Learning by Bidding*

Tanjim Hossain†
Hong Kong University of Science and Technology

September 2004

Abstract

We analyze a dynamic second-price auction similar to online auctions where the price in any period equals the current second-highest bid. In this auction, one bidder, known as the informed agent, knows her private valuation. The other bidder, known as the uninformed agent, does not know his private valuation. However, upon seeing a posted price, he learns whether his valuation is above that price. We show that, in any equilibrium, the informed agent bids in the first period if her valuation is below some cutoff and bids only in the last period otherwise. The uninformed agent bids in every period to optimally change the price unless the price is above his valuation or he is the high bidder. Thus, this model suggests that late bidding or *sniping* and multiple bidding or *nibbling* are interrelated and provides an explanation for their prevalence in online auctions. It also provides a rationale behind frequent use of secret reserve price auctions by showing that those may increase the seller's expected revenue. Moreover, this model shows that a secret reserve price reduces sniping.

JEL Classifications: D44, L86, D82, C73

Keywords: Internet auctions, dynamic auctions, asymmetric information, anchoring effect, contemplation cost

*I would like to thank Faruk Gul, Han Hong, Daniel Kahneman, Fahad Khalil, Jeffrey Livingston, Alexandre Mas, Eric Maskin, Shantanu Mukherjee, Elie Tamer, Thijs van Rens and seminar participants at Boston University, Cornell, HKUST, UPenn, Princeton, NYU and University of Washington for their helpful comments and suggestions and Matthew Langley for help in creating the data set. I am indebted to John Morgan and, especially, Wolfgang Pesendorfer for their encouragement, guidance and support all through the development of this paper. Financial support of the Woodrow Wilson Foundation is gratefully acknowledged. All remaining errors are mine.

†Department of Economics, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. E-mail: tanjim@ust.hk

1 Introduction

The trade volume of Internet auction sites has grown dramatically over the last few years. Established in 1995, *eBay*, "The World's Online Marketplace," now runs more than 16 million auctions on any given day. In 2002, the value of goods traded on eBay was about $14.87 billion. Around 48 million users actively participated on eBay websites worldwide during the financial year 2003-04.[1] It has 85% share of the US online auction market. In 2002, 48.2% of adults in the US made online purchases and eBay had the highest number of unique online purchasers. Manufacturers such as Xerox and IBM and retailers such as J.C. Penney and Gymboree have set up shops on eBay (Walker, 2001).

Auctions on eBay are dynamic second-price auctions. An eBay auction starts at an opening price and the price changes as bidders place bids. A bidder can submit a "maximum bid" which is kept undisclosed. At any point during the auction, the current price equals the second highest bid received so far.[2] The auction ends at a pre-announced time. At that time, the bidder who had placed the highest bid wins the object and pays the end price equaling the maximum bid of the second highest bidder. The winner, hence, does not pay the maximum bid she submitted.

It is an empirical regularity on eBay that many auctions receive *sniping* bids, *i.e.*, bids placed in the last few minutes of an auction. Many bidders frequently update their bids and bid slightly above the current price if they are not the highest bidder. We refer to this multiple bidding as *nibbling*.[3] Hossain and Morgan (2003) sold brand new popular music CDs and Xbox game cartridges on eBay. There was little uncertainty about the quality or popularity of these products and, therefore, the assumption of private values seems appropriate. In 76% of the auctions, at least one bidder had placed multiple bids. Bids were placed in the last five minutes in 30% of the auctions. Some bidders bid early in the auction, some bid late and others bid both early and late.

Many other studies found the same bid pattern in eBay auctions for a wide variety of goods. In a study of coin auctions by Bajari and Hortacsu (2003), 32% of the bids were submitted after 97% of the duration of the auction had passed. Roth and Ockenfels

[1] See eBay company overview on http://ebay.com/community/aboutebay/ for more details.

[2] The price, in fact, equals the second highest bid plus a small bid increment that we ignore. For a more detailed description of eBay auctions, see Roth and Ockenfels (2001).

[3] I thank John Morgan for suggesting this term.

(2001) gathered a sample of over 1000 auctions of various items. Of the auctions with two or more bidders, 18% received bids in the last minute and 74% had at least one bidder submitting multiple bids. A bidder submitted two bids on average. Hasker, Gonzalez and Sickles (2001) collected data from 7932 auctions of color computer monitors. Around 77% of them included bidders who placed multiple bids and over 11% of bids were received in the last minute.

To take a closer look at the effect of multiple bidding on auction outcomes, we collected a data set of 2026 completed auctions of "Titleist 975J" golf drivers conducted on eBay between February and April of 2003. Of these auctions, 1289 resulted in sale. Average revenue from these auctions was $164.39 with a standard deviation of $31.96. Table 1 presents summary statistics for some variables characterizing the auctions. These auctions experienced a significant amount of sniping and nibbling. On average, a bidder placed 1.89 bids. About 40% of all bidders placed more than one bid or *nibbled* (these bidders are referred to as "multiple-bidders" henceforth). We define sniping as bidding only in the last three minutes of an auction and 6.5% of all bidders *sniped*.

Did bid patterns have any impact on something of high economic significance; e.g., seller's revenue? Auctions with a higher ratio of multiple-bidders received higher revenue. Table 2 shows that, controlling for everything else, an auction where all bidders placed only one bid each received almost 34% less revenue than an auction with only multiple-bidders.[4] This result shows that multiple bidding phenomenon has significant effect on revenue. Moreover, this increase in revenue seems to be due to multiple-bidders. Multiple-bidders started with a lower bid than bidders who placed only one bid (referred to as "single-bidders" henceforth), but finally ended up bidding higher in the auction. In the data set, first bids by multiple-bidders were lower than those by single-bidders by $7 and last bids by multiple-bidders were higher than those by single-bidders by $20.[5] The average final bid of bidders was below $118. Therefore, there was a significant difference between single-bidders and multiple-bidders in terms of their final bids and their effect on revenue. Hence, explaining multiple bidding in eBay auctions is important.

In a standard private-value setting, it is a weakly dominant strategy for a bidder to bid

[4]Please see Hossain (2004) for a more detailed description of the empirical analysis.

[5]These results are presented in Table 3 and are discussed in more details in Hossain (2004). The results do not change when we account for the same bidder bidding in multiple auctions.

her valuation. Hence, any strategy profile where all bidders ultimately bid their private valuations is an equilibrium of an eBay-type auction. All these equilibria lead to the same outcome as the equilibrium where all bidders bid their valuations at the beginning of the auction. Thus, sniping and nibbling can be consistent with the standard theory. However, it will not explain the increase in revenue in auctions with a higher share of multiple-bidders and higher final bids by multiple-bidders without further assumptions. Moreover, equilibria where some snipe and some nibble cease to exist when some perturbation to the auction design is introduced. Suppose in every period of a discrete time auction, there is a small probability that a bidder will not be able to place any bid in the remainder of the auction. This probability creates a small opportunity cost of not bidding one's valuation immediately. Then, in the unique equilibrium in weakly undominated strategies, a bidder bids her valuation in the first period.

Some have used common-value objects to explain sniping. Bajari and Hortacsu (2003) introduced a two-stage auction where an open-exit continuous time ascending auction is followed by a discrete last period. If the object is common-value, then there is a unique symmetric Nash equilibrium where all biders bid only in the last period. Hence, there will be sniping but there will not be any nibbling or any early bidding.

Roth and Ockenfels (2001) have a model similar to that of Bajari and Hortacsu. However, objects are private-value and a last period bid does not get accepted with a positive probability. For some parameter values, there exist multiple equilibria with an equilibrium in which all bidders bid only in the last period. In other equilibria, bidders go to "war;" that is, they bid their valuations before the last period. Although sniping can be an equilibrium phenomenon, both early and late bids will not arise in the same equilibrium. In the equilibrium where all bidders bid early, there should not be any systematic difference between the last bids of a multiple-bidder and of a single-bidder. Moreover, empirical analysis by Hasker, Gonzalez and Sickles (2001) rejected the hypothesis that bidders follow "snipe-or-war" strategies suggested by Roth and Ockenfels.

In Rasmusen (2003), one of the two bidders is informed and knows her private-valuation. The other bidder, the uninformed one, does not know her private valuation and can learn it exactly any time by paying a cost. That model can explain multiple bidding, but fails to explain sniping unless the uninformed bidder is naive and the informed bidder is sophisticated. Moreover, multiple-bidders will not end up bidding higher than single-bidders. Wang (2003) has a model with two consecutive auctions where all bidders other than the winner of the first auction participate in the second auction. Thus, there will be one fewer bidder in the second auction and the two auctions are ex-ante asymmetric. In the first auction, all bidders will snipe so that other bidders do not get information about their valuations while the first auction is still going on. In that auction, there will be sniping but no earlier bidding. Sniping in the first auction depends on the assumption that the exact same bidders participate in both auctions and no new bidders enter. Moreover, the second auction is the "last" auction, so information revealed in that auction does not hurt bidders unlike the first auction. These assumptions are inconsistent with eBay auctions where consecutive auctions are ex-ante symmetric because there are future auctions for both and buyers enter and exit the market frequently.

To provide a plausible explanation of the empirical evidence in an independent private-value setting, we ask the question "[H]ow do people learn how much they value a good?" A rational agent is usually assumed to know her valuation for a private-value object. We suggest that some agents are non-standard in the sense that they may not be aware of how much they like an object. The cognitive cost of calculating their exact valuation is prohibitively high for these agents. In many real life situations, people do not need to know their exact valuations. For example, when a person buys milk in the supermarket, she only needs to know whether the price is low enough. That is, knowing her preferences around the posted price is enough for decision-making in this situation. There is also evidence that people do not understand how much they like an object before they have to pay for it. If there is a posted price then the person knows that she has to pay at least that amount to get the good. A posted price thus facilitates the need to know whether a person's valuation is above or below the price.

Similar to Ergin (2003), we assume that individuals learn their preferences by contemplating about their tastes. To model this learning, we assume agents ask themselves if they would be willing to pay some hypothetical amount for the good. If one can contemplate about any hypothetical price without any cost, she can calculate her valuation like a standard rational agent. However, when contemplation or self-introspection is costly, calculating the exact valuation is infinitely costly if the agent has infinitely many possible valuations. We assume that for some agents, contemplation at any price is costless and they know their valuations. For other agents, the cost of contemplating at a posted price is costless and contemplating at any other price is infinitely costly. For such an agent, knowing the minimum amount she has to pay to get that object makes it easier to think about the relation between her valuation and the amount. She can immediately compare her taste for the good and a posted price and learn whether her valuation is above that.

The idea that a posted price can trigger a revision of a consumer's willingness to pay is somewhat related to the "anchoring effect." Ariely, Loewenstein, and Prelec (2003) showed products such as bottles of wine and computer keyboards to subjects and asked whether they would buy these products for a dollar figure equaling the last two digits of their social security numbers. The subjects then stated their maximum willingness to pay for these products. The impact of the social security number on willingness to pay was significant in every product category. Subjects with above-median social security numbers stated values from 57% to 107% greater than did subjects with below-median social-security numbers. Many other studies suggest that people's estimate of valuations can be shaped by an asked price. In this paper, agents are Bayesian learners and costless learning at the posted price does not perfectly replicate the results from "anchoring effect" experiments. Nevertheless, with either, an agent's expected valuation for a good can be significantly affected by the price even if it is not correlated with the quality of the product. A high anchor is needed to get a high estimation from an agent in both cases.

This paper analyzes the effect of learning from a posted price in a dynamic second-price auction with a structure similar to those of Bajari and Hortacsu (2003) and Roth and Ockenfels (2001). However, our auction is in discrete time with a fixed number of periods T. A discrete time auction takes into account that bidders may need a "reaction time" to respond to an action by another bidder. Thus, it better models the decision problem of choosing an optimal maximum bid given the time remaining in the auction. It is similar to Vincent's dynamic auction (1990) except that the seller cannot accept any offer before period T and the winner pays the second highest bid instead of her own bid.

In this auction, there are two bidders. A bidder can submit a bid in every period. Her highest bid is considered her current bid and earlier bids are discarded. The auction starts at an opening price and in any future period, the current price equals the second highest bid. Hence, the price does not change if only one bidder places bids. At the end of period T, the highest bidder wins the object and pays the second highest bid. Bidders have independently drawn private valuations for the object. In every period, there is a small probability ϵ that a bidder will exogenously drop out from the auction. If she drops out, she cannot place any more bids in the auction and her current bid will be her last bid. In real life, this may correspond to a bidder being unable to go back to an auction because of unanticipated personal commitments, computer malfunctions or forgetfulness.

When both bidders know their valuations, both bidding their valuations in the first period is the unique equilibrium in weakly undominated strategies. Next we consider the case where one bidder, referred to as the informed bidder, knows her valuation exactly. The other bidder, referred to as the uninformed bidder, learns about his valuation by comparing it with the current price. At the beginning of every period, he receives a positive or negative signal denoting whether his valuation is above the current price. In the rest of the paper, we use female and male pronouns for the informed bidder and the uninformed bidder respectively.

A simple example: Suppose there are 2 periods, the opening price is zero and bidders' valuations are uniformly distributed on $[0, 1]$. Bidder 1 knows her valuation v_1 and bidder 2 knows whether his valuation v_2 is above the current price p_t in each period. In weakly undominated strategies, bidder 1's final bid is v_1 and bidder 2 bids the conditional expected value of v_2 in the last period. That is, if v_2 is above price p_2, then he bids $\frac{1+p_2}{2}$.

Let us compare two strategies of bidder 1:

Strategy 1: bidding v_1 in period 1.

Strategy 2: bidding zero in period 1 and v_1 in period 2.

For any v_1, no other strategy is strictly preferred to both strategies. If bidder 1 bids in period 1, bidder 2 gets new information about v_2 in period 2. If she does not bid in period 1, bidder 2 does not get any new information in period 2 as the price stays at zero. He bids $\frac{1}{2}$ in period 2 in that case. Suppose bidder 2's bid in period 1 is $x_1 \leq \frac{1}{2}$.

If $v_1 \leq x_1$ then bidder 1 gets zero payoff from both strategies. If $v_1 \in \left(x_1, \frac{1}{2}\right]$ then p_2 equals x_1 if bider 1 follows strategy 1. With probability x_1, Bidder 2 gets a negative signal in period 2 and does not place any more bid. If bidder 2 gets a positive signal then he bids $\frac{1+x_1}{2}$ and bidder 1 loses the auction. Hence, bidder 1 gets an expected payoff of

$x_1(v_1 - x_1)$ from strategy 1. She gets an expected payoff of zero from strategy 2 as $v_1 \leq \frac{1}{2}$. If $v_1 \in \left(\frac{1}{2}, \frac{1+x_1}{2}\right]$ then bidder 1 gets an expected payoff of $x_1(v_1 - x_1)$ from strategy 1. She gets an expected payoff of $v_1 - \frac{1}{2}$ from strategy 2. There is a value $v \in \left(\frac{1}{2}, \frac{1+x_1}{2}\right)$ such that bidder 1 gets the same payoff from the two strategies. If $v_1 \geq \frac{1+x_1}{2}$ then bidder 1 wins if she follows strategy 1 even when bidder 2 gets a positive signal in period 2. Bidder 1's expected payoff is $v_1 - \frac{1}{2} - \frac{x_1{}^2}{2}$. Her expected payoff from strategy 2 is $v_1 - \frac{1}{2}$. Hence, she prefers strategy 2 to strategy 1 if $v_1 \geq \frac{1+x_1}{2}$.

To learn about v_2, bidder 2 chooses his period 1 bid optimally. Bidder 2's best response to bidder 1's strategy is choosing $x_1 = \frac{v}{2}$. In the equilibrium,

$$x_1 = \frac{1}{2+\sqrt{2}}, \ v = \frac{2}{2+\sqrt{2}}.$$

Therefore, bidder 1 bids v_1 in period 1 if $v_1 \leq \frac{2}{2+\sqrt{2}}$ and bids zero in period 1 and v_1 in period 2 if $v_1 > \frac{2}{2+\sqrt{2}}$. Bidder 2 bids $\frac{1}{2+\sqrt{2}}$ in period 1 and conditional expected value of v_2 in period 2 if he gets a positive signal.

Bidder 1 bids either in period 1 or in period 2. In contrast, bidder 2 bids in both periods if she gets a positive signal in period 2. Thus, bidder 2 nibbles to learn about v_2 and this nibbling leads to sniping by bidder 1. For an arbitrary T and a general distribution F from which v_1 and v_2 are drawn, there exists an analogous equilibrium.

Theorem 2, in section 3, shows that, for small perturbations, there exists an equilibrium (σ^v, σ^x) with the following properties:

i) Bidder 1 bids above the opening price either only in the first period or only in the last period. There exists a cutoff value $v < 1$ such that if $v_1 \leq v$, then she bids v_1 in period 1. If $v_1 > v$, then she bids v_1 in period T.

ii) Bidder 2 bids in every period unless he is the high bidder or learns that $v_2 < p_t$. His bid in period t is denoted by x_t and the sequence of x_t is denoted by x.

Results from theorem 2 illustrate that sniping and nibbling are equilibrium phenomena in a private-value setting. This suggests that a significant fraction of bids placed in the dying moments of an eBay auction will be placed by bidders who bid only towards the end of an auction. In our data set, 67% of the bidders who placed a bid in the last three minutes of an auction did not place any bid earlier in that auction.

Theorem 3 shows that this auction has a unique class of equilibria characterized by v and x. Suppose a sequence of equilibria of the ϵ-perturbed auction converges to (σ_1, σ_2)

as ϵ goes to zero. When (σ_1, σ_2) is played, the informed bidder bids in period 1 if her valuation is below v and bids only in period T otherwise. The uninformed bidder chooses his bids when he is not the high bidder using the sequence x. Moreover, the winner and transaction price pair of (σ_1, σ_2) is the same as that of (σ^v, σ^x) with probability 1.

Learning by bidding may also explain the extensive use of "secret reserve price auctions." In a standard private-value setting, a secret or a public reserve price lead to the same outcome. If all bidders are informed, then paying the fee for keeping the reserve price secret is suboptimal for a seller. However, an uninformed bidder wins the object when his valuation is below the reserve price with a positive probability when the reserve price is secret. Proposition 6 shows that a secret reserve price auction may generate higher expected revenue than a public reserve price auction.

The rest of the paper is organized as follows: The next section introduces the benchmark model where both bidders are informed. Section 3 analyzes the auction with an informed bidder and an uninformed bidder. Section 4 discusses the auction with two uninformed bidders. Section 5 analyzes secret reserve price auctions and section 6 concludes the paper. All proofs are in the appendix.

2 The Benchmark Model

The auction is a dynamic version of conventional sealed-bid second-price auctions and is similar to eBay auctions for a single object. The seller, indexed by 0, auctions off an indivisible object to two bidders. There are T periods indexed by $t \in \{1, 2, \ldots, T\}$. Nature independently draws each bidder's private valuation from the distribution F which is strictly increasing and twice differentiable on $[0, 1]$. Bidder $i \in \{1, 2\}$ has a private valuation v_i for the object and is risk-neutral. If she wins the object and pays a price p, her payoff is $v_i - p$. If she does not win, she pays zero and her payoff is zero.

A bidder can place a bid in each period. Bidder i's action in period t is denoted by $b_t^i \in [0, 1]$. We define β_t^i as $\max_{\tau \leq t} b_\tau^i$. In period t, the second highest bid up to period $t-1$ is posted as the current price p_t. The auction starts at an opening price of $m \in [0, 1)$. Hence, $p_1 = m$ and $p_t = \max\left[m, \min\left[\beta_{t-1}^1, \beta_{t-1}^2\right]\right]$ for $t \geq 2$. After period T, if both β_T^1 and β_T^2 are below m then the object stays unsold. Otherwise, if $\beta_T^1 \geq \beta_T^2$ then bidder 1 wins and if $\beta_T^2 > \beta_T^1$ then bidder 2 wins. The winner pays the maximum of m and the second highest bid. Therefore, the transaction price is max $\left[m, \min\left[\beta_T^1, \beta_T^2\right]\right]$.

The high bidder in period t is denoted by $w_t \in \{0,1,2\}$. If $t = 1$ or $\beta_{t-1}^i < m$ for all i then $w_t = 0$. If $\beta_{t-1}^i \geq m$ for some i then $w_t = 2$ if $\beta_{t-1}^1 < \beta_{t-1}^2$ and $w_t = 1$ if $\beta_{t-1}^1 \geq \beta_{t-1}^2$. The set of bidders whose period t bid is above p_t and β_{t-1}^i is denoted by $A_t \subseteq \{1,2\}$. Hence, $A_0 = \emptyset$ and $A_t = \left\{i \in \{1,2\} \,|b_t^i > \max\left[p_t, \beta_{t-1}^i\right]\right\}$ for $t \geq 2$. In period t, bidder i observes p_t, w_t and A_{t-1} but does not observe b_τ^j for any $\tau < t$ where $j \neq i$. The public history at the beginning of period t, h_t is the sequence $\{(p_\tau, w_\tau, A_{\tau-1})\}_{\tau=1}^t$. A terminal history is denoted by h_{T+1} where w_{T+1} is the winner and p_{T+1} is the final price. The set of all possible h_t is denoted by H_t. A strategy σ is a sequence $\{\sigma_t\}_{t=1}^T$ where $\sigma_t : H_t \times [0,1] \to [0,1]$.

Now we introduce a perturbation to the auction design. At the beginning of period $t \in \{2, \ldots, T\}$, there is a small positive probability ϵ that bidder i exogenously "drops out" off the auction. Then she will not be able to place any more bids in that auction. Formally, if bidder i drops out in period t, then $\beta_T^i = \beta_{t-1}^i$. A bidder has an ex-ante probability of $(1-\epsilon)^{t-1}$ of being able to place a bid in period t. We are interested in equilibrium analysis when ϵ approaches zero. In the remainder of the paper, bidder i's period t action b_t^i refers to the bid she places in period t provided she has not dropped out yet in period t. The fact that bidder i has dropped out is private information.

We restrict our attention to weakly undominated strategies in this paper. An *equilibrium* of this auction is defined as a Perfect Bayesian Equilibrium where none of the bidders play a weakly dominated strategy. Lemma 1 shows that a bidder bids her valuation in some period in any weakly undominated strategy.

Lemma 1 *If bidder i follows a weakly undominated strategy, then for any terminal history h_{T+1}, $b_t^i \leq v_i$ for all $t \leq T$ and there exists a $\tau \leq T$ such that $b_\tau^i = v_i$.*

Theorem 1 shows that the only equilibrium of this auction has all bidders bidding their valuations in the first period.

Theorem 1 *In any equilibrium, $b_1^i = v_i$ if $v_i \geq m$.*

In any undominated strategy, a bidder bids her valuation at some period. Since ϵ is positive, bidding one's valuation in the first period is the only equilibrium of this discrete-time game. This implies that late and multiple bidding cannot exist in an equilibrium in the benchmark model where all bidders know their private valuations.

Bidders participating only in one auction is a simplifying assumption. On eBay, many concurrent auctions of identical or very similar goods go on and new auctions of similar products start every day. Therefore, bidders can participate in many auctions that are ex-ante identical. Moreover, the good can usually be purchased in the retail market. When we account for these outside opportunities, the optimal bid of a bidder can be below her private valuation. Still, it will be a function of her private valuation and her optimal strategies will be type-dependent. Given many concurrent auctions, a bidder may place more than one bid in an auction as she gets new information about her outside opportunities as the auction progresses. Nevertheless, one would not see a bidder placing 5 or 6 bids within a minute on the same auction without placing bids in any other auctions, a pattern common in eBay auctions. More importantly, concurrent auctions will make sniping less profitable than bidding early. Since there are many concurrent auctions and bidders do not participate in all auctions, coordination among bidders on which auction to participate in is an issue. Bidders prefer an auction with fewer bidders (and a lower price) and bidding early reduces competition. Sniping, on the other hand, will increase competition and will be sub-optimal. Thus, a retail market and concurrent and future auctions do not explain the pervasiveness of sniping on eBay.

3 Auctions with an Uninformed Bidder

In this section, bidder 1 is a standard agent who knows her valuation v_1. Bidder 2 learns about his valuation only by costlessly comparing it with the current price. More precisely, he learns whether v_2 is above p_t in every period t. For clarity, we will use female and male pronouns for bidders 1 and 2 respectively. We refer to bidders 1 and 2 as the informed bidder and the uninformed bidder respectively.

Before the auction starts, bidder 1 learns v_1 and bidder 2 learns whether v_2 is above m. At the beginning of period $t \in \{2, \ldots, T\}$, he receives a private signal that tells him whether v_2 is as large as p_t. The signals are always accurate and are either positive or negative. A positive signal implies that $v_2 \geq p_t$. Since p_t is non-decreasing in t, once the uninformed bidder receives a negative signal, all his future signals will be negative.

Receiving a signal, bidder 2 updates his prior according to the Bayes' rule.

In any weakly undominated strategy, bidder 2's bid in the last period equals the conditional expected value of v_2 if v_2 is above p_T. For $a, c \in [0,1]$ with $a < c$, let

$$\mathbf{E}[Q|[a,c]] = \frac{\int_a^c Q(y)\, dF(y)}{F(c) - F(a)}.$$

Lemma 2 *If $v_2 \geq p_T$, then $b_T^2 = \mathbf{E}[y|[p_T, 1]]$ in any weakly undominated strategy.*

There is an equilibrium where bidder 1 bids v_1 when she places a bid above the minimum price. She bids v_1 in the first period if v_1 is below a cutoff value v and bids v_1 in the last period otherwise. In this equilibrium, bidder 2 places bids in every period until he becomes the high bidder or gets a negative signal. Hence, bidder 1 snipes when v_1 is high enough and bidder 2 nibbles when bidder 1 bids early. Bidder 2 uses his bids as potentially costly experiments to learn v_2. The cost of bidding is that he may become the high bidder and pay a price higher than v_2. Bidder 1 uses her bids to decide how much experimentation to allow. If she bids in period 1, bidder 2 can learn about v_2 by bidding. If she bids only in the last period, bidder 2 gets no information about v_2.

Let us define bidder 1's strategy σ^v to be such that

$$\sigma_t^v(h_t, v_1) = \begin{cases} v_1 \text{ for all } t & \text{if } v_1 \leq v \\ m \text{ for } t < T \text{ and } v_1 \text{ for } t = T & \text{if } v_1 > v. \end{cases}$$

If bidder 2 gets a negative signal in period t, he does not bid any more as the price is higher than v_2. By lemma 2, he bids the conditional expected value of v_2 in the last period. In period $t < T$, if he gets a positive signal and is not the high bidder, he places a bid that maximizes his expected payoff given history h_t. The sequence of these bids is of most interest to us. Let us define $x = \{x_t\}_{t=1}^{T-1}$ to be a sequence of $T-1$ scalars where bidder 2's bid in period t is $\max[x_t, \underline{v}(h_t)]$ if $v_2 \geq p_t$, $t < T$ and $w_t \neq 2$. Here $\underline{v}(h_t)$ is bidder 2's belief of the lowest possible value of v_1 given h_t. If bidder 2's belief system is μ_2, then $\underline{v}(h_t) = \inf(\text{supp}\mu_2(h_t))$. Bidder 2's strategy σ^x is such that

$$\sigma_t^x(h_t, v_2) = \begin{cases} 0 & \text{if } v_2 < p_t \text{ or } (t < T \text{ and } w_t = 2) \\ \max[x_t, \underline{v}(h_t)] & \text{if } v_2 \geq p_t,\ t < T \text{ and } w_t \neq 2 \\ \mathbf{E}[y|[p_T, 1]] & \text{if } v_2 \geq p_T \text{ and } t = T. \end{cases}$$

Strategies σ^v and σ^x are weakly undominated.

If (σ^v, σ^x) is an equilibrium strategy profile then we refer to (v, x) as the equilibrium cutoff-action pair. There is a unique equilibrium cutoff-action pair of this auction when ϵ is small enough. Equilibrium cutoff-action pairs converge to the equilibrium cutoff-action pair of the unperturbed auction ($\epsilon = 0$) as ϵ approaches zero.

Bidder 1's expected payment is $P(t, v_1)$ and her probability of winning if she bids only in period t is $W(t, v_1)$. Using $x_0 = m$ and $x_T = v$ for notational convenience,

$$\begin{aligned} P(1, v_1) &= \sum_{t=0}^{T-1} x_t \left(F(x_t) - F(x_{t-1})\right) \\ P(T, v_1) &= mF(m) + \int_m^1 y dF(y). \end{aligned}$$

Further, $W(1, v) = F(x_{T-1})$ and $W(T, v) = 1$. To characterize the equilibrium cutoff-action pair of the unperturbed auction, we use the following two equations:

$$x_t = \mathbf{E}\left[y | \left[x_{t-1}, x_{t+1}\right]\right] \text{ for } t \in \{1, \ldots, T-1\} \tag{1}$$

$$v = \frac{P(T, v) - P(1, v)}{W(T, v) - W(1, v)}. \tag{2}$$

Theorem 2 *(i) There is an $\bar{\epsilon} > 0$ such that for all $\epsilon \in [0, \bar{\epsilon})$ there exists a unique equilibrium cutoff-action pair (v^ϵ, x^ϵ). Moreover, $v^\epsilon < 1$ and $x_t^\epsilon - x_{t-1}^\epsilon > 0$ for all t.*
(ii) The equilibrium cutoff-action pair when $\epsilon = 0$ is (v, x) if and only if $\lim_{\epsilon \to 0} (v^\epsilon, x^\epsilon) = (v, x)$ and (v, x) satisfies equations 1 and 2.

In the static or 1-period game, bidder 1 bids v_1 and bidder 2 bids $\mathbf{E}[y | [m, 1]]$ if $v_2 \geq m$ in the unique equilibrium. In the dynamic auction, bidder 2 gets multiple opportunities to change the price with his bids and learn about v_2. When choosing a bid, he faces the trade-off between winning at a price higher than v_2 by bidding too high and not learning much about v_2 by bidding too low. Bidder 2's equilibrium bids reflect optimal experimentation when he is not the high bidder. Bidder 1's equilibrium behavior is a strategic response to this learning. Since the current price stays unchanged if one of the bidders does not bid and the auction ends after a fixed number of periods, bidder 2 does not get any opportunity to learn about v_2 if bidder 1 does not bid until the last period. Bidder 1's optimal strategy determines how much learning by bidder 2 she should allow and bidder 2's optimal strategy determines exactly how much to learn in each period. In the equilibrium discussed in theorem 2, bidder 1 either lets bidder 2 experiment as much as possible by bidding in the first period or she does not let bidder 2 experiment at all by bidding only in the last period.

If bidder 1 could induce perfect learning of v_2 by bidder 2, her expected payoff, when $\epsilon = 0$, would be

$$(v_1 - m) F(m) + \int_m^{v_1} (v_1 - y) \, dF(y) . \tag{I}$$

If she bids only in the last period, her expected payoff is

$$(v_1 - m) F(m) + \int_m^{1} (v_1 - y) \, dF(y) . \tag{II}$$

Since (I)>(II), if she could induce perfect learning by bidding in period 1, bidder 1 would bid in period 1 for any v_1. However, when bidder 2 learns only by judging whether $v_2 \geq p_t$, his learning of v_2 is not perfect for any finite T. Then bidder 1 can be better off by not bidding before the last period for some values of v_1. This will be clear in a 2-period auction example. Suppose bidder 2 bids x_1 in period 1. If $v_1 > \mathbf{E}[y | [x_1, 1]]$, bidder 1's expected payoff from bidding v_1 in period 1 is

$$(v_1 - m) F(m) + (v_1 - x_1)(F(x_1) - F(m)) + \int_{x_1}^{1} (v_1 - y) \, dF(y) . \tag{III}$$

Her expected payoff if she bids only in period 2, given by (II), is above (III) if $x_1 > m$. Hence, bidding v_1 in period 1 is better for bidder 1 if $v_1 \leq v$ where

$$v = \frac{\int_m^1 y dF(y) - x_1 (F(x_1) - F(m))}{1 - F(x_1)} .$$

If $v_1 > v$, it is optimal for her to bid only in the last period. Thus, imperfect learning by bidder 2 leads to sniping by bidder 1.

If bidder 2 is not the high bidder in period 2, then $p_2 = x_1$. If $v_2 > x_1$, then he bids $\mathbf{E}[y | [x_1, 1]] > v$. Bidder 1 wins for sure if $v_1 > v$. In period 1, bidder 2 maximizes

$$\int_m^1 \int_m^{x_1} (z - y) \, dF(y) \, dF(z) + \int_{x_1}^1 \int_{x_1}^{v} (z - y) \, dF(y) \, dF(z) .$$

Hence, he chooses x_1 to be $\mathbf{E}[y | [m, v]]$.

The equilibrium cutoff-action pair is unique for small enough ϵ. There is a unique cutoff $v^\epsilon < 1$ for which there exists an action sequence x^ϵ such that (v^ϵ, x^ϵ) leads to an equilibrium. The optimal x^ϵ given this v^ϵ is unique. As ϵ goes to zero, this pair converges to (v, x), the unique equilibrium cutoff-action pair of the unperturbed auction.

This auction, essentially, has a unique equilibrium. Any equilibrium at the limit as ϵ goes to zero belongs to a class of equilibria characterized by (v, x). Since bidder 2 makes a discrete bid jump $(x_t - x_{t-1})$ if he is not the high bidder and gets a positive signal, bidder 1 has multiple equilibrium strategies that give her the same payoff as that from following σ^v. If v_1 is in $[x_{k-1}, x_k)$ for some k and bidder 1 chooses any bid from $[x_{k-1}, v_1]$ in period 1 then, in equilibrium, she loses the auction if and only if bidder 2 bids in period k. Hence, bidder 1 gets the same expected payoff from any bid in $[x_{k-1}, v_1]$ in period 1. Once bidder 2 becomes the high bidder, he may bid below or equal to the conditional expected value of v_2 in any period. Nevertheless, in any equilibrium, bidder 2's final bid is the expected value of v_2 conditional on his signal in period T. All these strategies lead to the same partition of v_1 according to which bidder 1 decides whether or not to snipe. Bidder 2 uses the same sequence of bids x for bids in periods he is not the high bidder.

In the two period example, the following strategies characterize any equilibrium strategy by bidder 1. If $v_1 \in [m, x_1)$, she bids any value in $[m, v_1]$ in period 1 and v_1 in period 2. Similarly, if $v_1 \in [x_1, v]$, she bids from $[x_1, v_1]$ in period 1 and v_1 in period 2. If $v_1 > v$, she bids m in period 1 and v_1 in period 2. Bidder 2 bids x_1 in period 1 if $v_2 \geq m$ and conditional expected value of v_2 in period 2 in any equilibrium strategy. As a result, for all practical purposes, bidder 2's information set in any period is the same in any equilibrium and all equilibria lead to the same final outcome. Theorem 3 formalizes this result.

We need to introduce some new notations for theorem 3. Suppose $b = \{b_t\}_{t=1}^{T-1}$, where b_t is a scalar, is such that $b_t < b_{t+1}$. Then $d(a, b) = b_\tau$ if $b_\tau \leq a < b_{\tau+1}$. That is, $d(a, b)$ returns the largest number in b that is smaller than or equal to a. We define $C(v, x) = \{(\sigma_1, \sigma_2) \,|\sigma_1 \in \Omega_1^v, \sigma_2 \in \Omega_2^x\}$ where Ω_1^v is such that $\sigma \in \Omega_1^v$ if and only if

$$\sigma_t(h_t, v_1) = \begin{cases} y \in [d(v_1, x), v_1] \text{ for } t = 1 \text{ and } v_1 \text{ for some } t \in \{1, 2, \ldots, T\} & \text{if } v_1 \leq v \\ m \text{ for } t < T \text{ and } v_1 \text{ for } t = T & \text{if } v_1 > v. \end{cases}$$

We define Ω_2^x to be a set of strategies for bidder 2 such that

$$\Omega_2^x = \{\sigma | \sigma_t(h_t, v_2) = \sigma_t^x(h_t, v_2) \text{ for } h_t \text{ such that } w_t \neq 2\}.$$

If (σ_1, σ_2) is an equilibrium and $(\sigma_1, \sigma_2) \in C(v, x)$ then, in both (σ^v, σ^x) and (σ_1, σ_2), bidder 1 bids in period 1 if and only if $v_1 \leq v$. In (σ_1, σ_2), she bids at least $d(v_1, x)$ in period 1 if $v_1 \leq v$ and bids v_1 in some period of the auction. If bidder 1 bids in period 1, then if bidder 2's bid ever surpasses bidder 1's bid, he will be the winner for sure. Both

(σ_1, σ_2) and (σ^v, σ^x) lead to the same history along the path where bidder 2 is not the high bidder. If bidder 2 wins, then the final price is $\max[m, v_1]$ with probability 1. If bidder 1 does not bid in period 1, then (σ_1, σ_2) and (σ^v, σ^x) have the same (w_t, p_t) for all $t \leq T$ and bidder 1 bids v_1 in period T if $v_1 > v$. Therefore, although bidder 1 may place more than one bid in (σ_1, σ_2) if $v_1 \in (m, v]$, for any t, the current high bidder w_t in (σ_1, σ_2) and (σ^v, σ^x) will be the same. If bidder 2 is not the high bidder in period t, he follows the same strategy in any equilibrium. Hence, the price pattern when bidder 1 is the high bidder or she has not placed a bid yet will be the same. Occurrences of sniping by bidder 1 and nibbling by bidder 2 will be the same. Moreover, the winner and final price pair will be the same.

Lemma 3 *If (σ_1, σ_2) is an equilibrium and $(\sigma_1, \sigma_2) \in C(v, x)$, then the winner and transaction pair from (σ_1, σ_2) is the same as that from (σ^v, σ^x) with probability 1.*

Two equilibrium strategy profiles belonging to $C(v, x)$ lead to the same final outcome.

Theorem 3 *Let $(\sigma_1^\epsilon, \sigma_2^\epsilon)$ be a sequence of equilibria converging to (σ_1, σ_2) as ϵ goes to zero. Then, $(\sigma_1, \sigma_2) \in C(v, x)$.*

Bidder 1 bids at least $d(v_1, x)$ when she bids above m in an equilibrium strategy. If bidder 1 bids above m for the first time in period t when $v_1 = v^*$, then, in that equilibrium, bidder 1 does not bid before period t if $v_1 > v^*$. Hence, there will be a partition of possible values of v_1 according to which bidder 1 decides when to place her first bid. There is no equilibrium where the partition has more than two elements. Because ϵ is positive, there is no equilibrium where none of the bidders place a bid in period 1. Then, in any equilibrium, bidder 1 bids in period 1 if v_1 is below a cutoff value and bids in period T otherwise. Bidder 2 chooses bids to maximize his expected payoff knowing that bidder 1 snipes if v_1 is above the cutoff value v. If bidder 1 does not snipe, then she bids close enough to v_1 in period 1 such that bidder 2 faces the same history in periods he is not the high bidder. Since the equilibrium cutoff-action pair is unique, it is optimal for bidder 2 to bid according to σ^x. As a result, any equilibrium can be characterized by (v, x). Final bids of bidders in (σ_1, σ_2) are the same as those in (σ^v, σ^x) with probability 1. Hence, the winner and the transaction price pairs are the same with probability 1.

Since the equilibrium outcome is unique, we can characterize equilibrium properties of this auction by the characteristics of (v, x). We provide an example of the equilibrium cutoff-action pair for uniform F to illustrate (v, x) for an arbitrary T:

$$\begin{aligned} v &= \frac{T + m\sqrt{T}}{T + \sqrt{T}} = 1 - \frac{1-m}{\sqrt{T}+1} \\ x &= \left\{ \frac{t + (T-t)\, m + m\sqrt{T}}{T+\sqrt{T}} \right\}_{t=1}^{T-1} \\ &\Rightarrow \; x_t - x_{t-1} = \frac{1-m}{T+\sqrt{T}}. \end{aligned}$$

As T approaches infinity, $x_t - x_{t-1}$ gets close to zero and v gets close to one. The winner and the transaction price of this action is the same as that of the benchmark model with probability approaching one. Proposition 1 generalizes this result for all F.

Proposition 1 *For all t, $x_t - x_{t-1} > 0$ with $\lim_{T\to\infty} x_t - x_{t-1} = 0$ and $v < 1$ with $\lim_{T\to\infty} v = 1$. The winner and transaction price pair is the same as that of the benchmark model with probability 1 when T approaches infinity.*

Since bidder 2 does not learn anything about v_2 if the price is unchanged, he makes a positive bid increment if he gets a positive signal and is not the high bidder. Because higher bids increase the risk of paying above v_2, bid increments approach zero as T approaches infinity. As a result, bidder 1 can induce almost perfect learning by bidder 2 if she bids in period 1 and, hence, v approaches 1.

A characteristic of the equilibrium is that bidder 2 bids in every period unless he gets a negative signal. If bidder 1 bids in period 1 and wins, she pays a price above v_2. As a result, v approaches 1 only when the expected overpayment approaches zero. This means that the necessary conditions for $v \to 1$ are: i) $x_t - x_{t-1} \to 0$ for all t and ii) $x_{T-1} \to 1$. As T goes to infinity, the bidder with the higher v_i wins with probability approaching 1 and bidder 2's final bid if he loses approaches v_2. As a result, the winner and transaction price pair is the same as that of the benchmark model with probability 1.

In equilibrium, bidder 2 potentially places many bids. Equation 1 shows that his first bid is below the expected value of v_2 conditional on $v_2 \geq m$. He places bids until he gets a negative signal which implies his final bid is above v_2. If he does not get any negative signal in the auction, his final bid equals the conditional expected value of v_2 given $v_2 \geq p_T$.

Hence, bidder 2's final bid overshoots v_2 on average. On the other hand, bidder 1's final bid equals v_1. Thus, the uninformed bidder's final bid will be, on average, higher than the informed bidder's final bid in any equilibrium. An informed bidder cannot benefit by placing multiple bids. Therefore, this model is consistent with the empirical fact that a multiple-bidder's first bid is lower than a single-bidder's first bid, but a multiple-bidder's last bid is higher than a single-bidder's last bid as presented in table 3.

The expected revenue from the auction equals the expected revenue from (σ^v, σ^x). The seller's expected revenue from an auction with an informed and an uninformed bidder is different than that from an auction with two informed bidders. The seller gets a higher expected revenue if bidder 2 is uninformed.

Proposition 2 *The seller's expected revenue if bidder 2 is uninformed is higher than the benchmark model revenue. The revenue difference goes to zero as T approaches infinity.*

The seller gets a higher revenue when bidder 2 is uninformed as his last bid overshoots v_2 on average. Distribution of the final price converges to that of the benchmark model as T goes to infinity and the revenue difference goes to zero. The result that uninformed bidders increase expected revenue for finite T is consistent with the empirical result that auctions with a higher share of multiple-bidders receive higher revenue.

Section 2 showed that sniping does not occur in equilibrium if bidder 2 is an informed bidder. However, if there is a small possibility that bidder 2 is uninformed then sniping occurs in equilibrium. Suppose bidder 1 is informed with certainty and bidder 2 is uninformed with probability $\alpha \in (0, 1)$ and is informed with probability $1 - \alpha$. Bidder 1 does not know whether bidder 2 is uninformed. This auction has an equilibrium where bidder 1 bids his valuation when she places a bid above m. She bids v_1 in period 1 if v_1 is equal to or below some cutoff value v and bids v_1 in period T otherwise. If bidder 2 is informed, he bids v_2 in the first period. If he is uninformed then he tries to learn v_2 by his bids using strategy σ^x as defined earlier in this section. Suppose (v, x) is the equilibrium cutoff-action pair at the limit as ϵ goes to zero if $\alpha = 1$. Then (v, x) leads to an equilibrium for all $\alpha \in (0, 1]$.

Proposition 3 *At the limit as ϵ goes to zero, the equilibrium cutoff-action pair is the same for any $\alpha \in (0, 1]$. The expected revenue is increasing in α.*

When bidder 2 is informed, bidder 1 gets the same expected utility from bidding in periods 1 or T. For all $\alpha \in (0,1]$, the optimal cutoff for a given x is the same. Similarly, for a given v, optimal x is the same. Hence, the same (v, x) is the equilibrium cutoff-action pair at the limit as ϵ goes to zero for any $\alpha \in (0,1]$. Seller's expected revenue conditional on bidder 2 being uninformed is independent of α. Using proposition 2, the expected revenue is increasing in α.

The probability of sniping approaching zero as T approaches infinity depends on the fact that bidder 2 gets infinite number of signals as T approaches infinity. Suppose signals arrived stochastically. In period t, bidder 2 learns whether $v_2 \geq p_t$ with some probability and does not learn anything new about v_2 with some probability. Therefore, the signals are still binary but the arrival is stochastic. Before the auction starts, bidder 2 learns whether v_2 is above m. At the beginning of period $t \in \{2, \ldots, T\}$, private signals arrive following a binomial process with probability $\frac{\lambda}{T-1}$. If he receives a signal, he learns whether $v_2 \geq p_t$ and updates his prior of v_2 using the Bayes' rule. If a signal does not arrive, he does not learn anything about v_2 in that period. Therefore, the expected number of signals is λ which is finite.

Equilibrium bidding behavior, discussed earlier in this section, is robust to uncertain arrival of signals. One can show that the main results from theorems 2 and 3 hold even when signals do not arrive in every period. Any equilibrium can be characterized by a unique equilibrium cutoff-action pair.[6] Sniping and nibbling occur in any equilibrium. Moreover, bidder 2's final bid is, on average, above bidder 1's final bid as before.

Characteristics of the equilibrium cutoff-action pair depend on the expected number of signals λ. For any λ, the equilibrium cutoff stays bounded away from one and the bidder with the lower v_i wins with nonzero probability even when T approaches infinity for small ϵ. That is, sniping occurs with nonzero probability. The winner and final price of this auction are different from those of the benchmark model with nonzero probability. The analysis with T approaching infinity and signals arriving with probability less than one can be thought of as a bidder being able to place as many bids as he wants; however, when he is placing so many bids, he cannot compare v_2 and p_t every time the price changes. He expects to get finite number of opportunities, which arrive stochastically, to learn the

[6] Please see section 4 in Hossain (2004) for a detailed discussion of the stochastic signal arrival case.

relation between v_2 and the price. This leads to nonzero probability of sniping for any T.

Sniping vanishing as T goes to infinity also depends on the assumption that bidder 2 costlessly learns whether $v_2 \geq p_t$. Suppose bidder 2 can learn the relationship between v_2 and p_t in any period, but at a small positive cost κ. This can be thought of as a cost for introspecting one's own preferences around the current price. Then, for any positive κ, there is a finite τ such that even if $T > \tau$, bidder 2 will decide to get a signal at most τ times in any equilibrium. This implies that the probability of sniping will not approach zero as T approaches infinity unlike the case when self-introspection at the price is costless. Interestingly, when κ is positive, there may not be a unique equilibrium outcome. For example, suppose F is uniform, m equals zero, κ equals 0.005 and T is at least 3. Then, there exists an equilibrium where bidder 2 learns whether v_2 is greater than p_t at most twice and bidder 1 snipes if v_1 is greater than $\frac{2}{2+\sqrt{2}}$ which is the equilibrium cutoff when contemplation is costless and T equals 2. There also exists an equilibrium where bidder 2 learns whether v_2 is greater than p_t at most three times and bidder 1 snipes if v_1 is greater than $\frac{3}{3+\sqrt{3}}$. However, for any T, there is no equilibrium where bidder 2 learns whether $v_2 \geq p_t$ four or more times.

4 Auctions with Two Uninformed Bidders

This section analyzes the auction when both bidders are uninformed. One informed player is necessary for sniping. Nevertheless, this game also has interesting properties. Both uninformed bidders nibble to learn about their valuations. The range of experimentation for each bidder depends on the other bidder's experimentation strategy.

Both bidders 1 and 2 are uninformed and they get signals in every period. There is no equilibrium of this auction where either player uses pure strategies in every period. Suppose, for some h_t, bidder 1 bids x_t in period $t < T$ if $v_1 \geq p_t$. Bidder 2's best response is to bid $x_t - \eta_2$ if $v_2 \geq p_t$ where $\eta_2 \to 0^+$. But bidder 1's best response to this is bidding $x_t - \eta_2 - \eta_1$ where $\eta_1 \to 0^+$. Hence, bidders will not follow pure strategies. To allow mixed strategies, we redefine strategy $\sigma = \{\sigma_t\}_{t=1}^{T-1}$ with $\sigma_t : H_t \times [0,1] \to \Delta([0,1])$ where $\Delta(X)$ denotes the set of all probability measures over X.

Without loss of generality, we assume $T = 2$ and $m = 0$ for simplicity. Theorems 4 and proposition 4 can be extended to any T and any opening price. Let strategy σ^G be such that bidder i chooses his period 1 bid from $\left[\underline{b}_1, \bar{b}_1\right)$ according to the distribution G for some positive $\underline{b}_1$. He bids $\mathbf{E}\left[y|\left[p_2, 1\right]\right]$ in period 2 if $v_i \geq p_2$. If $\left(\sigma^G, \sigma^G\right)$ is an equilibrium then G is called an equilibrium distribution. Equilibrium distribution G is such that both bidders get the same expected utility from any bid on $\left[\underline{b}_1, \bar{b}_1\right)$. For all ϵ smaller than some $\bar{\epsilon} > 0$, there exists an equilibrium distribution. All equilibria of this game are symmetric.

Theorem 4 *There is an $\bar{\epsilon} > 0$ such that for all $\epsilon \in [0, \bar{\epsilon})$, there exists an equilibrium distribution G^ϵ such that $\lim_{\epsilon \to 0} G^\epsilon = G$. Then G is an equilibrium distribution of the unperturbed auction.*

An interesting feature is that the auction becomes sort of a coordination game under this equilibrium. Suppose G is an equilibrium distribution. Let $\overline{G}$ be the distribution G truncated from left at $\underline{b}_1 + \psi$ where $\psi < \bar{b}_1 - \underline{b}_1$. Then, both bidders using the distribution $\overline{G}$ will be another equilibrium of this auction.

Proposition 4 *Distribution $\overline{G}^\epsilon$ also leads to an equilibrium of the ϵ-perturbed auction.*

Here $\lim_{\epsilon \to 0} G^\epsilon = G$ implies $\lim_{\epsilon \to 0} \overline{G}^\epsilon = \overline{G}$ and $\overline{G}$ is an equilibrium distribution of the unperturbed auction. If bidder 1 does not bid from the interval $[\underline{b}_1, \underline{b}_1 + \psi]$ then the change in expected payoff for bidder 2 when he bids $y \in \left[\underline{b}_1 + \psi, \bar{b}_1\right)$ is independent of y. Thus, we can get a new equilibrium distribution by truncating an equilibrium distribution from the left. Aggressive bidding by one bidder thus leads to aggressive bidding by the other. Support of G is open on the right and it can be truncated only from the left. Hence, any equilibrium distribution has an open support. This feature would go away if there were an informed bidder in addition to two uninformed bidders because an informed bidder's equilibrium bids are more closely tied to her true valuation.

5 Secret Reserve Price Auctions

In online auctions, the opening price serves as the reserve price. In eBay auctions, the seller can set a secret reserve price r_s in addition to the opening price. The seller announces that there is a secret reserve price but keeps that price a secret. At all time during the auction, she announces whether the secret reserve is met. If the highest bid received in the auction is below the secret reserve then the object stays unsold. Otherwise, the highest bidder pays the higher value of the secret reserve price and the second highest bid.

We model a secret reserve as a bid placed by the seller at the beginning of the auction. When a bid below the secret reserve is received, that bid becomes the current price. If $\max_i \beta^i_{t-1} \leq r_s$ then $p_t = \max_i \beta^i_{t-1}$. Thus, p_t can change even when only one bidder participates. The object stays unsold if $\max_i \beta^i_T < r_s$.

In this section, we assume that the seller has an outside option and she chooses the value of this outside option as the reserve. This value r, unknown to the buyers, is drawn from $[0, R]$ with $R \leq 1$ according to a publicly known distribution. For simplicity, we assume that this distribution is the distribution F truncated from right at R. The seller decides whether to keep this reserve price secret or public.

Suppose there are an informed bidder and an uninformed bidder and $T = 2$. Bidder 1 knows v_1 and bidder 2 learns whether $v_2 \geq p_t$ in every period. The equilibrium cutoff-auction pair of the unperturbed auction is (v, x). If there is a secret reserve price drawn from $[m, R]$ in addition to the opening price, suppose the equilibrium cutoff-action pair is (v^s, x^s). We assume that R equals bidder 2's highest possible equilibrium bid.[7] If F is uniform then $R = \frac{2+m}{3}$. Proposition 5 shows that a secret reserve price reduces the probability of sniping and raises bidder 2's bid in period 1. That is, $v^s > v$ and $x^s_1 > x_1$.

Proposition 5 *The probability of sniping is lower in the secret reserve price auction.*

With two informed bidders and an opening price of m, an auction without any secret reserve has the same bid pattern as an auction that has a secret reserve price in addition. However, if bidder 2 is uninformed, secret reserve price auctions lead to higher cutoffs for bidder 1 as the price can change even if she does not place any bid. Hence, a secret reserve price leads to reduced sniping.

Controlling for the opening price, secret reserve price auctions experienced lower level of sniping than auctions without a secret reserve price in our data set. Columns (1) and (2) of Table 4 report the marginal effects coefficients from probit regressions of the dummy for sniping activity. The marginal effect of having a secret reserve price is negative and quite large. A secret reserve price auction was about one-third less likely to experience

[7] The value of R is such that $\mathbf{E}[y | [\mathbf{E}[y | [m, R]], 1]] = R$.

sniping than an auction with the same opening price and no secret reserve. Proposition 5 is consistent with this finding.

Secret reserve prices also let us test an alternative theory that may explain the empirical result that the first bids of multiple-bidders were lower than the first bids of single-bidders, but the last bids of multiple-bidders were higher. Suppose some bidders obtain a small enjoyment from placing a bid. Such a bidder may divide her final bid into multiple bids and she may place a bid even when the price is above her valuation. In that case, the expected value of her first bid is lower than that of a standard bidder on average while the expected value of her last bid is higher.

Before a bid has been received in an auction, the current price equals the opening price whether or not there is a secret reserve price. Therefore, in our model, having a secret reserve price will not affect the probability of receiving at least one bid in the auction. Similarly, a secret reserve price will have no effect on whether a bid was received in the benchmark model. Between two auctions with the same opening price and the same bidder arrival process, a bid is less likely to be the winning bid in the secret reserve price auction as the real reserve price in that auction is higher. If a bidder obtained utility from placing bids, she is more likely to place a bid even when the opening price is above her valuation if there is a secret reserve. Hence, a secret reserve price will increase the probability of at least one bid being received if bidders obtained utility from placing bids. The last two columns of Table 4 report the marginal effects coefficients from probit analysis of the dummy for at least one bid received.[8] The coefficients of the dummy for secret reserve price auctions are insignificant. Hence, secret reserve price auctions were not less likely to receive no bids at all.

When all bidders are informed, keeping the reserve price secret or public leads to the same final outcome. If the opening price equals the reserve r, only bidders with valuation above r bid their valuations. When the opening price m is below r and the secret reserve r_s equals r, bidders with valuations above m bid their valuations. Any bid between m and r is, basically, ignored. The final price and the winner of the two auctions will be the same. The seller has to pay an extra fee for keeping the reserve

[8]The share of multiple bidders equals number of multiple bidders divided by the total number of bidders if the dependent variable equals 1 and zero otherwise. Hence, we do not use it in columns (3) and (4).

secret on eBay. Therefore, keeping the reserve secret is not rational when all bidders are informed. Nevertheless, secret reserve price auctions are quite common on eBay. Bajari and Hortacsu (2002) found that sellers frequently chose to keep the reserve price secret for objects with relatively high book values. In our data set, more than a quarter of the auctions were secret reserve price auctions.

In our model, an uninformed bidder may win the object even when his valuation is below the reserve if it is secret. That does not happen if the reserve is public. As a result, a secret reserve price may lead to a higher revenue for the seller. Suppose F is uniform and $m = 0$. The expected revenue is higher if the reserve is secret.

Proposition 6 *If F is uniform, the expected revenue is higher when r is kept secret.*

Thus, in some cases, having a secret reserve price can be profitable for the seller even after paying a fee for keeping the reserve secret. This also implies that the revenue of auctions depend on both the effective reserve and whether the reserve is public or secret.

6 Conclusion

This paper suggests that people do not always know their exact private valuation for a good. Applying the idea of spontaneous learning at a posted price, we suggest a new approach in explaining how an agent learns her own type. We introduce a dynamic second-price auction where one bidder knows her private valuation and the other bidder can only learn whether his valuation is above the current price any time during the auction. In any equilibrium of this auction, there will be sniping and nibbling as evidenced in eBay auctions. We also show that secret reserve price auctions experience less sniping and can be more profitable than public reserve price auctions.

This paper justifies the differences in bid patterns between auctions held by eBay and by the auction unit of *Amazon.com,* another major online auctioneer. Amazon's auction format is similar to eBay's with the difference that Amazon auctions have "soft" closing times. In Amazon auctions, the length of the auction is increased if there is a bid within the last ten minutes of the pre-announced closing time. After that, the auction closes when there has not been any activity for ten minutes. Roth and Ockenfels (2001) found that auctions on Amazon had frequent occurrences of multiple bids but infrequent occurrences of late bids. Ariely, Ockenfels and Roth (2004) ran laboratory experiments to mimic eBay and Amazon auctions for private-value objects. They found significant presence of multiple bids and late bids in eBay style auctions and significant presence of only multiple bids in Amazon style auctions.[9] If we incorporate soft closing time in our model, there will not be any benefit from sniping as it will just lengthen the auction. Costless comparison of the posted price and valuation will still give rise to nibbling.

The main idea proposed in this paper is, in a dynamic mechanism such as an ascending auction, people get new information about their own preferences in addition to information about other players' preferences as they participate in the game. The paper suggests a method to model this learning. The current price any time during the auction plays a key role in the learning process. A natural extension of this model will be analyzing various auction mechanisms when bidders can contemplate at any hypothetical price for a small cost. This may lead to an interesting decision theoretic exploration of the cognitive process of learning one's type. Learning about one's private type through the process of negotiation can also be used in bargaining or principal-agent problems.

[9] However, in that study, all subjects were told their monetary valuations of winning before the auctions started.

A Appendix

Proof of Lemma 1

Proof. Suppose $\beta_T^2 = \beta^*$. If $v_1 \geq \beta^*$ then bidder 1 strictly prefers $\beta_T^1 = v_1$ over $\beta_T^1 < \beta^*$ and is indifferent between $\beta_T^1 = v_1$ and $\beta_T^1 \geq \beta^*$. If $v_1 < \beta^*$ then she strictly prefers $\beta_T^1 = v_1$ over $\beta_T^1 \geq \beta^*$ and is indifferent between $\beta_T^1 = v_1$ and $\beta_T^1 < \beta^*$. Similarly bidder 2 weakly prefers $\beta_T^2 = v_2$ over $\beta_T^2 \neq v_2$ for all β_T^1. Therefore, for all $h_T \in H_T$, if σ_i is weakly undominated then in that strategy, for all t, $b_t^i \leq v_i$ and there exist a τ such that $b_\tau^i = v_i$. ■

Proof of Theorem 1

Proof. It is an equilibrium where $b_1^i = v_i$ if $v_i \geq m$ for all $i \in \{1, 2\}$. Suppose there is an equilibrium where there is a $t_i \leq T$ such that $b_t^i < v_i$ for all $t < t_i$ and $b_{t_i}^i = v_i$ if $v_i \geq m$. Moreover, without loss of generality, $t_1 \geq t_2$ and $t_1 \geq 2$. If bidder 1 chooses $b_{t_1-1}^1 = v_1$ instead, her expected payoff increases because the probability of dropping out is positive. Then she prefers $b_{t_1-1}^1 = v_1$ to $b_{t_1-1}^1 < v_1$. By backward induction, $b_1^i = v_i$ if $v_i \geq m$ is the only equilibrium. ■

Proof of Lemma 2

Proof. In the appendix, we will use $e(p)$ to represent $\mathbf{E}[y|\,[p,1]]$. Bidders 1 and 2 are informed and uninformed respectively. Suppose β_T^1 is drawn from the continuous distribution G with support $[a, c]$ where $p_T \leq a$ and $c \leq 1$. Bidder 2 knows $v_2 \in [p_T, 1]$ and chooses b_T^2 to maximize $\int_{p_T}^1 \int_a^{b_T^2} (v - y)\, dG(y)\, dF(v)$ if he has not dropped out of the auction. The first order condition requires,

$$\int_{p_T}^1 \left(v - b_T^2\right) g\left(b_T^2\right) dF(v) = 0.$$

Therefore, choosing $b_T^2 = e(p_T)$ is strictly better than any other strategy if $g(e(p_T)) > 0$. If $a > e(p_T)$ or $c < e(p_T)$ then any b_T^2 in $[m, a)$ and $[c, 1)$ respectively is optimal. Therefore, $b_T^2 = e(p_T)$ if $v_2 \geq p_T$ and $b_T^2 \leq p_T$ if $v_2 < p_T$ is weakly dominant. ■

Proof of Theorem 2

Proof. When $v_1 = v^\epsilon$, bidder 1's expected utility from bidding in period 1 equals

$$\begin{aligned}&(v^\epsilon - m) F(m) + \sum_{t=1}^{T-1} (1-\epsilon)^{t-1} (v^\epsilon - x_t^\epsilon) \left(F(x_t^\epsilon) - F\left(x_{t-1}^\epsilon\right) + \epsilon (1 - F(x_t^\epsilon))\right) \\ &+ (1-\epsilon)^{T-1} \left(1 - F\left(x_{T-1}^\epsilon\right)\right) \left(v^\epsilon - e\left(x_{T-1}^\epsilon\right)\right) \mathbf{1}_{\left\{v^\epsilon > e\left(x_{T-1}^\epsilon\right)\right\}}.\end{aligned}$$

Her expected utility from bidding in period T equals

$$v^\epsilon - mF(m) + (1 - F(m)) (1-\epsilon)^{T-1} \left(v^\epsilon - \left(1 - (1-\epsilon)^{T-1}\right) x_1^\epsilon - (1-\epsilon)^{T-1} e(m)\right).$$

If $\epsilon = 0$ then bidder 1 is better off by bidding in period T if $v_1 > e(x_{T-1})$. For small ϵ, she is indifferent between bids in periods 1 and T when $v_1 = v^\epsilon < 1$ where v^ϵ equals

$$\frac{\begin{array}{c}(1-\epsilon)^{T-1} \left((1-\epsilon)^{T-1} \int_m^1 y dF(y) + \left(1 - (1-\epsilon)^{T-1}\right) (1 - F(m)) x_1^\epsilon\right) \\ - \sum_{t=1}^{T-1} (1-\epsilon)^{t-1} x_t^\epsilon \left(F(x_t^\epsilon) - F\left(x_{t-1}^\epsilon\right) + \epsilon (1 - F(x_t^\epsilon))\right)\end{array}}{(1-\epsilon)^{T-1} (1 - F(m)) - \sum_{t=1}^{T-1} \left(F(x_t^\epsilon) - F\left(x_{t-1}^\epsilon\right) + \epsilon (1 - F(x_t^\epsilon))\right) (1-\epsilon)^{t-1}}. \quad (3)$$

Suppose the highest k such that $1 \in A_k$ equals k^*. Let bidder 2's belief system $\mu_2(h_t)$ for $t \geq 2$ be, if $k^* = 1$ then $v_1 \in (p_2, v^\epsilon]$ if $w_t = 1$ and $v_1 = p_t$ if $w_t = 2$. If $k^* > 1$ then $v_1 \in (\max[e(p_t), v^\epsilon], 1]$ if $w_t = 1$ and $v_1 = p_t$ if $w_t = 2$. If $1 \notin A_k$ for all $k < t$ then $v_1 \in [0, m] \cup (v^\epsilon, 1]$. No other strategy gives bidder 1 a higher payoff. Suppose $b_1^1 = b^* \in (x_\tau^\epsilon, v_1)$ for a $\tau \in \{0, 1, \ldots, T-1\}$. Given σ^x and μ_2, bidding v_1 before period T cannot be profitable for bidder 1. Now, since ϵ is positive, bidder 1 prefers $b_1^1 = v_1$ to $b_1^1 = b^*$ if $e(b^*) \geq v_1$ (preference is weak if $x_\tau^\epsilon \leq v_1 < x_{\tau+1}^\epsilon$). If $v_1 > e(b^*)$, then she prefers $b_1^1 = m$ to $b_1^1 = b^*$ as ϵ is small enough that paying $e(b^*)$ if $v_2 > b^*$ is less profitable although the probability that $\beta_T^1 = v_1$ is $(1-\epsilon)^{T-1}$. This is true because paying $\mathbf{E}[y| [m, b^*]]$ if $v_2 \leq b^*$ is more profitable than paying the right Riemann approximation of $\mathbf{E}[y| [m, b^*]]$ which is larger as long as one interval has nonzero measure. Thus, there is no profitable deviation and σ^v is a best response to σ^x.

Suppose, in period $t < T$, $v_2 \geq p_t$ and $w_t \neq 2$. Expressing v^ϵ by x_T^ϵ, x_t^ϵ maximizes

$$(1-\epsilon) \int_{x_t^\epsilon}^1 \int_{x_{t-1}^\epsilon}^{x_{t+1}^\epsilon} (z - y) dF(y) dF(z) + \epsilon \int_{x_{t-1}^\epsilon}^1 \int_{x_{t-1}^\epsilon}^{x_t^\epsilon} (z - y) dF(y) dF(z).$$

The first order conditions require

$$(1-\epsilon) \int_{x_{t-1}^\epsilon}^{x_{t+1}^\epsilon} (y - x_t^\epsilon) dF(y) + \epsilon \int_{x_{t-1}^\epsilon}^1 (y - x_t^\epsilon) dF(y) = 0. \quad (4)$$

Suppose x^ϵ satisfies the first order conditions. Then the second order conditions are also satisfied at x^ϵ and boundary points are not optimal. Using the Envelope theorem, we can use a second order condition on a single variable around x^ϵ. Bidder 2's ex-ante expected utility maximizer is dynamically consistent. From the first order conditions, it is easy to see that $x_t^\epsilon - x_{t-1}^\epsilon > 0$. Both σ^v and σ^x are optimal for all possible histories. Therefore, if such σ^v and σ^x exist then (σ^v, σ^x) is an equilibrium.

We can use Brouwer's fixed point theorem on $M : [m,1]^T \longrightarrow [m,1]^T$ to show existence. The first $T-1$ elements of M are the functions for bids derived from equation 4. The last element is the function $\max\left[M^C(v^\epsilon, x^\epsilon), m\right]$; $M^C(v^\epsilon, x^\epsilon)$ is given by expression 3. The domain is non-empty, convex, and compact and M is continuous. By Brouwer's fixed point theorem, there exists a fixed point of M. By construction, $x_t^\epsilon \leq x_{t+1}^\epsilon$ and $x_{T-1}^\epsilon \leq v^\epsilon$. Furthermore, $v^\epsilon > m$ at the fixed point. Any fixed point of the system generates an equilibrium bid schedule for bidder 2 and cutoff value for bidder 1.

This fixed point is unique. First, there is no $v^\epsilon \leq 1$ such that $(v^\epsilon, x_a^\epsilon)$ and $(v^\epsilon, x_b^\epsilon)$ that satisfy the equilibrium criteria. If $x_{a,1}^\epsilon < x_{b,1}^\epsilon$ then that will lead to $x_{a,2}^\epsilon < x_{b,2}^\epsilon$ as bidder 2 maximizes his expected payoff with the same cutoff but a higher signal in period 2. On the other hand, a higher bid in period 2 makes a higher bid in period 1 more profitable. In that sense there is strategic complementarity. Moreover, strict single-crossing property defined by Milgrom and Shannon (1994) is satisfied. If $x_{a,1}^\epsilon < x_{b,1}^\epsilon$ then $x_{a,t}^\epsilon < x_{b,t}^\epsilon$ for all t. Then, the cutoff corresponding to x_a^ϵ is smaller than the cutoff corresponding to x_b^ϵ.

Now suppose $(v_a^\epsilon, x_a^\epsilon)$ and $(v_b^\epsilon, x_b^\epsilon)$ are two equilibrium cutoff-action pairs with $v_a^\epsilon < v_b^\epsilon \leq 1$. Hence, $x_{a,t}^\epsilon < x_{b,t}^\epsilon$ for all t. Bidder 1's expected payoff from bidding in the first period is the right Riemann approximation of the area under the decreasing curve $v_1 - y$, $\int_m^{v_1} (v - y)\, dF(y)$ using x for the intervals. Therefore, if $v_1 = v_a^\epsilon$, bidder 1's expected payoff from bidding in period 1 is higher when bidder 2 follows x_a^ϵ instead of x_b^ϵ. She is indifferent between bidding in period 1 and T when bidder 2 follows x_a^ϵ. In the unperturbed game, she gets the same utility from bidding late whether bidder 2 follows x_a^ϵ or x_b^ϵ. Therefore, for small enough ϵ, the difference in payoff from bidding late with x_a^ϵ and x_b^ϵ is small enough that bidder 1 strictly prefers bidding late if bidder 2 follows x_b^ϵ. That would imply that $v_b^\epsilon < v_a^\epsilon$. Therefore, a unique equilibrium cutoff-action pair (v^ϵ, x^ϵ) exists for all ϵ smaller than some $\overline{\epsilon} > 0$.

In the unperturbed auction,

$$v=\frac{\int_m^1 ydF(y)-\sum_{t=1}^{T-1}x_t\left(F(x_t)-F(x_{t-1})\right)}{1-F(x_{T-1})}=\frac{P(T,v)-P(1,v)}{W(T,v)-W(1,v)}$$

and $x_t=\mathbf{E}\left[y|\left[x_{t-1},x_{t+1}\right]\right]$. Moreover, $v=\lim_{\epsilon\to 0}v^{\epsilon}$ and $x_t=\lim_{\epsilon\to 0}x_t^{\epsilon}$. ■

To prove theorem 3, we use several lemmas. Suppose, if bidder 1 bids v_1 in period 1, $\chi(\sigma_2)$ is bidder 2's action in σ_2 along a history where he is not the high bidder and does not get a negative signal. When (σ_1,σ_2) is played, bidder 1's first bid above m is between $d(v_1,\chi(\sigma_2))$ and v_1 where $d(a,b)$ is the largest scalar smaller than or equal to a in the sequence b. Therefore, (σ_1,σ_2) can be characterized by an equilibrium where bidder 1 bids v_1 when she places a bid. The possible values of v_1 can be divided into a 2-element partition such that if v_1 is in the first element, bidder 1 bids before period T, and she bids only in period T otherwise. In any equilibrium (σ_1,σ_2), both bidders bid in period 1 with positive probability. Since there is a unique equilibrium cutoff-action pair (v,x), bidder 1 bids in $[d(v_1,\sigma_2),v_1]$ in period 1 if $v_1<v$ and bids only in period T if $v_1>v$ in (σ_1,σ_2). Moreover, bidder 2 bids x_t along a history where he is not the high bidder and gets a positive signal as the history in that path is the same for any equilibrium σ_1.

Lemma 4 *When (σ_1,σ_2) is played, $b_t^1>\max\left[p_t,\beta_{t-1}^1\right]$ implies $b_t^1\in\left[d(v_1,\chi(\sigma_2)),v_1\right]$.*

Proof. Strategy σ_1 is weakly undominated; hence, this lemma is true for $t=T$. We will show that $b_t^1\in\left(\max\left[p_t,\beta_{t-1}^1\right],d(v_1,\chi(\sigma_2))\right)$ cannot be an equilibrium. If bidder 1 bids $d(v_1,\chi(\sigma_2))$ the first time she places a bid, her future bids do not affect her payoff. Essentially, any equilibrium strategy is characterized by the strategy where bidder 1 bids v_1 the first time she places a bid. Suppose $b^*\in(m,d(v_1,\chi(\sigma_2)))$. Without loss of generality, we consider the following strategies:

i) $b_1^1=b^*$ and $b_\tau^1>b^*$ for some $\tau<T$ if $h_\tau\in H_\tau^*\subseteq H_\tau$ and $b_T^1=v_1$ if $h_\tau\notin H_\tau^*$,

ii) $b_1^1=b^*$, $b_t^1\leq p_t$ for $t\in\{2,\dots,T-1\}$, and $b_T^1=v_1$ for all $h_T\in H_T$.

Case i) Suppose $H_\tau^*=H_\tau$. If $w_\tau=1$, then h_τ is the same as it would be had b_1^1 equaled v_1. Let $h_{t-1},\widetilde{h}_t$ stand for the list created by appending $\widetilde{h}_t$ at the end of h_{t-1}. If $\mu_2(h_\tau,(p_{\tau+1},w_{\tau+1},\{2\}))=\mu_2(h_\tau,(p_{\tau+1},w_{\tau+1},\{1,2\}))$, then bidder 1 gets a higher expected utility by choosing $b_1^1=v_1$ as $\epsilon>0$. If $\mu_2(h_\tau,(p_{\tau+1},w_{\tau+1},\{2\}))\neq\mu_2(h_\tau,(p_{\tau+1},w_{\tau+1},\{1,2\}))$, then

$$\inf\left(\mathrm{supp}\mu_2(h_\tau,(p_{\tau+1},w_{\tau+1},\{2\}))\right)<\inf\left(\mathrm{supp}\mu_2(h_\tau,(p_{\tau+1},w_{\tau+1},\{1,2\}))\right)$$

or

$$\sup\left(\text{supp}\mu_2\left(h_\tau,\left(p_{\tau+1},w_{\tau+1},\{2\}\right)\right)\right)<\sup\left(\text{supp}\mu_2\left(h_\tau,\left(p_{\tau+1},w_{\tau+1},\{1,2\}\right)\right)\right)$$

as bidder 1 can profitably deviate by pretending to have a lower v_1 otherwise. Then, b_t^2 will be weakly higher with $\mu_2\left(h_\tau,\left(p_{\tau+1},w_{\tau+1},\{1,2\}\right)\right)$. Bidder 1 is worse off with $b_1^1=b^*$ as her expected payment is higher and expected probability of winning is lower.

If $w_\tau=2$ and $p_\tau=b^*$, then suppose $b_\tau^2=x_t(b^*)<e(b^*)$ if bidder 2 has not yet received a negative signal. If $v_1\in[b^*,x_t(b^*)]$, then $b_1^1\geq d(v_1,\chi(\sigma_2))$ is better than $b_1^1=b^*$ for bidder 1. In equilibrium, a player has correct beliefs about other player's strategies. Therefore, when $w_\tau=2$ then bidder 2 knows either $v_1=b^*$ or $v_1>x_t(b^*)$. As a result, choosing $b_\tau^2>x_t(b^*)$ will be optimal for bidder 2 unless $x_t(b^*)=e(b^*)$. However, if $x_t(b^*)=e(b^*)$ then bidder 1 will be better off with $b_1^1\geq d(v_1,\chi(\sigma_2))$ because $\epsilon>0$ and $\beta_T^2\geq e(b^*)$ if $b_\tau^2>p_\tau$ when $w_\tau=2$.

If $H_\tau^*\subset H_\tau$ then if both b_1^1 and b_τ^1 are above the current price if $h_\tau\in H_\tau^*$ with $\tau<T$, then $v_1\geq e(p_\tau)$. This is true because if $h_\tau\notin H_\tau^*$, then $\beta_T^2=e(b^*)$ with probability 1 if $v_2\geq b^*$ and bidder 2 receive a signal at b^*. Then $1\in A_1$ and $1\in A_\tau$ implies $v_1\in[e(p_\tau),v]$ where $v\leq 1$. In that case, $b_{\tau+1}^2\geq e(p_\tau)$ is optimal for bidder 2 as $\beta_\tau^1>e(p_\tau)$. Then it will be optimal for bidder 1 to choose $b_t^1<p_t$ for $t<T$ and $b_T^1=v_1$. Hence, $b_1^1=b^*<v_1$ and $b_\tau^1=v_1$ for some $\tau<T$ if $h_\tau\in H_\tau^*$ is not an equilibrium strategy.

Case ii) Suppose bidder 1 follows the strategy $b_1^1=b^*$, $b_t^1\leq p_t$ for $t<T$ and $b_T^1=v_1$. If $v_1<e(b^*)$, she is better off by choosing $b_1^1=v_1$ and if $v_1\geq e(b^*)$ then she is better off with $b_1^1=m$ because $\epsilon\to 0$. ■

Now we can restrict attention to strategies where bidder 1 bids v_1 when she bids above m for the first time.

Lemma 5 *If in an equilibrium $b_\tau^1=v_1$ and $b_t^1\leq m$ for $t<\tau$ when $v_1=v^*$, then $b_t^1\leq m$ for $t<\tau$ when $v_1>v^*$.*

Proof. If this lemma is not true then there exists a value v^* such that $b_\tau^1=v_1$ if $v_1=v^*$ and $b_{\tau-k}^1=v_1$ if $v_1\in(v^*,v_b]$. If the lowest valuation for which bidder 1 bids in period $\tau-k$ is above v^* then bidder 2's best response to that is bidding at least v^*. Then if v_1 is slightly above v^*, bidder 1 gets zero payoff from bidding in period $\tau-k$. Then she would be better off by bidding in period τ as she would get a positive payoff.

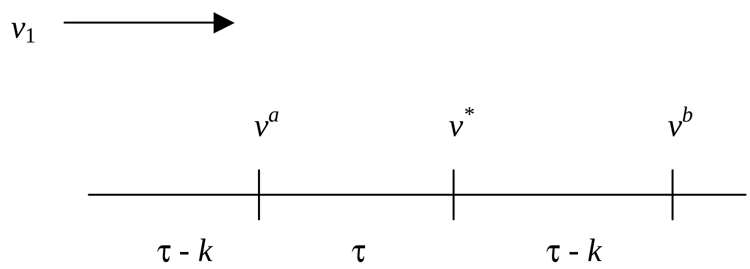

Figure 1: Illustration of the Suggested Contradiction

Now suppose bidder 1 bidding in period $\tau - k$ implies that $v_1 \leq v^a$ or $v_1 \in \left(v^*, v^b\right]$. Then bidder 2 will bid at least v^* when the price goes above v^a if bidder 1 bids in period $\tau - k$. As a result, if v_1 is slightly above v^*, bidder 1 loses the auction if v_2 is not below v^a. In equilibrium, if $v_1 = v^*$, bidder 1's expected utility from $b_\tau^1 = v_1$ is strictly higher than it will be if she chose the optimal strategy for $v_1 = v^a$. That implies that she will be better off by choosing $b_\tau^1 = v_1$ if v_1 is just above v^*. Therefore, this lemma holds true. ■

Hence, there is a partition of v_1 that determines when bidder 1 bids.

Lemma 6 *The size of the equilibrium partition of v_1 equals two.*

Proof. Suppose, in equilibrium, bidder 1 bids in period τ if $v_1 \in [m, v^a]$ for a $v^a > e(m)$ and bids in period $\tau + k$ if $v_1 \in \left(v^a, v^b\right]$ where $\tau + k < T$ and $v^b \leq 1$. Then, if $1 \in A_{\tau+k}$ then bidder 2 bids at least v^a in period $\tau + k + 1 < T$ unless he gets a negative signal. If $\tau + k + 1 = T$ then bidder 2 bids $e(p_T)$. If bidder 1 bids only in period T bidder 2 bids $e(m)$ in any weakly undominated strategy. For v_1 slightly above v^a, either of bidding in period τ or bidding in period T will lead to higher payoffs. Therefore, in any equilibrium, bidder 1 bids only in period τ or T. ■

Lemma 7 *In (σ_1, σ_2), both bidders bid in period 1 with positive probabilities.*

Proof. Using lemmas 5 and 6, there is a cutoff value $v \leq 1$ such that if $v_1 \in [0, v]$, bidder 1 bids in period $\tau < T$ and she bids only in period T otherwise. Suppose $\tau > 1$ in σ_1 and x_τ is bidder 2's first bid above m if $v_2 \geq m$. Since ϵ is positive and close to zero, bidder 1 will be better off by bidding at least x_τ in period $\tau - 1$ for $v_1 \in [x_\tau, x_{\tau+1}]$ where $x_{\tau+1}$ is bidder 2's smallest possible bid in period $\tau + 1$ if he does not get a negative signal. Therefore, $\tau = 1$ and bidder 1 bids in period 1 if v_1 is low enough and bids in period T

otherwise. If $v_2 \geq m$ then bidder 2 bids in period 1 irrespective of bidder 1's strategy since ϵ is positive. Hence, both bidders bid in period 1 with positive probabilities. ■

Proof of Lemma 3

Proof. In any equilibrium $(\sigma_1, \sigma_2) \in C(v, x)$, $\Pr(1 \in A_1)$ and w_t are the same and $\beta_T^1 = v_1$. Bidder 2 faces the same information sequence. Hence, his final bid, the winner and transaction price of (σ_1, σ_2) are same as those of (σ^v, σ^x) with probability 1. ■

Proof of Theorem 3

Proof. From lemmas 4 to 7 and uniqueness of (v, x), it follows that bidder 1 bids at least $d(v_1, \chi(\sigma_2))$ in period 1 if $v_1 < v$ and snipes if $v_1 > v$. Bidder 1's final bid equals v_1 in any weakly undominated strategy. Since the price pattern when $w_t = 1$ or $1 \notin A_1$ is the same for any σ_1, bidder 2 bids according to σ^x if $w_t \neq 2$ and his final bid, when he is not the winner, is the same for any optimal σ_2. He may choose from many optimal strategies different from σ^x when he is the high bidder. Nevertheless, even if $w_t = 2$, he never bids above $\mathbf{E}[y| [p_t, 1]]$ and his final bid equals $\mathbf{E}[y| [p_T, 1]]$ if $v_2 \geq p_T$. Therefore, if (σ_1, σ_2) is an equilibrium then $(\sigma_1, \sigma_2) \in C(v, x)$. ■

Proof of Proposition 1

Proof. Suppose (v, x) is an equilibrium cutoff-action pair and there exists a t such that x_t and x_{t-1} are arbitrarily close. Since F is atomless and strictly increasing, $F(x_t) - F(x_{t-1}) \to 0$. By using theorem 3, $x_t = \mathbf{E}[y| [x_{t-1}, x_{t+1}]]$. Thus, $x_t - x_{t-1} \to 0$ for some t implies $x_t - x_{t-1} \to 0$ for all $t < T$. Similarly, $x_t - x_{t-1}$ bounded away from zero for some t implies so is true for all $t < T$. Suppose $x_t - x_{t-1} \to 0$ for all t when T is small. Then, $x_{T-1} \to m$ and $v \to e(m)$ which implies $x_{T-1} \to \mathbf{E}[y| [m, e(m)]] > m$. Hence, $x_t - x_{t-1} > 0$ for all t in that case.
This implies that $v < 1$ because

$$v = \frac{\int_{x_{T-1}}^{1} y dF(y) + \sum_{t=1}^{T-1} \int_{x_{t-1}}^{x_t} (y - x_t) dF(y)}{1 - F(x_{T-1})} < \frac{\int_{x_{T-1}}^{1} y dF(y)}{1 - F(x_{T-1})} \leq 1 \ .$$

As T approaches infinity, $x_t - x_{t-1}$ approach zero for all t because x_{T-1} approaches infinity otherwise. Therefore, $\lim_{T\to\infty} (x_t - x_{t-1}) = 0$ for all t. Now suppose, $\lim_{T\to\infty} x_{T-1} = 1 - \eta$

for some $\eta > 0$. Since $\lim_{T\to\infty}(x_t - x_{t-1}) = 0$ for all t, $x_{T-1} = \mathbf{E}[y|[x_{T-2}, v]]$ implies $v \to 1-\eta$. However, equation 2 implies that $v \to e(1-\eta)$. This is impossible if $\eta > 0$. Therefore, $\eta = 0$ when T approaches infinity and $\lim_{T\to\infty} v = 1$. We also get $\lim_{T\to\infty} v = 1$ by using the l'Hôpital's rule with $x_t - x_{t-1} \to 0$ for all t and $x_{T-1} \to 1$.

The probability of the bidder with the higher v_i winning the auction equals

$$1 - \sum_{t=1}^{T} \int_{x_{t-1}}^{x_t} (F(y) - F(x_t))\, dF(y) - \int_v^1 (1 - F(y))\, dF(y).$$

This probability approaches 1 and bidder 2's last bid when bidder 1 wins converges to v_2 only as T approaches infinity. Bidder 1's last bid equals v_1 with probability 1. The final outcome is same as the unique benchmark model outcome as $T \to \infty$. ■

Proof of Proposition 2

Proof. When bidder 2 is uninformed, β_T^2 is either above v_2 or equals his conditional valuation given $v_2 \geq p_T$ with probability 1. That is, bidder 2 bids above v_2 in expectation. On the other hand, $\beta_T^1 = v_1$ with probability 1. Seller's expected revenue,

$$\begin{aligned}\pi_U &= 2mF(m)(1-F(m)) + \int_v^1 \int_m^1 y dF(y)\, dF(z) \\ &+ \sum_{t=1}^{T} \int_{x_{t-1}}^{x_t} \left(\sum_{k=1}^{t-1} x_k (F(x_k) - F(x_{k-1})) + y(1 - F(x_{t-1})) \right) dF(y) \qquad (5)\end{aligned}$$

where $x_0 = m$ and $x_T = v$. With two informed bidders, the seller's expected revenue,

$$\begin{aligned}\pi_I &= 2mF(m)(1-F(m)) + \int_v^1 \left(\int_m^z y dF(y) + z(1-F(z)) \right) dF(z) \\ &+ \int_m^v \left(\int_m^z y dF(y) + z(1-F(z)) \right) dF(z). \qquad (6)\end{aligned}$$

When T is small, $(x_t - x_{t-1}) > 0$ for all t and $v < 1$. Hence, the seller gets higher expected revenue when bidder 2 is uninformed. From proposition 1, $\lim_{T\to\infty}(x_t - x_{t-1}) = 0$ for all t and $\lim_{T\to\infty} x_{T-1} = \lim_{T\to\infty} v = 1$. As T approaches infinity, the seller's expected revenue with informed bidder 1 and uninformed bidder 2 approaches the expected revenue with two informed bidders. ■

Proof of Proposition 3

Proof. When $v_1 = v$, bidder 1's expected payoff from bidding v_1 in period 1 is

$$F(m)(v-m) + \alpha \left(v \left(F(x_{T-1}) - F(m) \right) - \sum_{t=1}^{T-1} x_t \left(F(x_t) - F(x_{t-1}) \right) \right)$$
$$+ (1-\alpha) \int_m^v (v-y)\, dF(y)$$

at the limit as ϵ approaches zero. Her expected payoff from sniping is

$$F(m)(v-m) + \alpha \int_m^1 (v-y)\, dF(y) + (1-\alpha) \int_m^v (v-y)\, dF(y).$$

For all α, bidder 1 is indifferent between bidding in periods 1 and T if $v_1 = v$ because

$$v = \frac{\int_m^1 y dF(y) - \sum_{t=1}^{T-1} x_t \left(F(x_t) - F(x_{t-1}) \right)}{1 - F(x_{T-1})}.$$

Since (v, x) is the unique equilibrium cutoff-action pair for $\alpha = 1$, x is optimal for bidder 2 when he is uninformed. Therefore, (v, x) is the equilibrium cutoff-action pair for all $\alpha \in (0, 1]$. Hence, expected revenue conditional on bidder 2 being uninformed equals π_U no matter what α is. Seller's expected revenue equals $\alpha \pi_U + (1-\alpha) \pi_I$ where π_U and π_I are given in equations 5 and 6 respectively. As π_U is greater than π_I for any given T, the expected revenue is increasing in α. ■

Proof of Theorem 4

Proof. Given that bidder 2 follows σ^G, bidder 1's expected utility from bidding $y \in \left[\underline{b}_1, \overline{b}_1\right)$ in period 1 equals

$$\int_{\underline{b}_1}^{y} \int_0^1 \left((1-\epsilon) F(u) + \epsilon \right) (z-u)\, dF(z)\, dG^\epsilon(u)$$
$$+ (1-\epsilon) \left((1-\epsilon) F(y) + \epsilon \right) \int_y^{\overline{b}_1} \int_y^1 (z-u)\, dF(z)\, dG^\epsilon(u) = K_1^\epsilon.$$

In equilibrium, K_1^ϵ is independent of y. Hence, G^ϵ solves

$$f(y) \left(\left((1-\epsilon) F(y) + \epsilon \right) \int_y^{\overline{b}_1} (u-y)\, dG^\epsilon(u) + (1-\epsilon) \int_y^1 \int_y^{\overline{b}_1} (z-u)\, dG^\epsilon(u)\, dF(z) \right)$$
$$+ g^\epsilon(y) \left((1-\epsilon) F(y) + \epsilon \right) \left(\int_0^y (z-y)\, dF(z) + \frac{\epsilon}{1-\epsilon} \int_0^1 (z-y)\, dF(z) \right) = 0 \quad (7)$$

with $G^{\epsilon}\left(\underline{b}_1\right)=0$ and $G^{\epsilon}\left(\bar{b}_1\right)=1$. To show existence, define $\Gamma(y)=\int_y^{\bar{b}_1} G^{\epsilon}(u)\,du$. Then, $G^{\epsilon}(y)=-\Gamma'(y)$ and $g^{\epsilon}(y)=-\Gamma''(y)$. Equation 7 becomes

$$\begin{aligned} & f(y)\left(\int_y^1 z dF(z)-\left(1-2F(y)-\frac{\epsilon}{1-\epsilon}\right)\bar{b}_1-y\left(F(y)+\frac{\epsilon}{1-\epsilon}\right)\right) \\ & +f(y)\left(1-2F(y)-\frac{\epsilon}{1-\epsilon}\right)\Gamma(y)+f(y)\int_y^1 (z-y)\,dF(z)\,\Gamma'(y) \\ -\left((1-\epsilon)F(y)+\epsilon\right) & \left(\int_0^y (z-y)\,dF(z)+\frac{\epsilon}{1-\epsilon}\int_0^1 (z-y)\,dF(z)\right)\Gamma''(y)=0. \qquad (8) \end{aligned}$$

In addition, $\Gamma\left(\bar{b}_1\right)=0$ and $\Gamma'\left(\bar{b}_1\right)=-1$. Equation 8 is a linear second-order differential equation with two initial conditions. Therefore, there exists a unique solution for $G^{\epsilon}(.)$ for given $\underline{b}_1$ and $\bar{b}_1$. We can find $\underline{b}_1$ using $G^{\epsilon}\left(\underline{b}_1\right)=0$. To find $\bar{b}_1$, notice that when $y\rightarrow\bar{b}_1$,

$$\int_{\underline{b}_1}^{\bar{b}_1}\int_0^1 \left((1-\epsilon)F(u)+\epsilon\right)(z-u)\,dF(z)\,dG^{\epsilon}(u)=K_1^{\epsilon}. \qquad (9)$$

From equation 9, we can see that $[\underline{b}_1,\mathbf{E}[y|\,[0,1]])$ is an equilibrium support of the first period bid. There is no equilibrium with $\bar{b}_1>\mathbf{E}[y|\,[0,1]]$. For a given support, a unique G^{ϵ} satisfies the indifference condition given by equation 7. There cannot be an equilibrium where the two bidders choose their first period bid from different supports. Thus, this auction has only symmetric equilibria.

The support of G^{ϵ} is closed on the left and open on the right. If it is closed on the right then

$$g^{\epsilon}\left(\bar{b}_1\right)\left(\int_0^{\bar{b}_1}\left(z-\bar{b}_1\right)dF(z)+\frac{\epsilon}{1-\epsilon}\int_{\bar{b}_1}^1\left(z-\bar{b}_1\right)dF(z)\right)=0.$$

That implies $g^{\epsilon}\left(\bar{b}_1\right)=0$ when $\epsilon<\bar{\epsilon}$ for some small $\bar{\epsilon}$. On the other hand, if it is open on the left then if bidder i bids $\underline{b}_1$ instead of $\underline{b}_1+\eta$ for $\eta\rightarrow 0^+$, his expected payoff goes up as his probability of being the high bidder reduces without reducing the information he gets. Therefore the support will be the interval $\left[\underline{b}_1,\bar{b}_1\right)$ where $\underline{b}_1>0$. Hence, G^{ϵ} is an equilibrium distribution which converges to G that solves

$$\frac{f(y)}{F(y)}\int_y^1\int_y^{\bar{b}_1}(z-u)\,dG(u)\,dF(z)+f(y)\int_y^{\bar{b}_1}(u-y)\,dG(u)+g(y)\int_0^y (z-y)\,dF(z)=0.$$

It is obvious that $\left(\sigma^G,\sigma^G\right)$ is an equilibrium of the unperturbed auction. ■

Proof of Proposition 4

Proof. Suppose G^{ϵ} solves the equation 7 for $y \in \left[\underline{b}_1, \bar{b}_1\right)$ and $\psi \in \left(0, \bar{b}_1 - \underline{b}_1\right)$. Define $\overline{G}^{\epsilon}$ to be truncated distribution G^{ϵ} on $\left[\underline{b}_1 + \psi, \bar{b}_1\right)$. Then $\overline{G}^{\epsilon}$ solves equation 7. This implies that $\overline{G}^{\epsilon}$ is an equilibrium distribution of the ϵ-perturbed auction. Thus, there is a continuum of equilibria. ■

Proof of Proposition 5

Proof. In equilibrium,

$$F\left(x_1^s\right)\int_m^{R_2}\left(z - x_1^s\right)dF\left(z\right) + \left(F\left(x_1^s\right) - F\left(m\right)\right)\int_m^{\min\left[v^s, e\left(x_1^s\right)\right]}\left(z - x_1^s\right)dF\left(z\right) = 0.$$

Since $R_2 \geq v$, this implies that $x_1^s > x_1$. Moreover, $x_1^s > x_1$ implies $v^s > v$. Therefore, $v^s > v$ and $x_1^s > x_1$. ■

Proof of Proposition 6

Proof. When $m = 0$ and $r_s = r$, $R_2 = \frac{2}{3}$, $x_1^s = \frac{1}{3}$ and $v^s = \frac{2}{3}$. The seller's expected revenue,

$$\pi_{SRP} = \begin{cases} \frac{r}{9} + \frac{r^2}{3} + \frac{10}{27} & \text{if } r \leq \frac{1}{3} \\ \frac{r}{6} - \frac{7r^2}{6} + \frac{10}{27} & \text{if } r > \frac{1}{3}. \end{cases}$$

If the reserve is public then $x_1 = \frac{1+r+r\sqrt{2}}{2+\sqrt{2}}$ and $v = \frac{2+r\sqrt{2}}{2+\sqrt{2}}$. The seller's expected revenue,

$$\pi_{PRP} = \frac{1-r}{8\left(1+\sqrt{2}\right)^4}\left(50 + 35\sqrt{2} + \left(36 + 26\sqrt{2}\right)r + \left(186 + 131\sqrt{2}\right)r^2\right).$$

For all $r \in [0, R_2]$, $\pi_{SRP} \geq \pi_{PRP}$. ■

References

[1] Ariely, D., G. Loewenstein, and D. Prelec, (2003): “Coherent Arbitrariness: Stable Demand Curves without Stable Preferences,” *Quarterly Journal of Economics*, 118, 73-105.

[2] Ariely, D., A. Ockenfels, and A. E. Roth, (2004): “An experimental Analysis of Ending Rules in Internet Auctions,” Working Paper, Sloan School of Management.

[3] Bajari, P., and A. Hortacsu, (2003): “Winner’s Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions,” *RAND Journal of Economics*, 34, 329-355.

[4] Ergin ,H. (2003): “Costly Contemplation,” Working Paper, MIT.

[5] Hasker, K., R. Gonzalez, and R. C. Sickles (2001): “An Analysis of Strategic Behavior and Consumer Surplus in eBay Auctions,” Working Paper, Rice University.

[6] Hossain, T. (2004): “Essays on Bounded Rationality in Auctions and Bargaining,” Chapter 1, Ph. D. Dissertation, Princeton University.

[7] Hossain, T., and J. Morgan (2003): “A Test of the Revenue Equivalence Theorem Using Field Experiments on eBay,” Working Paper, Princeton University.

[8] Rasmusen, E. (2003): “Strategic Implications of Uncertainty Over One’s Own Private Value in Auctions,” Working Paper, Indiana University.

[9] Milgrom, P., and C. Shannon (1994): “Monotone Comparative Statics,” *Econometrica*, 62, 157-180.

[10] Roth, A. E., and A. Ockenfels (2001): “Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet ,” NBER (Cambridge, MA) Working Paper No. 7729.

[11] Vincent, D., (1990): “Dynamic Auctions,” *Review of Economic Studies*, 57, 49-61.

[12] Wang, T. J., (2003): “Is Last Minute Bidding Bad?,” Working Paper, UCLA.

Table 1: Summary Statistics of Independent Variables

Variable	Mean
Opening price	90.410 [80.89]
Length in days	6.418 [1.72]
Dummy for secret reserve price auctions	0.255 [0.44]
Dummy for seller's ID less than a month old	0.053 [0.22]
Dummy for pictures on item page	0.938 [0.24]
Dummy for brand new clubs	0.268 [0.44]
Dummy for almost new clubs	0.027 [0.16]
Dummy for used clubs	0.314 [0.46]
Dummy for Paypal payments accepted	0.431 [0.50]
Dummy for a head cover sold with the club	0.016 [0.12]
Dummy for left-handed clubs	0.051 [0.22]
Dummy for auctions ending on weekends	0.265 [0.44]
Dummy for seller living in Canada	0.008 [0.09]
Standard deviations are shown inside brackets	

Table 2: Revenue from Public Reserve Price Auctions

Explanatory Variables	Revenue (OLS) (1)	Revenue (OLS) (2)	End Price (Censored Regression) (3)	End Price (Censored Regression) (4)
Opening price	0.051	0.071	-0.337*	-0.332*
	(0.75)	(1.01)	(-7.82)	(-7.59)
Square of opening price	-0.002*	-0.002*	0.002*	0.002*
	(-6.02)	(-6.01)	(7.83)	(7.55)
Share of multiple bidders	121.755*	120.444*	57.728*	57.206*
	(17.24)	(16.95)	(16.15)	(16.26)
Characteristic dummies	Yes	Yes	Yes	Yes
Day Specific Fixed Effects	No	Yes	No	Yes
Constant	Yes	Yes	Yes	Yes
Observations	1509	1509	1509	1509
Mean of the Dependent Variable	112.780	112.780	170.763	170.763
[Standard Deviation]	[81.04]	[81.04]	[34.48]	[34.48]
Adjusted/Pseudo R-squared	0.4730	0.4820	0.0675	0.0761
F-statistics	85.60	24.39	-	-
Likelihood Ratio Chi-squared	-	-	638.33	719.89

Notes:
1. Coefficients significant at 95% level are denoted by * and t/z-statistics are in parentheses.
2. Characteristics dummies are the last ten dummies in Table 2 and the dummies for auction lengths. The “Weekend” dummy was not included with “Day Specific Fixed Effects” (columns (2) & (4)).

Table 3: Censored Normal Regressions of First and Last Bids of All Bidders

Explanatory Variables	First Bid (1)	First Bid (2)	Last Bid (3)	Last Bid (4)
Opening price	0.484*	0.469*	0.481*	0.465*
	(44.49)	(42.81)	(41.98)	(40.25)
Multiple bidder dummy	-6.882*	-7.252*	21.170	20.086*
	(-6.10)	(-6.55)	(17.89)	(17.19)
Share of multiple bidders	0.213	0.787	2.571	2.588
	(0.07)	(0.27)	(0.82)	(0.83)
Secret reserve price auction dummy	2.096	0.799	3.026*	1.804
	(1.75)	(0.64)	(2.43)	(1.40)
Hours remaining in the auction	-0.090*	-0.117*	-0.102*	-0.126*
	(-29.19)	(-32.43)	(-29.34)	(-31.73)
Bidder's feedback rating × 10000	0.442*	0.458*	0.506*	0.527*
	(2.21)	(2.33)	(2.45)	(2.60)
Characteristic dummies	Yes	Yes	Yes	Yes
Day Specific Fixed Effects	No	Yes	No	Yes
Constant	Yes	Yes	Yes	Yes
Observations	9003	9003	9003	9003
Mean of the Dependent Variable	107.341	107.341	117.709	117.709
[Standard Deviation]	[54.00]	[54.00]	[54.09]	[54.09]
Pseudo R-squared	0.0353	0.0389	0.0388	0.0418
Likelihood Ratio Chi-squared	3237.81	3569.02	3394.54	3653.98

Notes:

1. Coefficients significant at 95% level are denoted by * and z-statistics are in parentheses.
2. Characteristics dummies are the last ten dummies in Table 2 and the dummies for auction lengths.
   The "Weekend" dummy was not included with "Day Specific Fixed Effects" (columns (2) and (4)).
3. Independent Variable × K, where K is a constant, means that the coefficients of that variable were multiplied by K for presentational purposes.

Table 4: Probit Analyses of Dummies for a Bid Received and Sniping

Explanatory Variables (Marginal Effects)	Dummy for Sniping (1)	Dummy for Sniping (2)	Dummy for a Bid Received (3)	Dummy for a Bid Received (4)
Opening price	-0.002*	-0.002*	-0.002*	-0.002*
	(-11.27)	(-11.25)	(-17.72)	(-16.99)
Share of multiple bidders	-0.006	-0.004	-	-
	(-0.15)	(-0.10)		
Secret reserve price auction dummy	-0.087*	-0.087*	0.021	0.014
	(-3.95)	(-3.84)	(1.51)	(1.21)
Characteristic dummies	Yes	Yes	Yes	Yes
Day Specific Fixed Effects	No	Yes	No	Yes
Constant	Yes	Yes	Yes	Yes
Observations	2026	2026	2026	2026
Mean of the Dependent Variable	0.235	0.235	0.752	0.752
[Standard Deviation]	[0.42]	[0.42]	[0.43]	[0.43]
Pseudo R-squared	0.0911	0.1125	0.5068	0.5391
Likelihood Ratio Chi-squared	201.34	248.49	1149.70	1223.13

Notes:
1. The marginal effects of the independent variables are reported in this table.
2. Coefficients significant at 95% level are denoted by * and z-statistics are in parentheses.
3. Characteristics dummies are the last ten dummies in Table 2 and the dummies for auction lengths. The “Weekend” dummy was not included with “Day Specific Fixed Effects” (columns (2) and (4)).