values auction, it is an equilibrium for all bidders to snipe in order to avoid revealing information about the item's true value. Roth and Ockenfels' (2000) model is again slightly richer, as it accounts for uncertainty about the transmission of snipes as well as the probability an item is fake. They also take into account the uninformed bidder's valuation and the minimum bid. The uninformed bidder uses information from the informed bidder's actions to determine if her own valuation of an object is accurate. The authors claim that an informed bidder will never bid above her valuation for an object, so the presence of an informed bid shows an uninformed bidder that the object must have some value. In order to avoid a bidding war, the informed bidder may submit her bid equal to her valuation in the second stage of the auction, preventing her information from being shared with uninformed bidders. The authors also define the case when it is profitable for the uninformed bidder to snipe. In particular, if the ratio of the probability of an item being "fake" to the probability of an item being "genuine" is greater than the ratio of the uninformed bidder's valuation to the minimum price, then the uninformed bidder should not snipe, assuming there is no signal from the informed bidder that the object is genuine.

Ely and Hossain's (2009) model takes a different approach by predicting sniping as well as squatting in the context of competing price effects and simultaneous auctions. They speculate that the presence of a squatting bid will encourage potential entrants to consider simultaneous auctions for the same good in order to avoid competition and to win at a lower price. Intuitively, the earlier a bid is placed, the longer it has this deterrent effect. This is what they call the "competition effect" that promotes squatting. This effect trades off with the "escalation effect," which hurts the case for squatting and instead is an incentive to snipe. The escalation effect holds that entry into an auction can provoke a bidding war. The authors acknowledge that this relies on the existence of naïve bidders who do not bid their valuations: A bidder could win against a naïve competitor with a higher valuation if the bidder snipes and prevents the competitor from raising his bid in response.

Further, Ely and Hossain replace the two-period models of Roth and Ockenfels and Bajari and Hortaçsu by dividing eBay auctions into N+2 periods, where N is the number of bidders. Each bidder bids in her own period. The authors create a final period, N+1, when sniping can occur. Sophisticated bidders will submit two bids in their model. First, they will incrementally raise their bids in different auctions for the same good to find the one with the lowest high bid. A sophisticated bidder succeeds once her incremental bid makes her the high bidder in an auction. She then bids her valuation in that auction, maximizing her chance of winning the auction at the lowest price. The authors claim that naïve bidders will do the same thing but will not realize that bidding need not be incremental. As a result, once the naïve bidder becomes the high bidder in an auction, she will do nothing until she is outbid. Thus if all bidders are naïve, the authors predict auctions with sniping to have a lower closing price, while they make no such prediction if all bidders are sophisticated and bid incrementally only to uncover their competitors' bids.

Moreover, as mentioned earlier, the authors show that sniping can lead to an inefficient allocation of goods. If a sophisticated bidder gradually raises her bid in two auctions, Auction 1 and Auction 2, she will stop bidding once she is the highest bidder in the auction with the lower standing bid, say Auction 2. If the sniper bids in and wins Auction 1, and the sophisticated squatter wins Auction 2, then inefficient allocation is possible. In this case, the bidder whom the sniper beats in Auction 1 has a higher valuation than the sophisticated squatter who wins Auction 2. Thus the bidders with the highest and *third* highest valuations win the two auctions. This is even more clearly the case with naïve bidders who not bid their valuations and do not have the chance to respond to a snipe.

3. Data collection and description

Like in much of the previous literature, my data comes directly from eBay. I used several of eBay's application program interfaces (APIs) to generate lists of unique auction identification numbers, known as Item IDs. I narrowed my search to include only listings that had ended, so the entire bid history for the item would be available. I also excluded items sold in formats other than the ascending price auction eBay typically uses, notably the very common "Buy It Now" listings. The APIs are designed primarily for third-party commercial use, and no one API gives enough data about each bid, the time it was placed, and the bidder who placed it for academic analysis. To get more information about each item, I wrote a Python code to scrape data from the webpage corresponding to each Item ID. This code also took auctions that had no bids out of my data set.

The auctions I studied were drawn from five categories: U.S. Indian Head pennies, one-ounce Canadian silver coins, DVDs of the *Harry Potter* series, Intel 20i7 processors, and used models of the iPhone 6. In general, I looked for large markets for relatively homogenous goods. I chose these categories of goods with the hope that private and common values would be of different importance. I also speculated that the types of bidders (e.g., businesses, collectors, first-time bidders) might differ among the categories. I searched for my particular categories of items using eBay's APIs, just as one might look for a particular category using the search bar on the regular eBay site. This method is imperfect, since the search results show auctions for similar, but not identical, products. As in many previous studies, I use data about each individual auction to control for the heterogeneity in search results.

I chose the two categories of coins in order to have a group of items whose value might not be known without some level of expertise or research, which would allow me to analyze the effects of costly information acquisition and asymmetrical information between bidders on the prevalence and effectiveness of alternative bidding strategies, as in Bajari and Hortaçsu (2003) and Ockenfels and Roth (2006). The Canadian coins are relatively homogenous in terms of age, but there are relatively few observations. The Indian Head pennies, on the other hand, have a much broader range in age but are a much larger market, allowing for more observations and more powerful results. *Harry Potter* DVDs are widely sold and typically have an easily knowable retail value. I chose a specific computer processor in order to have very homogenous group and one with very expensive items. Unfortunately, this traded off with the sample size; there are relatively few observations in this group. Finally, the iPhone 6 market is relatively large and, although the products are all essentially the same, they have varying degrees of wear and tear, adding a degree of uncertainty to the item's true value.

My final data set sampled thousands of auctions that occurred in January and February 2017. After eliminating Buy It Now auctions, auctions not denominated in U.S. dollars, and auctions where all bidder information was private, my data set contained 1,973 auctions with 9,830 bidders who placed 18,452 distinct bids (i.e., not counting as separate bids all the times a bid was automatically raised). None of the auctions in the final data set had reserve prices. It is important to have information about individual bidders in order to track whether they submit multiple bids and to consider the effect of their experience on their behavior. As in much of the literature, I use a bidder's feedback rating as a proxy for her experience on eBay.[7] Table 1 shows the distribution of bids and bidders into the categories discussed previously.

[7] Each time a bidder completes a transaction, the seller can rate her as good, bad, or neutral, which correspond to changes in her feedback rating of +1, -1, or 0. Previous literature has noted that bidders generally receive positive feedback, so their rating is a decent measure of how many transactions they have completed.

Table 1. *Number of bids by item type*

	Frequency	Percentage
Processor	234	1.268155
Harry Potter	668	3.620204
Canadian coin	1722	9.332322
Indian Head	9396	50.92131
iPhone 6	6432	34.85801
Total	18452	100
Observations	18452	

Sellers set the duration and starting price of the auction in advance, and these are known to all bidders. Auctions can be 1, 3, 5, 7, or 10 days long. All of these lengths are represented in my data, although 7-day auctions are by far the most common, as demonstrated in Table 2. To account for the fact that the auctions are of varying lengths, I usually consider the time elapsed in an auction as the percentage of the auction that has elapsed.

Table 2. *Number of auctions by duration*

	Frequency	Percentage
1	14	.7095793
3	57	2.889002
5	359	18.19564
7	1379	69.89356
10	164	8.312215
Total	1973	100
Observations	1973	

The data show extremely clear patterns of bidding, with most bidding occurring at the beginning or end of an auction and very little in the intermediate hours and days. There is also evidence of multiple bidding, that is, cases where bidders ignore eBay's advice to bid their maximum willingness to pay when they first bid and instead bid different amounts at different times. Of the 18,452 individual bids collected, 8,622, almost 47%, were *not* the first bid the bidder had placed for a particular item. In some cases, incremental bidding does not even appear to be a response to being outbid; some bidders place many bids within seconds of each other. It is possible that they do this in anticipation of their lower bid being outbid, causing the higher bid to come into effect, but at a price below what they would pay if eBay automatically raised their bid by the minimum increment.[8] This is consistent with Ariely, Ockenfels, and Roth's (2003) speculation that bidders might have an incentive to bid slightly under their valuations in order to pay less than the minimum increment over the second-highest price.

The most striking feature of the data is the pattern of bidding during particular segments of auctions. Figure 1 shows the percentage of time in an auction that has elapsed before a bid is placed. Very few bids are placed during the middle 80% of the auctions. There is a small spike in bidding at the very beginning of the auctions that declines immediately. By far the clearest trend is the extreme increase in bidding activity at the end of the auctions, particularly in the last 1% of bidding time.

[8] This takes advantage of the fact that bidders must observe the minimum increment with respect to only the highest bid at the time they bid. Consider, for example, an auction with a starting price of $10.00 and a minimum increment of $0.50. If bidder 1 places a maximum bid of $20.00, bidder 2 sees that she must bid at least $10.50. If bidder 2 bids $20.01, she wins the auction while paying only $0.01, rather than $0.50, more than the second-highest bid.

Figure 1: Arrival of all bids

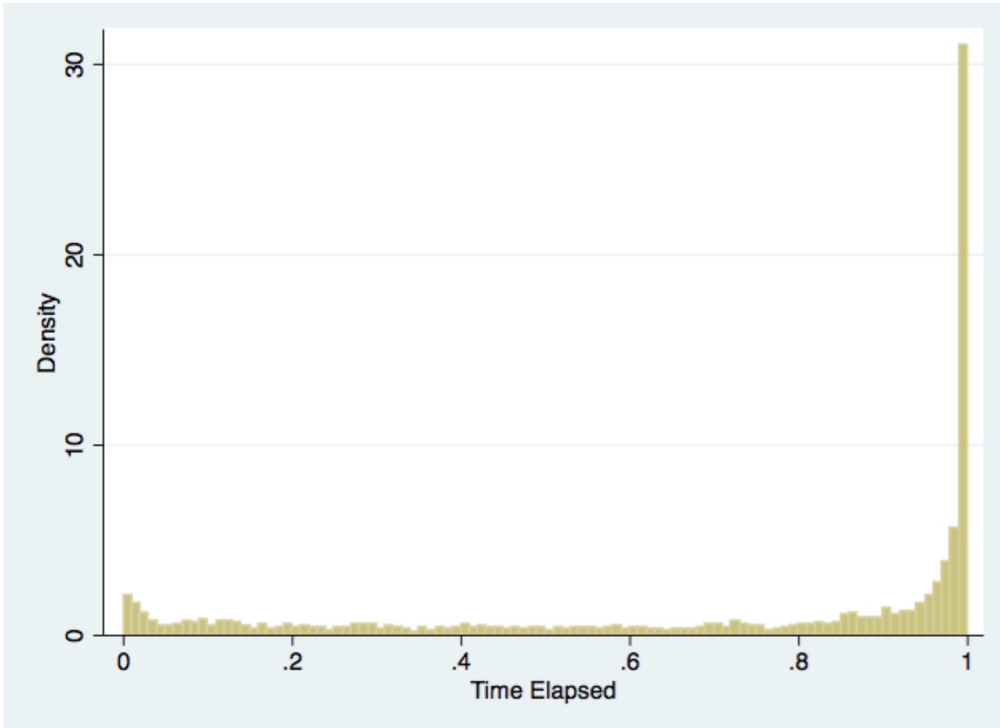

The average auction length was about 6.5 days, making the last 1% of the "average" auction about 1.5 hours. Most previous studies, however, have focused on sniping in the final minutes or seconds of auctions. Indeed, although there is a clear increase starting in around the last 20% of the time of most auctions, by far the largest spike comes in the last few minutes. Of 18,452 distinct bids, 3,353—just over 18%—came in the last ten minutes of an auction. 2,119 bids—over 11%—occurred in the last minute. 1,352—over 7%—came in the last ten seconds. 165—about 1%—came in the last second. In this paper, I define a snipe as a bid that is placed during the last ten seconds of an auction.

Figure 1 also shows a small but noticeable spike in activity in the first five percent of the auctions, but this spike declines quickly after the auction begins. At first glance, this appears to be evidence of squatting. Of all 18,452 bids, 2 were made in the first minute of an auction, and 54 were made in the first ten minutes. 263 bids came in the first hour of bidding. I believe it is appropriate to use percentages rather than fixed times for squatting. Squatting must be defined with a large enough window that some bidders have a chance to see the listing; it would be unreasonable to define a squatting bid as any bid that comes in during the first ten seconds of a listing. Moreover, squatters benefit not from minimizing the time gap between the beginning of the auction and their bid, but rather from being one of the relatively earlier bidders. In the following analysis, a squat is a bid placed in the first 5% of an auction.

4. Predicting sniping and squatting

4.1. Evidence of sniping and squatting

The spikes in bidding at the beginning and end of auctions shown in Figure 1 are preliminary evidence of sniping and squatting. If there were no benefit to bidding at a particular time, we may expect bids to arrive at a uniform rate throughout the auction. We would also expect the time the first bidder enters to be close to zero in most auctions. Indeed, for any given time interval, we would expect later intervals to be associated with the arrival of fewer first bidders, as bidding has started in progressively more auctions. Figure 2 shows the number of auctions whose bidding starts in a given time interval. As in Figure 1, time is measured from zero to one, as the proportion of time in the auction that has elapsed.

Figure 2: Arrival of first bidder

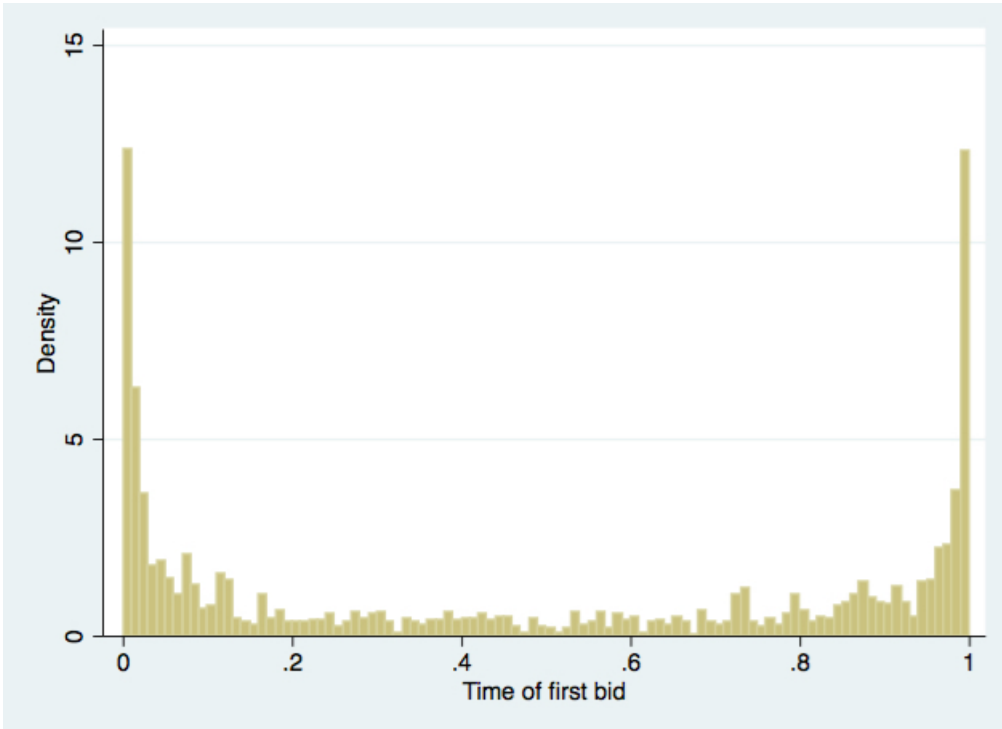

It is intuitive that we would see a spike in the arrival of the first bidder at the very beginning of an auction. While the extreme downward slope of the spike is striking, it is not surprising if one assumes that some bidders search eBay on a frequent schedule and bid as soon as they see an item. The surprising feature of Figure 2 is the spike on the right side of the chart, showing that there are many auctions that go without any bidders until the last few hours of bidding. This could be explained if there are committed potential bidders who monitor auctions with no bidders and enter at the last moment, perhaps hoping to be the only bidder. This is the first evidence we see that auctions with apparently little competition could attract snipers.

While we see clear graphical evidence of sniping, it is less obvious that squatting is a widely used strategy. Figure 1 shows a relatively small increase in bidding at the beginning of the auctions, and, as mentioned above, the large spike in Figure 2 should not be taken as evidence that bidders arrive early for strategic reasons. Figure 3 is very similar to Figure 1, but while Figure 1 shows the time that all *bids* are placed, Figure 3 shows the time that all *bidders* place their first bid. In theory, the presence of squatting should be clearer in Figure 3, since any early bidder could be considered a squatter. I demonstrate the existence of a bias toward bidding at the beginning of an auction by rejecting the null hypothesis that bidders arrive at a uniform rate. A simple way to test this hypothesis is to compare the frequency of bidders arriving in the first and second halves of the auction. If the number of bidders arriving in the first half is statistically significantly greater than the number of bidders arriving in the second half, I can reject the null that bidders arrive at a uniform rate or favor bidding at the end.

Figure 3: Arrival of all bidders

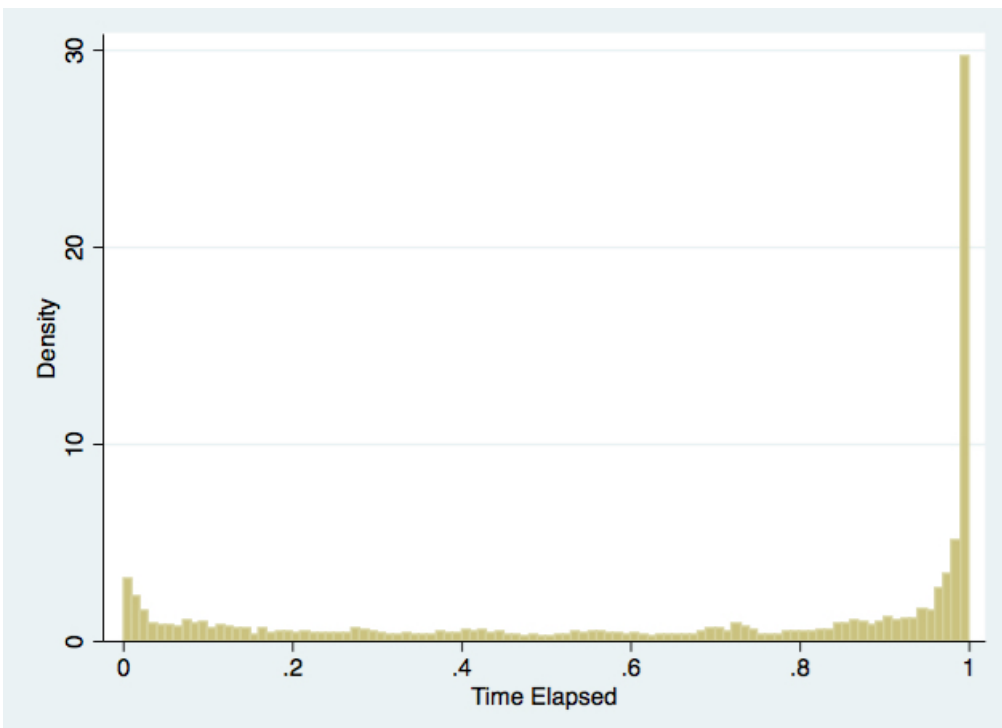

Figure 3 makes it clear that breaking the sample into halves will not show that more bidders arrive in the first half than in the second half, since so many bidders—indeed, apparently snipers—arrive in the last few percentiles of an auction. In order to exclude sniping from this part of the analysis, I break the sample into thirds and compare the number of arrivals in the first third of the auction to the number of arrivals in the second third of the auction. This method has the potential to be misleading, as the researcher could break up the data into segments of whatever size yields the desired result. As the sample is broken down into more parts, the possibility of type one error increases. This difficulty is addressed by Wolak (1989), who proposes a more rigorous method of testing multiple inequalities. Although I do not follow Wolak's method here, I believe my method is relatively sound given that I break the data into the largest possible segments that exclude sniping and compare only the frequencies in the first and second thirds. As Figure 3 suggests, breaking the data into thirds yields a clear result: 2,414 distinct bidders arrived in the first third of their respective auctions, while only 1,237 arrived in the second third. (By comparison, 6,179 arrived in the last third.) A t-test of the difference in the probabilities that a bidder arrives in the first third versus the second third gives a t-statistic of 19.1634, which is significant at a level lower than 1%.

While the apparent bias to bid close to the beginning of an auction could suggest that bidders may do so strategically, the hypothesis of strategic squatting is not supported by the fact that squatting bids are relatively unlikely to win auctions, as shown in Figure 4. Indeed, later analysis will demonstrate that squatting does not appear to be a successful strategy in increasing the probability of wining or lowering the sale price, either. It is therefore doubtful that such a disproportionate number of bids come in during the first few hours of an auction for strategic reasons.

Figure 4: Time that winner places first bid

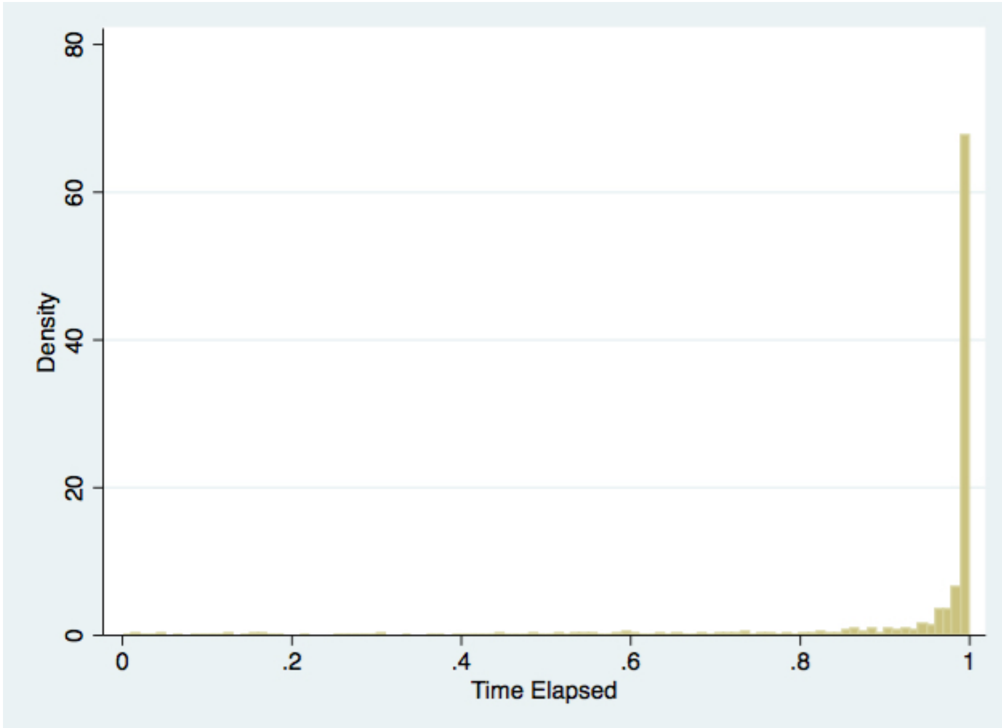

4.2. Sniping and squatting for different types of goods

I now turn to predicting when these behaviors occur. Much of the past literature has proposed two types of characteristics that are related to the presence of sniping and squatting: characteristics of the auction itself and characteristics of individual bidders. The presence of private or common values is one of the best examples of an auction-level characteristic. As Bajari and Hortaçsu (2003) and Ockenfels and Roth (2006) suggest, sniping may be particularly attractive in common values auctions. Another auction-level characteristic that could lead to sniping is the presence of incremental bidders or of other snipers, as suggested by Ockenfels and Roth (2006) and Ely and Hossain (2009). Ely and Hossain also propose that incremental bidding followed by squatting may be common in markets where there are concurrent auctions for many substitutable goods. On the bidder level, many past studies have found that experienced bidders tend to snipe more than inexperienced bidders. In this section I test some of these theories, beginning with whether the prevalence of sniping and squatting is different for different categories of goods. Table 3 summarizes my findings.

Table 3. *Squatting and sniping by item type*

	Frequency	Percentage
Processor		
Squats	28	11.96581
Snipes	4	1.709402
All bids	234	100
Harry Potter		
Squats	66	9.88024
Snipes	35	5.239521
All bids	668	100
Canadian coin		
Squats	183	10.62718
Snipes	48	2.787456
All bids	1722	100
Indian Head		
Squats	548	5.832269
Snipes	711	7.56705
All bids	9396	100
iPhone 6		
Squats	309	4.804104
Snipes	554	8.613184
All bids	6432	100
All categories		
Squats	1134	6.145675
Snipes	1352	7.327119
All bids	18452	100
Observations	18452	

I test the statistical significance of category-based differences with respect to sniping using an analysis of variance test. This yields an F statistic of 31.30, indicating that there is a highly statistically significant difference among the categories in terms of the prevalence of sniping. The differences in squatting by category are also statistically significant, with an F statistic of 28.04.

Sniping is by far the most common for iPhones and for Indian Head pennies, while squatting is relatively rare for these items. It is interesting to note that sniping is a good deal more common for Indian Head pennies than it is for the Canadian coins, which is surprising given that these items are relatively similar. However, the Indian Head pennies are antiques, while the Canadian coins are more recently minted collectors' items. This difference may create more uncertainty about the value of the Indian Head pennies. The difference in the prevalence of sniping could be explained if Indian Head pennies have a stronger common values element, as is explored in the following subsection.

4.3. Sniping in private and common values auctions

In this section, I test whether private or common values frameworks apply to each of the five categories of goods I observe. As Bajari and Hortaçsu (2003) note, both the qualitative characteristics of the object being auctioned and the relationship between the value of the bids and the number of bids can give clues as to the appropriateness of a private or common values model. Milgrom and Weber (1982) propose a number of qualitative considerations that can call into question the validity of a private values model, including the possibility that the item being auctioned has an unknown resale value, that it is not certain to be authentic, and that bidders may associate winning the item with some kind of "prestige." While Milgrom and Weber use the example of a painting to demonstrate the plausibility of these considerations, the same criteria can be applied to items such as coins and antiques, items that many studies use as classic examples for common values. Bajari and Hortaçsu (2003), for example, propose—and find supporting empirical evidence of—a common values framework for coins.

Following Bajari and Hortaçsu, I speculate that both the Canadian coins and Indian Head pennies will have strong common values elements. If most buyers of these coins are collectors, then Milgrom and Weber's "prestige" value may play an important role in bidding. These coins are relatively well known and widely traded, so it is not very costly for buyers to determine a fair price, although such collectibles could change in value significantly over time. The coins are relatively common and none are extremely old, so the likelihood that a coin is fake may be low. However, fake coins are still a possibility, and it is entirely plausible that, especially for older Indian Head pennies, the quality of the item is unknown, creating a consideration similar to one about the realness or fakeness of the item.

The computer processors and *Harry Potter* DVDs, on the other hand, are likely to fit the private values models better. As Milgrom and Weber note, non-durable consumer goods are generally good fits for a private-values model. These items have clear retail prices and are quite similar within their groups. Unlike an especially rare mintage of one of the coins above, individual computer processors and DVDs are unlikely to have particularly important "prestige" values to collectors and have much less ambiguity in terms of resale value. For these reasons, I speculate that the processors and DVDs will have less strong common values elements than the coins.

A more complex case is that of used iPhones. While these, like computer processors and DVDs, are non-durable consumer goods, other properties of used iPhones make them an interesting case. First, the quality of used iPhones might vary much more than computer processors or DVDs. Put simply, there are many potential problems with used iPhones, whereas something like a used DVD has essentially one: whether it is scratched. Aside from the uncertain quality of the product, it is also possible that many buyers of iPhones on eBay refurbish the phones professionally and wish to resell them after, say, fixing a cracked screen. In this case, the bidder may have a rough idea of the resale value of the phone but still lacks perfect information.

For these reasons, I would expect a relatively strong common values element to bidding for iPhones.

As in Bajari and Hortaçsu (2003), I test my intuitions about private and common values empirically. Since the highest bid is never known in eBay auctions, I use the median, rather than the average, bid as a signal of the equilibrium bid. In order to prevent the sale price—which is not related to the highest valuation—from affecting the signal, I limit this part of my study to auctions with three or more bidders. In the case of multiple bidders, I use only the highest bid from any given bidder in my analysis. Again like Bajari and Hortaçsu, I use an instrumental variable to mitigate the endogeneity created by the fact that auctions with more bidders should generally close at higher prices as a result of increased competition. However, instead of using the minimum bid as an instrument for participation, as Bajari and Hortaçsu do, I use the day of the week on which an auction ends. I assume that the day of the week has no effect on a bidder's willingness to pay for an item. The regressions in Table 4 show the relationship between the number of bidders and the day on which an auction closes. In each of the five categories of goods, the day of the week is a significant predictor of the number of bidders.

Table 4. *Number of bidders and day auction ends*

	Processor	Harry Potter	Canadian coin	Indian Head	iPhone 6
Day of end					
Monday	-9.227***	-0.0794	-0.669**	-0.581**	1.435***
	(0.841)	(0.615)	(0.206)	(0.197)	(0.252)
Tuesday	-4.227***	2.192**	0.147	-0.752***	-0.407
	(0.742)	(0.682)	(0.182)	(0.179)	(0.312)
Wednesday	-2.701***	-0.312	-0.401+	-1.650***	1.429*
	(0.530)	(0.552)	(0.240)	(0.185)	(0.643)
Thursday	-5.132***	-0.404	-0.430+	-1.346***	-1.550**
	(0.512)	(0.745)	(0.257)	(0.208)	(0.520)
Friday	-7.227***	0.178	-0.0945	-0.998***	-0.635
	(0.912)	(0.675)	(0.269)	(0.216)	(0.498)
Saturday	-5.227***	-0.928	-0.0974	-0.805***	2.562***
	(0.786)	(0.586)	(0.212)	(0.172)	(0.231)
Constant	12.23***	7.222***	8.171***	8.154***	8.483***
	(0.291)	(0.427)	(0.140)	(0.106)	(0.153)
F statistic	38.47	3.940	3.810	16.53	32.16
Prob > *F*	3.41e-24	0.000848	0.000914	5.97e-19	1.28e-37

Standard errors in parentheses
The dependent variable is the number of bidders in a given auction.
The base case is an auction that ends on a Sunday.
Note: Excludes auctions with fewer than three bidders
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Despite the joint significance of the day of the week for all items, this instrument does have a drawback in that it is difficult to explain why the day an auction closes has the observed effect. One would expect a higher number of bidders on weekends, when bidders are less likely to be working, but this is not always the case. One might also expect the effect of closing day to be consistent across categories, yet only Thursdays have a consistent effect all five categories. It is, however, encouraging that across all categories, most closing days have fewer bids than do Sundays, a day one might expect to have particularly active bidders.

Table 5 summarizes the results of a two-stage least squares regression showing the relationship between the number of bidders and the median bid, using as an instrument the day of the week the auction ends.

Table 5. *Median bid and number of bidders*

	Processor	Harry Potter	Canadian coin	Indian Head	iPhone 6
Bidders	-5.882**	-2.997*	-27.72***	18.30***	-9.563***
	(-3.22)	(-2.19)	(-3.65)	(11.35)	(-13.31)
Constant	268.3***	41.65***	276.5***	-107.2***	277.5***
	(15.09)	(4.22)	(4.54)	(-8.94)	(40.90)
Observations	110	286	1132	4639	2800

t statistics in parentheses
The dependent variable is the median bid in an auction.
The instrument for *Bidders* is the day of the week the auction ends.
Note: Excludes auctions with fewer than three bidders
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

For all categories, the coefficients on the number of bidders are statistically significant at the 10% level or below. In four of the categories, I find negative coefficients, suggesting that I can reject the null hypothesis of private values for the computer processors, *Harry Potter* DVDs, Canadian coins, and iPhone 6. Indian Head pennies, on the other hand, appear to have private values. Although many of these results are inconsistent with my earlier guesses, they may indicate that many eBay auctions have a mixture of private and common values elements, as Goeree and Offerman (2002) suggest may be likely. Furthermore, even though the negative relationship between bidders and signals indicates that many auctions have some accounting for the winner's curse, we can also consider the scale of this effect. The confidence interval for iPhones is quite small, and the coefficient is quite large and positive. This gives us good evidence that iPhones have a relatively large common values effect, at least compared to the processors and DVDs, as we expected. The evidence of common values for iPhones is consistent with the high rate of sniping in iPhone auctions.

One of the most surprising results is the difference between the large, positive coefficient for Indian Head pennies and the large, negative coefficient for Canadian coins. This result may cast doubt on the validity of my instrument, or on my assumption that all types of coins have the same common values characteristics. Another complicating factor is that the prevalence of sniping in the markets for these types of coins is reversed from what theory would suggest. Indian Head pennies have very high rates of sniping, and yet we cannot reject the null of private values in their case. The opposite is true for Canadian coins, which have once of the lowest rates of sniping but do apparently have a common values element.

4.4. A prediction model for sniping

Past literature has suggested that the incentive to snipe or squat may be affected by factors other than the nature of the item being sold or existence of common values. In this section, I follow Ockenfels and Roth (2006) in estimating the probability that a given bid is a snipe or a squat. I use a probit regression to model the probability that a given bid is a snipe depending on auction-wide characteristics (the number of bidders and the starting price), bidder-specific characteristics (her feedback rating), and bid-specific characteristics (whether the bid was the bidder's first bid in the auction). Table 6 summarizes the results of the regressions, organized by item category.

Table 6. *Probability that a given bid is a snipe*

	Processor	Harry Potter	Canadian coin	Indian Head	iPhone 6
First time	0.00886	0.498**	0.219	0.245***	0.342***
	(0.02)	(2.73)	(1.53)	(5.89)	(7.22)
Rating	-0.0000263	-0.0000241	0.00000881	0.0000310***	0.000193***
	(-0.21)	(-0.52)	(0.24)	(5.28)	(4.72)
Bidders	-0.116	-0.0151	-0.00216	-0.0248***	-0.00586
	(-1.13)	(-0.54)	(-0.08)	(-4.76)	(-0.80)
Starting price	0.000430	-0.00390	-0.0000708	0.00281**	0.00310***
	(0.14)	(-0.35)	(-0.02)	(2.87)	(7.23)
Constant	-1.274	-1.830***	-2.057***	-1.479***	-1.917***
	(-1.22)	(-7.05)	(-7.92)	(-29.92)	(-16.84)
Observations	234	665	1717	9276	6389

t statistics in parentheses
The dependent variable is the probability that a given bid is a snipe.
Note: This table shows the results of the latent model, not the marginal effects. I examine only on the signs and the significance of the coefficients.
Note: Excludes observations with hidden ratings
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

For DVDs, Indian Head pennies, and iPhones, bids that are a bidder's first are much more likely to be snipes than other bids. In addition, Indian Head pennies and iPhones have a small, but still positive and significant, coefficient on the bidder's feedback rating, suggesting that experienced bidders in these categories may be slightly more likely to snipe. The relationship between first-time bidding and sniping may also reflect a relationship between experience and sniping, as experienced bidders may be more likely to heed eBay's advice not to bid multiple times. The fact that first-time bids are more likely to be snipers is also consistent with the theory that potential snipers monitor auctions and enter only those that they believe they have a high chance of winning.

Sniping in prevalent among experienced, non-incremental bidders in many of the categories that I *expected* to have a significant common values element. As with the common values analysis, auctions for iPhones have particularly consistent results, suggesting that informed bidders (perhaps experienced refurbishers) find that it is beneficial to hide their estimates of the phone's value from competing refurbishers. However, the inconsistency of the relationship between sniping, experience, and common values in the other categories casts doubt over the iPhone finding.

Finally, there is some evidence supporting Roth and Ockenfels' condition for the probability of a successful snipe that is necessary to make sniping beneficial. They claim that a higher starting price should lower the minimum probability of transmission for which it is beneficial to snipe, thus making sniping more common. In cases where the coefficient on the starting price is significant, it is small but positive, suggesting that higher starting prices are indeed associated with more sniping.

The clearest and most surprising result in Table 6 is that, for some goods, whether a bidder has bid previously is an important predictor of whether that bidder snipes. Below, I test the hypothesis that first-time bidders are more likely to be snipers because snipers monitor auctions and bid (very late) only in those with little competition. I do this by predicting the probability that an auction has at least one snipe based on the number of bidders who had arrived before 95% of the auction who had elapsed. The results are presented in Table 7.

Table 7. *Probability that an auction has at least one sniper*

	Processor	Harry Potter	Canadian coin	Indian Head	iPhone 6
Number of non-sniping bidders	-0.0487	-0.679***	-0.448***	-0.523***	-0.267***
	(-0.13)	(-3.42)	(-4.38)	(-13.33)	(-6.70)
Bidders	-0.0832	0.827***	0.417***	0.577***	0.298***
	(-0.22)	(4.24)	(4.63)	(15.99)	(8.75)
Starting price	-0.00193	0.0203	-0.0124	0.00402*	0.00198
	(-0.35)	(1.21)	(-0.94)	(2.23)	(1.62)
Constant	0.534	-2.260***	-1.178**	-1.404***	-1.086***
	(0.30)	(-5.11)	(-2.86)	(-17.08)	(-3.56)
Observations	15	94	167	1242	455

t statistics in parentheses
The dependent variable is the probability that an auction has at least one sniper.
Note: This table shows the results of the latent model, not the marginal effects. I examine only on the signs and the significance of the coefficients.
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

For each type of good but the processor, there is a negative and statistically significant relationship between the number of "non-snipers" in the first 95% and the chance that an auction has a sniper. The coefficient for computer processers is insignificant likely because of the relatively small sample size and limited incidence of sniping in this market. Overall, this result is consistent with the theory that snipers act as they do in order to take advantage of limited competition. It is also evidence against the effectiveness of squatting. Ely and Hossain (2009) predict that squatting is effective because it decreases the number of bidders in an auction. I show in Section 5 that squatting does not appear to decrease the number of bidders, but even if it did, decreasing the number of bidders may backfire by making the auction particularly attractive to snipers.

4.5. A prediction model for squatting

I perform an analysis similar to that in in Section 4.4 in order to find the probability that a bid is a squat and the probability that an auction has a squat, taking into account a bidder's rating and the starting price. Tests of squatting are by definition more limited than tests of sniping, since it does not make sense to predict squatting based on the total number of bidders in an auction, many of whom enter *after* the squatter. Neither model yields telling results. An increase in the starting price negatively affects the probability than an auction has a squat, but this may be due to the fact that higher starting prices will weed out some bidders if their valuations are too low to making competing worthwhile. A bidder's experience has no significant effect on whether she squats. Thus we still have little evidence that the cluster of bids we observe at the beginning of an auction occurs for strategic reasons. The strategic reasons for both sniping and squatting are explored further in Section 5, which attempts to define how successful these behaviors are in terms of winning auctions and doing so at lower prices.

5. Benefits of sniping and squatting

5.1 Sniping and the probability of winning

It is difficult to estimate the exact effect of sniping on the probability of winning. Indeed, the probability that a given bid wins an auction depends on when the bid was placed. After all, bids placed late in an auction are more likely to win because only participants with valuations high enough to beat all the competitors will bid late. Thus we should expect that sniping bids have a higher probability of winning even if there is no strategic benefit to sniping. This relationship is made more complicated by eBay's automatic bidding function. Although late bids must exceed the listed price at the time of bidding, they may or may not exceed the highest bid.

Thus, unlike in an English auction, the last bid placed in an eBay auction is not *guaranteed* to win. Squatting bids should face the opposite effect: since early bids are not as heavily selected to have high valuations, they should have a lower probability of winning than other bids, all else being equal.

Ely and Hossain (2009) estimate the effect of sniping and squatting on the probability of winning by placing squatting and sniping bids and seeing how often they win. This method allows the authors to find the effect of both strategies relatively accurately because they know their own valuations and therefore can control for the fact that late bids are associated with higher valuations. In the following analysis, I attempt to find the effect of sniping and squatting on the probability of winning *without* knowing the valuation of the winner, beginning with a simple probit regression of the probability that a bid wins an auction based on the time the bid is placed and whether it is a squat or a snipe.

Table 8. *Probability a given bid wins an auction*

	P(winner)	
Snipe	1.237***	(31.26)
Squat	-0.224^{+}	(-1.88)
Time	1.37e-09***	(16.55)
Close-Start	-0.00224***	(-10.36)
Auto-raised	-0.251***	(-8.07)
Rating	0.00000265	(0.42)
First time	0.188***	(6.39)
Harry Potter	0.0772	(0.49)
Canadian coin	0.00329	(0.02)
Indian Head	-0.0934	(-0.64)
iPhone 6	-0.359*	(-2.47)
Constant	-1.788***	(-11.80)
Observations	18281	

t statistics in parentheses
The dependent variable is the probability that a bidder is the winner of her auction.
Time is the percent of the auction elapsed before a bid is placed.
Close-Start is the difference in closing price and starting price of the auction.
Auto-raised is one if the bid was automatically-raised.
Note: This table shows the results of the latent model, not the marginal effects. I examine only on the signs and the significance of the coefficients.
Note: Excludes observations with hidden ratings
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Tables 8 shows roughly what we would expect. Even accounting for *Time*, the percentage of the auction that has elapsed before a bid is placed, sniping bids are much more likely to win than other bids. On the other hand, squatting bids are slightly less likely to win, but this is only weakly significant. [9] While the negative effect of squatting is not as strong as we may have expected, neither of these results allows us to separate the strategic effects of sniping and squatting from the effects of simply being an early or late bid.

To combat the selection for higher valuations among late bidders, I consider the effect of sniping in auctions with only two bidders. In this case, when the second bidder sees an auction,

[9] With all other variables at their means, the marginal effect of being a snipe on the probability that a given bid wins is estimated to be an increase of about 16 percentage points. The marginal effect for squatting is a decrease in the chance of winning by about 2 percentage points.

the listed price will be the starting price, although she will know she has a competitor. Thus she also knows the minimum price at which she can win the item is the starting price plus the minimum bid increment, which is usually quite low. This does not completely eliminate the selection for higher-value bidders that we see in auctions with many bidders, but the selection effect should be smaller for the second bidder than for any other bidder.

If there is no strategic benefit to sniping, then the second bidder should be no more likely to win when she snipes than when she does not. One potential problem with this model is that an early second bid could significantly raise the price of the item, deterring other bidders from entering. Another potential failing is that this model does not account for the benefits of sniping due to information asymmetries. In the case of common values, for example, early bids could signal that the item is valuable and trigger more entry. If an early second bid *does* trigger more bidding, that auction is excluded from the analysis. On the other hand, cases where an early bid *would have* triggered more bidding, but did not because the second bidder sniped, are included.[10]

What this model does capture clearly is the benefit of sniping in terms of not giving an opponent time to react to one's bid. If both bidders bid their valuations, the second bidder wins if she has a higher valuation, regardless of whether she snipes. However, if the first bidder does not bid his valuation, the second bidder is becomes the high bidder as long as her bid is higher than that of the first bidder. Once the first bidder is notified that he has been outbid, he may bid a higher value closer to his true valuation. This is extremely difficult to do if a sniper outbids him: Not only must the first bidder be watching his computer at the end of the auction in order to learn that he has been outbid, but he must also react in a very short period of time. I test my model

[10] This is only a problem if we believe that that having two bids instead of one makes a meaningful difference in the signal communicated to the other bidders.

with a probit regression predicting whether a bidder wins depending on whether she was the first or second bidder and whether she snipes. The results are shown in Table 9.

Table 9. *Probability of winning a two-bidder auction*

	P(win)	
Winner is:		
First bidder & sniper	1.000^{**}	(2.90)
Second bidder & non-sniper	0.425^{***}	(4.27)
Second bidder & sniper	1.192^{***}	(7.87)
Constant	-0.820^{***}	(-11.38)
Observations	847	

t statistics in parentheses
The dependent variable is the probability that a bidder is the winner of her auction.
Note: Includes only auctions with two bidders
Note: This table shows the results of the latent model, not the marginal effects. I examine only on the signs and the significance of the coefficients.
Note: Base case is the probability that the first bidder is a non-sniper and wins
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Even by simply comparing the signs of the coefficients in Table 9, we see that the second bidder is generally more likely to win than the first bidder and that snipers are generally more likely to win than non-snipers. This provides evidence that it is beneficial for both bidders to snipe. My finding is consistent with the hypothesis that sniping is beneficial because it does not give an opponent the chance to respond if he did not initially bid his valuation. The fact that snipers are more likely to win also supports the hypothesis that potential snipers monitor an auction with relatively little competition, and then bid in those where there has been little competition up until the last few seconds. The first bidder may also benefit from monitoring the auction. He may have placed one low bid in the hope that no other bidders would arrive, but in cases of little competition, he may anticipate a sniper and monitor the end of the auction himself. In general, this simple model provides some evidence that at least part of the benefit of sniping comes from robbing a naïve bidder who has not bid his valuation the chance to react. This could also apply to bidders who learn by bidding. As Hossain (2008) suggests, bidders who are aware of their private values may wish to snipe in order to prevent other bidders from learning their own private values and increasing their bids.

5.2. Squatting and the probability of winning

The preliminary regression in Table 8 showed that squatting bids are weakly significantly less likely to win an auction than other bids are, but we cannot separate a potential benefit of squatting from the overall decreased probability that an early bid wins. A better way to test the effectiveness of squatting is to test whether it has the effect that makes it attractive: Squatting is thought to deter other bidders from entering and driving up the price. Thus if auctions with squatting bids also have fewer bidders, we might infer that bidders are in fact deterred.

My findings do not support this theory. For each category of good, the presence of a squatting bid raises the expected number of bidders by two to four bidders. Across all auctions, those with a squatter had a median of eight distinct bidders, while those without a squatter had a median of three distinct bidders. Indeed, this is what we would expect: Auctions with later entry times are likely to have fewer bidders, assuming bidders arrive uniformly. This may also be explained in part by the phenomenon observed earlier that *sniping* is more common in auctions where there are few bidders during the bulk of the auction. Even if squatting deters some bidders from entering in the middle period of the auction, this could actually *attract* competition at the end of the auction in the form of sniping.

5.3 Sniping, squatting, and closing price

Finally, I consider the relationship between sniping, squatting, and the closing price. Previous studies have also considered these relationships, but many, such as Ely and Hossain (2009), have attempted to calculate the surplus of the winner under different strategies. I do not calculate surplus, since this would require knowing the valuation of the winner. However, on average, changes in surplus due to bidding strategies should equal changes in price due to bidding strategies.

Below, I consider two models of the relationship between the time the winning bid is placed and the closing price of the auction. The first models time discretely, with indicator variables for whether the winning bid was a snipe or a squat. The second models time as a cubic function but does not have specific variables denoting sniping or squatting. The cubic function does, however, give more weight to times very close to the end of the auction, which could account for sniping. Both models have the same controls, including the number of bidders, the minimum bid, and the percentage of time elapsed before the first bid was placed. In all cases where continuous time is relevant, it is measured from zero to one, as the percentage of the time in the auction that has elapsed.

Table 10. *Sniping, squatting, and closing price*

	Discrete Time		Continuous Time	
Snipe	-15.62***	(2.977)		
Squat	13.31	(16.18)		
First time	3.322	(2.787)	0.549	(2.765)
Time of first bid	6.671	(4.489)	11.28*	(4.711)
Starting price	0.857***	(0.0336)	0.855***	(0.0336)
Bidders	12.36***	(0.536)	12.48***	(0.549)
Duration	-2.623**	(0.935)	-2.811**	(0.937)
Rating	-0.000366	(0.000532)	-0.000644	(0.000530)
Harry Potter	-96.20***	(16.77)	-96.46***	(16.79)
Canadian coin	-97.34***	(16.07)	-95.55***	(16.11)
Indian Head	-92.70***	(15.80)	-93.40***	(15.83)
iPhone 6	-38.46*	(15.07)	-40.73**	(15.07)
Time			-222.1^{+}	(113.8)
Time squared			526.1*	(229.9)
Time cubed			-344.2**	(131.1)
Constant	78.66***	(17.19)	110.9***	(22.57)
Observations	1942		1942	
Adjusted R^2	0.746		0.745	

Standard errors in parentheses
The dependent variable is the closing price of the auction.
Note: Excludes observations with hidden ratings
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Both models have similar predictive power and show similar results. As with previous models, there does not appear to be a benefit to squatting, as auctions won by squatting bids typically close at higher prices. The sniping indicator, however, has a significant negative coefficient, showing that auctions won by a sniper are expected to close more than $15 lower than those won by an ordinary bidder. Moreover, this model controls for the number of bidders. Therefore, while some of the benefit of sniping may come from snipers simply seeking out auctions with few bidders and little competition, this result shows benefits related to sniping even when there are more bidders and presumably more competition. In turn, this result supports the theories that snipers act in order to hide information and prevent bidding wars. It is important to note, however, that this theory is not mutually exclusive with the theory that snipers look for uncompetitive auctions.

Similarly, the continuous time variables have moderately to highly significant effects. Cubic time is strongly negative, indicating that bids at times closer to the end of an auction are expected to close at lower prices. I believe a cubic model is justified because any higher-order effects of time beyond the cube cause the lower-order effects to become insignificant. In this model, the first order effect is barely insignificant at the 5% level, with a p-value of -1.95. The cubic function is graphed in Figure 5, where the *y*-axis is the predicted closing price and the *x*-axis is the proportion of the auction that has elapsed before the winning bid is placed; this ranges from zero to one.[11]

Figure 5: Relationship between time of winning bid and closing price

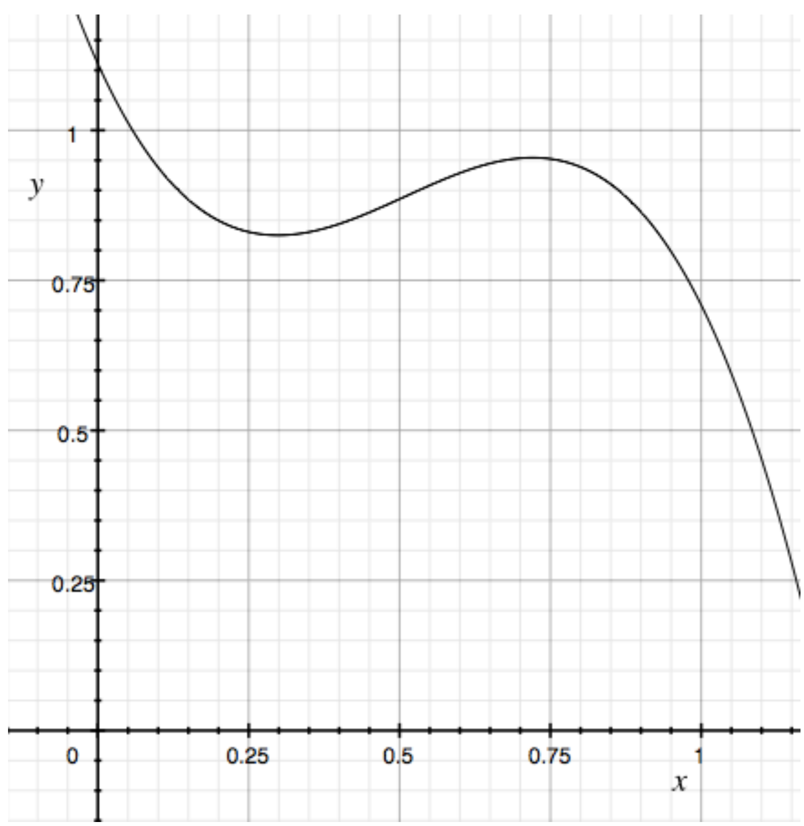

[11] The closing price is scaled down by a factor of 100 for readability.

As Figure 5 shows, there is an overall trend of later bids commanding lower closing prices. This "discount" in closing price as a result of time appears to increase in roughly the first quarter of the auction, perhaps evidence of a benefit to squatting. During the middle 50% of the auction, the discount appears to shrink, which could suggest that it is disadvantageous to bid during the middle of the auction. This is consistent with the pattern seen in Figure 1, where the middle of the auction saw the least bidding. In the last quarter of the auction, the discount increases again, such that the lowest prices are associated with winning bids placed at the very last minute of an auction. This may be evidence not only of the benefit of sniping, but may also show that sniping is a more effective method of winning an auction at a lower price than squatting is. However, it is important to remember that this general shape is expected as a result of cubic model I proposed. The confidence intervals on the coefficients of the cubic function are very wide, so this relationship should not be taken as conclusive evidence.

Still, the strongest result shown in Table 10 is that sniping bids win auctions at significantly lower prices, even controlling for the total number of bidders. While this suggests that snipers benefit even in auctions with many bidders (and presumably more competition), a variation of the model also supports the theory that snipers seek auctions in which there are few bidders, making them likely to win at a favorable price. Table 11 summarizes the results of a model similar to the discrete time model summarized Table 10. Here, though, I consider only auctions with at least one sniper and investigate the relationship between the number of snipers and the closing price.

Table 11. *Number of snipers and closing price*

	Closing Price	
First time	1.523	(4.529)
Time of first bid	3.711	(7.271)
Starting price	1.017***	(0.0499)
Bidders	10.72***	(0.819)
Duration	-4.975***	(1.444)
Rating	-0.000615	(0.000825)
Number of snipers	-6.407*	(3.054)
Number of squatters	8.193***	(0.840)
Harry Potter	-26.36	(33.66)
Canadian coin	-31.52	(32.66)
Indian Head	-9.188	(31.60)
iPhone 6	16.62	(30.88)
Constant	14.80	(33.44)
Observations	835	
Adjusted R^2	0.755	

Standard errors in parentheses
The dependent variable is the closing price of the auction.
Note: Excludes observations with hidden ratings
Note: Conditional on auction having at least one sniper
$^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Of particular importance is the large and significant negative coefficient on the number of squatters in the auction. This result is initially surprising. One would guess that any increase in the number of bidders would increase the expected closing price. Indeed, the variable *Bidders* has a large, positive, and significant coefficient. One explanation for this observation is that snipers do in fact seek auctions where they perceive little competition. As we have seen, a lower number of non-snipers corresponds to a relatively higher number of snipers. Thus we expect auctions with little apparent competition to attract more snipers. One can imagine multiple potential snipers watching an item. At the last moment, more than one of the potential snipers decides to enter. Auctions that attract a particularly high number of snipers may have such little non-sniping activity that the price when the flurry of sniping begins is low enough to outweigh the effect of the increased competition from sniping at the end.

6. Conclusion

This paper examines the prevalence and effectiveness of sniping and squatting. Like most previous studies, I find clear evidence of sniping. Evidence of squatting is less clear, and it is possible that bidders who appear to be squatters are simply frequent eBay bidders who see items soon after they are posted and bid as soon as they see an item. The prevalence of sniping and potential prevalence of squatting offer preliminary evidence that bidders find these strategies useful, and I test different theories as to why this might be.

First, I test the hypothesis that sniping is beneficial in common values auctions and find mixed evidence. Although there are some cases, like that iPhones, where there is intuitive and statistical evidence for common values as well as a particularly high rate of sniping, many other items have less clear relationships. Another common feature of past studies—that more experienced bidders are more likely to be snipers—is supported by my findings in only a few cases. It is therefore difficult to make a general claim about whether snipers are aware of the strategic considerations of their actions or whether they snipe because they misunderstand the nuances of bidding on eBay.

A new predictor of sniping that other literature has not examined is the number of bidders in an auction before the last few minutes of the auction begin. I find that auctions with many non-snipers tend to have fewer snipers. This can be compared to Ely and Hossain's competition effect, which holds that early bids may deter other bidders by incentivizing them to look for auctions for a similar good with fewer bidders. My results are consistent with a more limited version of this case: while it is unclear whether the threat of competition applies to all bidders, I find that potential snipers may be deterred from auctions where there are many non-sniping bids. While snipers may have different aims than ordinary bidders, the appearance of a competition effect for snipers is consistent with Ely and Hossain's broader theory that bidders seek auctions with the least competition.

While a clear pattern to the prevalence of sniping is decent evidence that snipers act strategically, stronger evidence for strategic sniping is the success of the strategy in terms of the probability of winning and the closing price. It is possible that sniping appears to be successful because potential snipers simply look for auctions that they are likely to win at a good price. It is also possible, however, that snipers have more complex motives. They may be calculating bidders who wish to hide information from their opponents. While I do not find conclusive evidence that common values auctions are significantly more likely to have sniping, this is not the only case where bidders may wish to hide information. For example, a bidder may snipe if she believes one of her competitors has not bid his valuation. In this case, sniping may prevent the competitor from responding with a higher bid, and the sniper may win at a lower price than she otherwise would have. This is consistent with most theories of sniping, which have in common the principle of hiding the information contained in a bid in order to prevent an opponent for raising his bid in response.

On the other hand, tests for the benefit of squatting are inconclusive at best. Although the presence of what appear to be squatting bids could suggest that bidders find the strategy useful, it is possible that activity at the beginning of an auction is a result of very regular bidders joining an auction without being particularly invested in winning, as Ariely and Simonson (2003) note that bidders are wont to do. Moreover, the lack of a significant relationship between early bids and the mechanism by which they should be effective—deterring other bidders from entering—casts doubt on the usefulness of squatting. Finally, squatting bids are significantly less likely to win auctions than other bids are, and winning squatters do not appear to pay significantly lower prices than other winners do.

While my results largely suggest that sniping is not the work of irrational bidders but rather is an intentional—and indeed effective—strategy, more research is needed to understand the mechanism by which it is successful. In particular, it would be interesting to distinguish between the effect of snipers choosing easy-to-win auctions and the effect of sniping even in very competitive auctions. My results suggest that sniping is beneficial in both cases, but knowing the relative sizes of these effects may be enlightening. If the necessary data becomes available, future studies could also examine how many snipes are placed in person and how many are placed by automated websites. Researchers could attempt to distinguish between the success of human and robot snipers. Also, with appropriate data, researchers could compare an individual user's bids across auctions to test whether apparent squatters are in fact frequent bidders who happen to bid early.

References

Anwar, Sajid, Robert McMillan, and Mingli Zheng. 2006. "Bidding Behavior in Competing Auctions: Evidence from eBay." *European Economic Review* 50 (2): 307-22.

Ariely, Dan, Axel Ockenfels, and Alvin E. Roth. 2003. "An Experimental Analysis of Ending Rules in Internet Auctions." CESifo Working Paper, No. 987.

Ariely, Dan, and Itamar Simonson. 2003. "Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions." *Journal of Consumer Psychology* 13 (1-2): 113-23.

Athey, Susan, and Philip A. Haile. 2002. "Identification of Standard Auction Models." *Econometrica* 70 (6): 2107-40.

Bajari, Patrick, and Ali Hortaçsu. 2004. "Economic Insights from Internet Auctions." *Journal of Economic Literature* 42 (2): 457-86.

———. 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical insights from eBay Auctions." *RAND Journal of Economics*: 329-55.

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2015. "Assessing Sale Strategies in Online Markets Using Matched Listings." *American Economic Journal: Microeconomics* 7 (2): 215-47.

Ely, Jeffrey C., and Tanjim Hossain. 2009. "Sniping and Squatting in Auction Markets." *American Economic Journal: Microeconomics* 1 (2): 68-94.

Goeree, Jacob K., and Theo Offerman. 2002. "Efficiency in Auctions with Private and Common Values: An Experimental Study." *The American Economic Review* 92 (3): 625-43.

Gray, Sean, and David Reiley. 2007. "Measuring the Benefits to Sniping on eBay: Evidence from a Field Experiment." *Unpublished Manuscript*.

Haile, Philip A., Han Hong, and Matthew Shum. 2003. "Nonparametric Tests for Common Values at First-Price Sealed-Bid Auctions." NBER Working Paper 10105.

Hossain, Tanjim. 2008. "Learning by Bidding." *The Rand Journal of Economics* 39 (2): 509-29.

Ku, Gillian, Deepak Malhotra, and J. Keith Murnighan. 2005. "Towards a Competitive Arousal Model of Decision-Making: A Study of Auction Fever in Live and Internet Auctions." *Organizational Behavior and Human Decision Processes* 96 (2): 89-103.

Milgrom, Paul R., and Robert J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica: Journal of the Econometric Society*: 1089-122.

Ockenfels, Axel, and Alvin E. Roth. 2006. "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction." *Games and Economic Behavior* 55: 297-320.

Paarsch, Harry J. 1992. "Deciding between the Common and Private Value Paradigms in Empirical Models of Auctions." *Journal of Econometrics* 51 (1-2): 191-215.

Pinske, Joris, and Guofu Tan. 2005. "The Affiliation Effect in First-Price Auctions." *Econometrica* 73 (1): 263-277.

Rasmusen, Eric. 2006. "Strategic Implications of Uncertainty over One's Own Private Value in Auctions." *Advances in Theoretical Economics* 6 (1): 7.

Roth, Alvin E., and Axel Ockenfels. 2000. "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet." NBER Working Paper 7729.

Titcomb, James. "eBay Turns 20: Here Are 20 Things You May Not Know." *The Daily Telegraph*. September 3, 2015.

Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance* 16 (1): 8-37.

Wolak, Frank A. 1989. "Testing Inequality Constraints in Linear Econometric Models." *Journal of Econometrics* 41 (2): 205-35.