RAND Journal of Economics
Vol. 39, No. 2, Summer 2008
pp. 509–529

Learning by bidding

Tanjim Hossain*

We analyze a dynamic second-price auction with an informed bidder and an uninformed bidder who, upon seeing a posted price, learns whether his valuation is above that price. In the essentially unique equilibrium, an informed bidder bids in the first period if her valuation is below some cutoff and bids only in the last period otherwise. An uninformed bidder bids in every period to optimally change the price unless the price is above his valuation or he is the high bidder. This model also provides a rationale behind the use of a secret reserve price in private-value settings.

1. Introduction

■ In most economic models, an agent with complete preferences is assumed to be able to formulate her preferences perfectly. For example, a rational agent usually knows her valuation in a private-value setting. In this article, we introduce nonstandard agents who are not aware of exactly how much they like an object. In many real-life situations, people actually do not need to know their exact valuations for efficient trades to take place. When a person buys milk in the supermarket, he does not need to know exactly how much he values a gallon of milk. He only needs to know whether he values it more than the price. That is, knowing his preferences around the posted price is enough for decision making in this situation. Using this intuition, we introduce a stylized model where, when confronted with a price, an agent who does not know his exact valuation for a good easily learns whether his valuation is above or below the price. We apply this model of preference elicitation in a dynamic second-price auction such as those held on eBay. We show that this model explains the prevalence of late and multiple bidding in an independent private-value (IPV) setting. Moreover, the model provides a rationale behind the use of secret reserve prices in second-price auctions and is consistent with many other stylized facts from online auctions.

An eBay auction is essentially a second-price auction with a fixed closing time where a bidder can submit as many bids as she wants. A bidder's latest bid has to be higher than all of her previous bids and this bid is basically considered as her active bid. At any point during the auction, the current price equals the second-highest bid received so far, but the exact bids are kept undisclosed. At the pre-announced time when the auction ends, the highest bidder wins and pays

*Hong Kong University of Science and Technology; tanjim@ust.hk.
I would like to thank Faruk Gul, Han Hong, Daniel Kahneman, Fahad Khalil, Jeffrey Livingston, Alexandre Mas, Eric Maskin, Shantanu Mukherjee, Nicola Persico, Elie Tamer, Thijs van Rens, seminar participants at various universities, the Editor Philip A. Haile, and two anonymous referees for their valuable comments and suggestions and Matthew Langley for help in creating the data set. I am greatly indebted to John Morgan and, especially, Wolfgang Pesendorfer for their guidance and support. Financial support from the Woodrow Wilson Foundation is gratefully acknowledged.

the bid of the second-highest bidder plus a small bid increment. The winner, hence, does not pay her own bid. For a more detailed description of eBay auctions, see Roth and Ockenfels (2001). It is an empirical regularity that online auctions receive a significant number of bids in the last few minutes of an auction. In this article, we define *sniping* as bidding in the last three minutes of the auction and not placing any earlier bid.[1] Moreover, many bidders frequently update their bids by placing a new bid slightly above the current price if they are not the highest bidder. We refer to this multiple bidding as *nibbling*.[2] Figure 1 presents all the individual bids submitted in an eBay auction of a golf driver. The bidder "dsgn101" submitted her only bid 17 seconds before the end of the auction, or *sniped*. On the other hand, all other bidders placed multiple bids. Each entry denotes submission of a new bid. When bidder "kalanisurf" entered, she saw that the opening price was $9.99 and the highest bidder was "fredburt," but did not know the current highest bid. She took two bids to become the current highest bidder with a bid of $120. Then bidder fredburt came back and placed three separate bids to overtake kalanisurf's bid. Bidders fredburt, kalanisurf, and later bidders "ogus" and "rbwaugh@aol.com" all *nibbled*. Bidder fredburt placed eight separate bids to bid $125. As eBay never raises the price above the bid of the second-highest bidder, placing a single bid of $125 would lead to the same response by other bidders with private values, but would eliminate the risk from a possible failure to place a future bid due to some unforeseen event. The fact that the price in eBay reflects only the second-highest bid makes nibbling, at best, superfluous in a private-value setting. Many of the late bidders also bid earlier in the auction, and this suggests that sniping does not only arise from buyers' desires to know the auction outcomes soon after they place a bid. Section 2 presents further evidence of prevalence of sniping and nibbling and some other anomalous empirical regularities in online auctions.

In Section 3, we introduce a second-price auction with a fixed number of periods where bidders can bid in every period. The current price at the beginning of a period equals the second-highest bid received in the previous periods. At the end of the auction, the bidder who placed the highest bid wins and pays the second-highest bid as the price. In an IPV setting where all bidders know their valuations, or are *informed bidders*, any strategy profile where bidders ultimately bid their private valuations is an outcome-invariant equilibrium. Thus, sniping and nibbling can be consistent with the standard theory. However, all bidders bid their valuations in the first period in the unique equilibrium in weakly undominated strategies if there is a small opportunity cost of waiting. Such an opportunity cost may arise if there is always a small probability that a bidder will not be able to place any bid in the remainder of the auction. In reality, this may correspond to a bidder being unable to go back to an auction because of unanticipated personal commitments, computer malfunctions, or forgetfulness.

The main innovation of this article is introducing a model of bidders who do not know their private valuations for the good exactly but learn more about those during the auction using a boundedly rational learning process. Specifically, when such a bidder is confronted with a minimum price that he has to pay if he wins, either as a posted price or as a minimum bid in an auction, he costlessly learns whether his valuation is above or below that price. As the minimum bid in our dynamic second-price auction can be changed many times, when a bidder does not know his valuation, he can experiment strategically by using his bids and learn more about his preferences. A point to note is that we do not assume that an *uninformed bidder*, a bidder who does not know his valuation, has to change the price with his own bids to get more information about his valuation. He gets more information any time the price changes whether the change results from his own bid or some other bidder's bid. New prices give him new information and this incentive leads him to place multiple bids. On the other hand, in some cases, it may be optimal for an informed bidder to restrict such *learning by bidding* by the uninformed bidder. This leads to late bidding or sniping by her.

[1] Traditionally in the literature, all late bids are considered sniping bids.

[2] We thank John Morgan for suggesting this term.

FIGURE 1

BID HISTORY FROM A TYPICAL eBAY AUCTION

eBayBid History for

TITLEIST 975J-VS 8.5 GRAFALLOY PRO-LITE STIFF (Item # 3603774740)

Currently	**$167.50**	First bid	**$9.99**
Quantity	**1**	# of bids	**21**
Time left	**Auction has ended.**		
Started	Apr-12-03 11:28:25 PDT		
Ends	Apr-19-03 11:28:25 PDT		
Seller (Rating)	**shop4golf (15)**		

Bidding History (Highest bids first)

User ID	Bid Amount	Date of Bid
dsgn101 (68)	$167.50	Apr-19-03 11:28:08 PDT
rbwaugh@aol.com (1)	$165.00	Apr-19-03 10:00:18 PDT
ogus (16)	$162.70	Apr-19-03 10:26:35 PDT
ogus (16)	$157.69	Apr-19-03 10:26:17 PDT
ogus (16)	$151.69	Apr-19-03 09:40:43 PDT
rbwaugh@aol.com (1)	$150.00	Apr-18-03 21:17:15 PDT
ogus (16)	$145.69	Apr-19-03 09:40:28 PDT
ogus (16)	$140.00	Apr-19-03 09:39:23 PDT
ogus (16)	$135.00	Apr-19-03 09:39:09 PDT
kalanisurf (1)	$130.00	Apr-18-03 14:46:21 PDT
fredburt (0)	$125.00	Apr-18-03 11:59:52 PDT
kalanisurf (1)	$125.00	Apr-18-03 14:40:36 PDT
kalanisurf (1)	$120.00	Apr-18-03 06:51:27 PDT
fredburt (0)	$120.00	Apr-18-03 11:59:42 PDT
fredburt (0)	$115.00	Apr-18-03 11:59:30 PDT
fredburt (0)	$110.00	Apr-15-03 15:33:38 PDT
kalanisurf (1)	$105.00	Apr-18-03 06:50:44 PDT
fredburt (0)	$100.00	Apr-14-03 15:06:35 PDT
fredburt (0)	$99.00	Apr-14-03 15:06:26 PDT
fredburt (0)	$80.00	Apr-14-03 15:06:16 PDT
fredburt (0)	$60.00	Apr-14-03 15:05:52 PDT

Remember that earlier bids of the same amount take precedence.

Bid Retraction and Cancellation History

There are no bid retractions or cancellations.

Theorem 1 shows that, when both bidders bid in the first period of the auction with positive probability, there essentially is a unique equilibrium 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 $v_2 < p_t$.

This result illustrates that sniping and nibbling are equilibrium phenomena in a private-value setting. Some bidders will bid only early in an auction, some will bid both early and late, and some will bid only late. Moreover, both early and late bidding may occur in the same equilibrium. A significant fraction of bidders who bid toward the end of an auction will not bid earlier in that auction. A bidder who places multiple bids may raise her bid frequently until she becomes the high bidder. Occurrence of sniping and nibbling in equilibrium is robust to many modifications of the original learning model. This model explains many other stylized facts presented in Section 2 quite well. We then show, in Proposition 4, that when an uninformed bidder can also learn by comparing their valuations to hypothetical prices or prices in other auctions before the auction in which he participates, sniping and nibbling will still occur.

Sniping and nibbling with uninformed bidders is not an artifact of having exactly one informed bidder and one uninformed bidder. We concentrate on this case because of the simplicity of the model. In Section 4, we extend the model to auctions with arbitrary numbers of bidders. We show that sniping and nibbling occurs in any auction with at least one informed bidder and at least one uninformed bidder.

Learning by bidding may also explain the extensive use of a secret reserve price in second-price auctions. If all bidders are informed, secret or public reserve prices lead to the same outcome in an IPV setting. On the other hand, an uninformed bidder may win the object even when his valuation is below the reserve price if the reserve is secret but not if it is public. Proposition 7 shows that a secret reserve price auction may generate higher expected revenue than a public reserve price auction.

The rest of the article is organized as follows: the next section reports some stylized facts that led to the questions addressed in this article and relates this article to the existing literature. Section 3 introduces the theoretical model with two bidders where one is informed and the other is uninformed. Section 4 discusses auctions with more than two bidders. Section 5 analyzes secret reserve price auctions and Section 6 concludes the article. All proofs are in the Appendix.

2. Some stylized facts and relation to the literature

■ We collected a data set of 2,026 completed auctions of " Titleist 975J" golf drivers conducted on eBay between February and April of 2003. In these auctions, 9,003 bidders placed 17,057 separate bids, on average placing 1.89 bids each. About 40% of all bidders placed more than one bid or *nibbled*. Of all bidders, 9.6% placed a bid in the last three minutes and 67% of these bidders did not place any bid earlier in the auction. Some bidders bid early in the auction, some bid late, and others bid both early and late.

Many studies found the same bid pattern in eBay auctions for a wide variety of goods. Roth and Ockenfels (2001) gathered a sample of over 1,000 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. In a study of coin auctions on eBay by Bajari and Hortacsu (2003), 32% of the bids were submitted after 97% of the duration of the auction had passed. Hossain and Morgan (2006) sold brand new popular music CDs and Xbox game cartridges on eBay. There was little uncertainty about the quality or popularity of these products; still, at least one bidder placed multiple bids in 76% of the auctions. Bids were placed in the last five minutes in 30% of the auctions. Ariely et al. (2005) found evidence of extensive sniping and nibbling in laboratory experiments of eBay-type auctions. We report some additional empirical regularities or stylized facts found in the existing literature or in our data set.

(i) Many late bidders on eBay bid *only* in the last few minutes of an auction. In our data set, more than two thirds of the bidders who placed a bid in the last three minutes of an auction did not bid earlier in that auction. A third of the late bidders placed bids both early and late.

(ii) Many bidders *self-nibble*; that is, they bid repeatedly below the highest outstanding bid (the exact outstanding bid is unknown to them) while the highest bidder stays unchanged. Moreover, many bidders place a new bid every time they are displaced from the highest bidder position. Almost 80% of all nibblers in our data set placed consecutive bids while they were not the high bidder. A nibbler placed more than 3.26 bids on average.

(iii) Roth and Ockenfels (2001) find that more experienced bidders are more likely to snipe than less experienced bidders. One may argue that experienced bidders are more proficient in thinking counterfactually, asking themselves repeatedly how much they would be willing to bid, and are thus more likely to be informed bidders. In our data set, a bidder with a feedback rating of at least 20 was one-third less likely to place multiple bids than a bidder with a lower feedback rating.

(iv) Now-defunct *Amazon.com* had a "soft" closing time: the length of the auction would be increased if there was a bid within the last ten minutes of the pre-announced closing time. After that, the auction would close when there had not been any activity for ten minutes. Roth and Ockenfels (2001) find that nibbling was common but sniping was relatively uncommon in Amazon auctions.

(v) Controlling for auction characteristics, an auction where all bidders placed only one bid each received almost 17% less revenue than an auction where all bidders placed multiple bids in our data set. According to our model, informed bidders are likely to bid only once and are also more likely to snipe. On the other hand, uninformed bidders are likely to bid multiple times. Thus, multiple bidding can be an indicator of being uninformed. However, whether a bidder places multiple bids depends on the opening price leading to an endogeneity problem in the regressions. Combined with the model's predictions, the third stylized fact suggests that an experienced bidder is more likely to be informed of her own private valuation. Therefore, we use bidder experience, specifically the share of bidders with feedback rating above 20, as an instrument for the share of informed bidders.[3]

(vi) In our data set, nibblers started with a lower bid than bidders who placed only one bid, but finally ended up bidding higher in the auction. The first bids of bidders who placed multiple bids were lower than those by bidders who place only one bid by $7 and their last bids were higher by $24. We use a dummy for bidders with feedback rating above 20 as an instrument for informed bidders.

(vii) Bajari and Hortacsu (2003) find 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. When all bidders are informed, keeping the reserve price secret or public leads to the same final outcome, but the seller has to pay an extra fee for keeping the reserve secret on eBay.

(viii) A secret reserve price auction was about one-third less likely to receive bids in the last few minutes by bidders who did not place any earlier bid than an auction with the same opening price but no secret reserve. The first bids of bidders who placed multiple bids were higher by $5 if there was a secret reserve price.[4]

[3] An underlying assumption here is a bidder's willingness to pay is uncorrelated with her experience on eBay.

[4] The relevant coefficients in regularities (v), (vi), and (viii) are significant at the 1% significance level.

As sniping is difficult to explain in a standard private-value second-price auction, many have suggested alternative models. In one of the first papers on online auctions, Roth and Ockenfels (2001) present a two-stage second-price auction where a continuous time auction is followed by a discrete last period. The object for sale is 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 or snipe. 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. Bidders have no incentive to nibble in this model. Hasker et al. (2004) reject the hypothesis that bidders follow "snipe-or-war" strategies using a data set of color computer monitors and Ariely et al. (2005) reject this hypothesis using laboratory experiments.

Some have used common-value objects to explain sniping. In another pioneering paper on eBay auctions, Bajari and Hortacsu (2003) analyze an auction for common-value objects with a structure similar to that of the Roth and Ockenfels (2001) model. However, bids placed in the last period are accepted with certainty. When the object is common value, then there is a unique symmetric Nash equilibrium where all bidders bid only in the last period. Hence, there will be sniping but there will be no nibbling.

Like in the current model, the model in Rasmusen (2006) has an informed and an uninformed bidder. The uninformed bidder can learn his private valuation *perfectly* during the auction by paying a cost. For certain cost levels, there is an equilibrium where the informed bidder places a bid below her valuation early in the auction using a mixed strategy to hide her true valuation and then bids her valuation late in the auction. In that equilibrium, no bidder bids more than twice and both bidders have incentives to place multiple bids. A sniper always places an earlier bid in this model, contrary to the first stylized fact. The uninformed bidder will not self-nibble or bid repeatedly while the identity of the high bidder does not change. Once the uninformed bidder's bid is overtaken by some other bidder in that model, he will expend the cost to learn his valuation and place at most one more bid. This contradicts the second stylized fact.

In Compte and Jehiel (2004), bidders also do not know their own valuations and get imperfect information during the auction. However, bidders receive information exogenously during the auction. In our model, the uninformed bidder optimally chooses what kind of information to get by strategic experimentation (his bid determines at what price he gets his next information). Moreover, bid activity by some bidders is needed in our model for the uninformed bidder to get more information about his valuation, as he gets new information only when the price changes.

On eBay, many concurrent auctions of identical or very similar goods go on and new auctions of similar products start every day. 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 five or six bids within a minute on the same auction without placing bids in any other auctions, a pattern common in eBay auctions. More importantly, as Ely and Hossain (2006) show, concurrent auctions actually make sniping less profitable than bidding early when all bidders are rational. In Ariely et al.'s (2005) experiments, sniping and nibbling occurred even though there was no concurrent or future auctions of identical or similar goods. Thus, a retail market and concurrent and future auctions do not explain the pervasiveness of sniping in online auctions.

Song (2004) estimates an empirical model of eBay auctions with unknown number of potential bidders using semi-nonparametric techniques. That model allows for both sniping and nibbling. In this article, we use simple reduced-form estimations instead of semi-nonparametric estimation. The main difference between the two models is that the current article provides a strict motivation behind sniping and nibbling behavior by bidders instead of these merely being an artifact of one of many outcome-invariant equilibria.

The learning model proposed in this article can be related to the 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 he has to pay at least that amount to get the good. It facilitates the need to know whether his valuation is above or below the price and makes this comparison cognitively easy. Thus, uninformed bidders can also be viewed as individuals who learn their preferences by contemplating about their tastes at some cost, as suggested by Ergin (2003). An uninformed bidder can contemplate the relation between his valuation and the posted price costlessly and a similar contemplation for any other hypothetical price is infinitely costly.

We can also draw some motivation for learning by comparing valuation with the price from psychology literature. More than five decades ago, Herbert Simon (1956, 1957) pointed out that evolutionary pressures guide individuals to local ("better than") rather than global ("best possible") optimization. Asking the self whether the good is better than some price is similar to Simon's local optimization. Recently, psychologists have adopted the "two-systems" perspective on human decision making (Sloman, 2002). According to that theory, judgments are made by the interaction of two mental systems. The first one, known as system 1, is analogous to intuition. The other one, known as system 2, is analogous to reason. Comparing a price and the valuation can be thought of as judgment using system 1 and knowing the exact valuation as judgment using system 2. Uninformed bidders can be viewed as agents with extremely high cost of using system 2. Informed bidders can use system 2 without any cost to learn their valuations exactly.

3. The main 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 $b_t^i \in [0, 1]$. Bidder i's highest bid up to period t is denoted by $\beta_t^i = \max_{\tau \le t} b_\tau^i$. The auction starts at an opening or period 1 price of $p_1 = m \in [0, 1)$. In period t, the second-highest bid up to period $t - 1$ is posted as the current price p_t. The high bidder is denoted by $w_t \in \{0, 1, 2\}$. If none of the bidders places a bid as great as m in periods 1 to $t - 1$, then $w_t = 0$ and $p_t = m$. Otherwise, $w_t = 2$ and $p_t = \max\{m, \beta_{t-1}^1\}$ if $\beta_{t-1}^1 < \beta_{t-1}^2$ and $w_t = 1$ and $p_t = \max\{m, \beta_{t-1}^2\}$ if $\beta_{t-1}^1 \ge \beta_{t-1}^2$. The highest bidder wins and pays the second-highest bid at the end of the auction. If none of the bidders places a bid as great as m in periods 1 to T, then the object stays unsold. Otherwise, if $\beta_T^1 \ge \beta_T^2$ then bidder 1 wins and pays $\{m, \beta_T^2\}$ and if $\beta_T^2 > \beta_T^1$ then bidder 2 wins and pays $\{m, \beta_T^1\}$.

A bidder never observes the other bidder's exact bids. The public history at the beginning of period t, h_t, is the sequence $\{(p_\tau, w_\tau)\}_{\tau=1}^t$. 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 $\sigma_t : H_t \times [0, 1] \to [0, 1]$. We restrict our attention to pure strategies in this section and Section 5. An *equilibrium* is defined as a Perfect Bayesian Equilibrium where none of the bidders plays a weakly dominated strategy.

If all bidders know their private valuations, then any strategy profile where each bidder ultimately bids her private valuation is an equilibrium. All these equilibria lead to the same outcome as the equilibrium where all bidders bid their valuations at the beginning of the auction. However, if there is a small probability that a bidder will not be able to place any bid in the remainder of the auction, then a bidder bids her valuation in the first period in the unique equilibrium in weakly undominated strategies.

In this article, we assume that 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.

At the beginning of period $t \in \{1, \ldots, T\}$, the uninformed bidder 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 $v_2 \geq p_t$ whereas a negative signal implies the opposite. Because p_t is nondecreasing in t, once the uninformed bidder receives a negative signal, all his future signals will be negative. Receiving a signal, bidder 2 updates his belief about v_2 according to Bayes' rule. Then, in any weakly undominated strategy, bidder 1 bids v_1 in one of the periods and never bids above v_1 and bidder 2 bids the conditional expected value of v_2 in the last period if v_2 is above p_T. To provide the intuition behind the main results from this model, we first analyze a simple example in the next subsection.

□ **A simple example.** Suppose there are two periods and the valuations of the two bidders are independently drawn from uniform distributions on [0,1]. The opening price m is zero. Because p_2 equals the lower of the two players' first-period bids, p_2 equals zero if one of the players does not bid in the first period. Informed bidder 1 knows her valuation v_1. Uninformed bidder 2 knows that $v_2 \in [0, 1]$ and learns whether $v_2 \geq p_2$ at the beginning of the second period.

In weakly undominated strategies, bidder 1 bids v_1 in one of the periods and bidder 2 bids the conditional expected value of v_2 in the last period. The key point is that, using her first-period bid, bidder 1 can control the possible values of the conditional expected value of v_2 given p_2. If bidder 1 does not place any bid in the first period, then bidder 2 only knows that $v_2 \in [0, 1]$ even in the second period and bids $\frac{1}{2}$ in that period. If bidder 1 places a bid in the first period, the price in the second period is strictly positive if bidder 2 places a first-period bid. If $v_2 < p_2$ then the conditional expected value of v_2 is below p_2 and bidder 2 does not place any bid in the second period. In that case, bidder 2's final bid is lower than what it would be if bidder 1 did not place a bid in the first period. However, if $v_2 \geq p_2$ then bidder 2's final bid is higher than what it would be if bidder 1 did not place a bid in the first period. As a result, for low values of v_1, it is optimal for bidder 1 to bid v_1 in the first period in the hope that bidder 2's valuation is below p_2 so that he does not bid in the second period. For higher values of v_1, it is optimal for bidder 1 only to bid in the last period, winning the object for sure at a price of $\frac{1}{2}$.

On the other hand, a first-period bid helps bidder 2 to get more information about v_2 from p_2. However, a large bid increases the probability of him bidding above v_2. He chooses his first-period bid to balance this tradeoff and places a bid smaller than $\frac{1}{2}$, the unconditional expected value of v_2. This pair of strategies of the two bidders essentially characterizes all equilibria of this game.

Now we characterize bidder 2's first-period bid and a cutoff value such that if v_1 is below that cutoff then bidder 1 bids v_1 in the first period and snipes otherwise. We will 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.
Theorem 1 shows that comparing these two strategies ensures that we find an equilibrium.

Suppose bidder 2 bids x in period 1 and bids $\frac{1+p_2}{2}$ in period 2 if $v_2 \geq p_2$. If $v_1 \leq x$ then bidder 1 gets zero payoff from both strategies. If $v_1 \in (x, \frac{1}{2}]$ then p_2 equals x if bidder 1 follows strategy 1. With probability x, bidder 2 gets a negative signal in period 2. A negative signal implies $v_2 < x$ and he does not place any more bids. If bidder 2 gets a positive signal, then he bids $\frac{1+x}{2}$ and bidder 1 loses the auction. Hence, bidder 1 gets an expected payoff of $x(v_1 - x)$ from strategy 1. She gets an expected payoff of zero from strategy 2 as $v_1 \leq \frac{1}{2}$. Therefore, following strategy 1 is a best response for bidder 1 for any $x \geq 0$. If $v_1 \in (\frac{1}{2}, \frac{1+x}{2}]$ then bidder 1 gets an expected payoff of $x(v_1 - x)$ from strategy 1. She gets an expected payoff of $v_1 - \frac{1}{2}$ from strategy 2. Bidder 1's best response is strategy 1 if $v_1 \leq v$ and strategy 2 if $v_1 > v$ where $x(v - x) = v - \frac{1}{2}$; that is, $v = \frac{\frac{1}{2} - x^2}{1 - x}$. If $v_1 \geq \frac{1+x}{2}$ 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 $v_1 - \frac{1}{2} - \frac{x^2}{2}$. Her expected payoff from strategy 2 is $v_1 - \frac{1}{2}$. Thus, her best response is strategy 2 for all $v_1 \geq v$. Given player 1's strategy, bidder 2's expected payoff from bidding x in the first period is $\int_x^1 (\int_0^v (z - y)\, dy)\, dz$. He maximizes his expected payoff by choosing $x = \frac{v}{2}$. Solving for v and x, we get

$$x = \frac{1}{2 + \sqrt{2}} \text{ and } v = \frac{2}{2 + \sqrt{2}}.$$

In this equilibrium, bidder 1 bids either in period 1 or 2. In contrast, bidder 2 may bid in both periods. Thus, bidder 2 nibbles to learn about v_2 and this nibbling leads to sniping by bidder 1. There exists an analogous equilibrium for any arbitrary number of periods T and a general distribution F from which v_1 and v_2 are drawn.

□ **The main result.** Now we formally present the above result in a general model. Theorem 1 shows that the winner and final price pair of any equilibrium where both bidders place a bid in the first period with positive probability is the same as that of a specific equilibrium (σ^v, σ^x). Thus (σ^v, σ^x) characterizes all the equilibria and the equilibrium outcome is unique. In this equilibrium, bidder 1 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. Bidder 2 places bids in every period as long as he is not the high bidder and v_2 is above the current price. Hence, bidder 1 snipes with positive probability and bidder 2 nibbles.

Here σ^v is such that

$$\sigma_t^v(h_t, v_1) = \begin{cases} v_1 \text{ for all } t & \text{if } v_1 \leq v \\ 0 \text{ for } t < T \text{ and } v_1 \text{ for } t = T & \text{if } v_1 > v. \end{cases}$$

If bidder 2 gets a negative signal in period t, he does not bid anymore as the price is already higher than v_2. 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 $v_2 \geq p_t$ 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. Strategy σ^x is such that

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

For notational convenience, we use $\mathbf{E}[Q | [a, c]] = \frac{\int_a^c Q(y)\, dF(y)}{F(c) - F(a)}$ for $a, c \in [0, 1]$.

If (σ^v, σ^x) is an equilibrium strategy profile, then we refer to (v, x) as an *equilibrium cutoff-action* pair. Noting $x_0 = m$ and $x_T = \mathbf{E}[y | [x_{T-1}, 1]]$, we use the following two equations to characterize the unique equilibrium cutoff-action pair of this game:

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

$$v = \frac{\int_m^1 y\, dF(y) - \sum_{t=1}^{T-1} x_t (F(x_t) - F(x_{t-1}))}{1 - F(x_{T-1})}. \tag{2}$$

Theorem 1. There exists a unique equilibrium cutoff-action pair (v, x) where (v, x) satisfies equations 1 and 2. Any equilibrium where both bidders place a bid in the first period with positive probability leads to the same outcome as that of the equilibrium (σ^v, σ^x).

In this dynamic auction, bidder 2 gets multiple opportunities to learn about v_2 by changing the price with his bids. When choosing a bid, he faces the tradeoff 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 when he is not the high bidder reflect optimal experimentation by him. Bidder 1's equilibrium behavior is a strategic response to this learning. Because 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 this game's essentially unique equilibrium, 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.

The equilibrium cutoff-action pair is unique. There is a unique cutoff $v < 1$ for which there exists an action sequence x such that (v, x) leads to an equilibrium. Moreover, the optimal x given this v is unique. Note that we are not restricting bidder 1's strategy to bidding v_1 whenever she bids above m. However, any equilibrium where both players place a bid in the first period with positive probability leads to the same outcome as the outcome from (σ^v, σ^x). For example, if bidder 1 follows a strategy where $\beta_t^1 \geq x_t$ for all t such that $v_1 \geq x_t$ whenever $b_1^1 > m$ leads to the same information structure for bidder 2 if bidder 1 followed the strategy that $b_1^1 = v_1$ if and only if $b_1^1 > m$. Then the public history after all periods will be the same as that if bidder 1 followed σ^v and bidder 2's best response can be characterized by σ^x, as the equilibrium cutoff-action pair is unique. All these equilibria are essentially the same as (σ^v, σ^x). There is no equilibrium where bidder 1 places a bid in the first period and bidder 2 becomes the high bidder in period t even though $v_1 \geq x_{t-1}$ occurs on the equilibrium path. Moreover, there is no equilibrium where bidder 1 bids for the first time in period $\tau \in \{2, 3, \ldots, T-1\}$ for some values of v_1. As a result, any equilibrium strategy can be characterized by σ^v and σ^x, thus leading to a unique equilibrium outcome.

Theorem 1 restricts attention to strategy profiles where both bidders place bids in the first period with positive probability. If we remove this restriction, there will be equilibria where bidder 1 places no bid in the first k periods irrespective of what v_1 is and then both players play an equilibrium using the equilibrium cutoff-action pair of a $T - k$ period auction. Nevertheless, if there is a small probability that a bidder will not be able to place any bid in the remainder of the auction, then all these equilibria vanish and we have a unique equilibrium outcome in weakly undominated pure strategies. Theorem 3 in Hossain (2004) shows this result formally.

An informed bidder has incentives to bid only early or only late and an uninformed bidder bids both early and late. Thus, the model is consistent with the first stylized fact that some bidders bid early, some bid late, and some bid both early and late. An uninformed bidder nibbles repeatedly while he is not the high bidder to get more information about his valuation. This explains self-nibbling as evidenced in the second stylized fact. If experienced bidders are more likely to be informed then the model suggests that experienced bidders are more likely to snipe and less likely to nibble, as found in the third stylized fact. If we incorporate Amazon's "soft" closing time in our model, there will not be any benefit from sniping as bidders always get time to react to any bid. Costless comparison of the posted price and valuation will still give rise to nibbling, as evidenced in the fourth stylized fact.

Equilibrium characteristics and some variations. 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:

$$v = \frac{T + m\sqrt{T}}{T + \sqrt{T}} = 1 - \frac{1 - m}{\sqrt{T} + 1}$$

$$x = \left\{ \frac{t + (T - t)m + m\sqrt{T}}{T + \sqrt{T}} \right\}_{t=1}^{T-1}$$

$$\Rightarrow x_t - x_{t-1} = \frac{1 - m}{T + \sqrt{T}}.$$

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 auction is the same as that of a model where both bidders are informed 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 bidder with higher valuation wins and pays the second-highest valuation with probability 1 when T approaches infinity.

Because 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. If bidder 1 could induce perfect learning by bidding in period 1, bidder 1 would bid in period 1 for any v_1. Hence, as T approaches infinity, the probability of sniping approaches zero.

An equilibrium characteristic is that bidder 2, when not the high bidder, bids in every period unless he gets a negative signal implying that he has bid above v_2. If bidder 1 bids in period 1 and wins, she pays a price above v_2. As a result, v approaches one only when the expected overpayment approaches zero, which happens only if T approaches infinity. The necessary and sufficient conditions for $v \to 1$ are: (i) $x_t - x_{t-1} \to 0$ for all t and (ii) $x_{T-1} \to 1$. At the limit when T approaches infinity, with probability 1, the bidder with the higher v_i wins and pays a price of min $\{v_1, v_2\}$.

Bidder 2's 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 places only one bid equaling v_1. Thus, this model is consistent with the sixth stylized fact. The seller gets a higher expected revenue when bidder 2 is uninformed than the case when both players are informed, as an uninformed bidder's last bid overshoots v_2 on average. The fifth stylized fact that auctions with a higher share of bidders who place multiple bids receive higher revenue is consistent with the overbidding by uninformed bidders predicted by this model.

Proposition 2. The expected revenue with one informed and one uninformed bidder is higher than that with two informed bidders. This revenue difference goes to zero as T approaches infinity.

Distribution of the final price converges to that of the benchmark model as T goes to infinity and the revenue difference goes to zero. This result strictly depends on the assumption that the uninformed bidder costlessly learns whether $v_2 \geq p_t$ in every period. We consider two alternative models where the probability of sniping does not go to zero as T approaches infinity.

First suppose bidder 2 cannot learn the relation between his valuation and the price in every period. Rather, at the beginning of any period t, he learns whether $v_2 \geq p_t$ with probability $\frac{\lambda}{T-1}$ and he gets no new information about v_2 with probability $1 - \frac{\lambda}{T-1}$. The signals are still binary but the signal arrival process is stochastic following a binomial process. Equilibrium bidding behavior, discussed earlier in this section, is robust to such uncertain arrival of signals. Theorem 1 holds even when signals do not arrive in every period—any equilibrium can be characterized by a unique equilibrium cutoff-action pair. Please see Section 4 in Hossain (2004) for a detailed discussion of the stochastic signal arrival case. Sniping and nibbling occur in any equilibrium and bidder 2's final bid, on average, is above v_2. When signals arrive stochastically, for any finite 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. Bidder 2 gets finitely many opportunities to compare his bid to the current price in expectation. Hence, he places large enough bid increments, compared to the price when he received his last signal, so that his next signal is informative. This creates friction in learning even when T approaches infinity and bidder 1 snipes with nonzero probability.

Next, 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. 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, $m = 0$, $\kappa = 0.005$, and $T \geq 3$. 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 is an equilibrium where bidder 2 learns whether $v_2 \geq p_t$ at most three times and bidder 1 snipes if v_1 is greater than $\frac{3}{3+\sqrt{3}}$. However, there is no equilibrium where bidder 2 learns whether $v_2 \geq p_t$ more than thrice.

Another potentially interesting variation of the model could be that bidder 1 does not know whether bidder 2 is informed or uninformed but only knows the probability of bidder 2 being uninformed. If there is a small possibility that bidder 2 is uninformed then sniping occurs in equilibrium. Suppose bidder 1 only knows that bidder 2 is uninformed with probability $\alpha \in (0, 1)$ and is informed with probability $1 - \alpha$. 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 from Theorem 1. Then (v, x) leads to an equilibrium for all $\alpha \in (0, 1]$.

Proposition 3. The equilibrium cutoff-action pair is the same for any $\alpha \in (0, 1]$. The expected revenue is increasing in α.

If bidder 2 is informed, bidder 1 gets the same expected utility from sniping or not. For all $\alpha \in (0, 1]$, the optimal cutoff for a given x is the same. Similarly, for a given v, the optimal x is the same. Hence, the same (v, x) is the equilibrium cutoff-action pair 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.

□ **Learning by contemplation.** We now analyze the model when the uninformed bidder can learn by contemplating about his preferences around other values in addition to the current price of the auction in which he is participating. Suppose, in addition to comparing his valuation with the price, the uninformed bidder can write down a hypothetical price b on a piece of paper and then can learn whether $v_2 \geq b$ if he contemplates hard enough. Of course, if he can contemplate like this infinitely many times, he will learn v_2 perfectly and become informed. Although this kind of learning makes sniping less effective by reducing bidder 2's dependence on learning by bidding, sniping and nibbling will still occur in this game. The uninformed bidder will nibble, as a new price gives him another benchmark to compare v_2 with. Nibbling leads to some friction in his learning of v_2 and he will overbid on average. This, in turn, will lead to sniping by bidder 1 if v_1 is high enough.

This will be clear in a simple two-period auction where bidder 2 can costlessly learn whether $v_2 \geq b$ for any b before the auction begins. As before, he can also learn whether $v_2 \geq p_t$ during the auction. Proposition 4 shows that sniping and nibbling occur when the uninformed bidder can learn by contemplating this way. This result will still go through if he can contemplate around finitely many hypothetical prices before the auction or the auction length is more than two periods.

Proposition 4. In the equilibrium, bidder 2 will contemplate whether $v_2 \geq b$ and then bid x_l if $v_2 < b$ and x_h otherwise in period 1. He will bid the conditional expected value of v_2 given his information in period 2. Bidder 1 will bid v_1 in the first period if $v_1 \leq v$ and will snipe otherwise. Here, $x_l = \mathbf{E}\,[y \mid [m, \mathbf{E}\,[y \mid [x_l, b]]]]$, $x_h = \mathbf{E}\,[y \mid [b, v]]$, $b = \mathbf{E}\,[y \mid [\mathbf{E}\,[y \mid [x_l, b]], x_h]]$, and

$$v = \frac{\int_m^1 y\,dF(y) - x_l\,(F\,(x_l) - F(m)) - \int_{x_l}^b y\,dF(y) - x_h\,(F\,(x_h) - F(b))}{1 - F\,(x_h)}.$$

For an example, consider uniform F with $m = 0$ and $T = 2$. In equilibrium, bidder 2 chooses $b = 0.45$. If $v_2 < b$ then he bids 0.15 in period 1 and if $v_2 \geq b$ then he bids 0.59 in period 1. In period 2, he bids his conditional expected value. Bidder 1 bids v_1 in period 1 if $v_1 \leq 0.74$ and snipes otherwise.

4. Arbitrary number of bidders

■ This section analyzes auctions with more than two bidders. There are some fundamental differences between auctions with one uninformed bidder and more than one uninformed bidder. To illustrate this, we first analyze an auction with two uninformed bidders and no informed bidder. Next we analyze a game with at least one informed bidder and at least two uninformed bidders. The dynamics of an auction with many informed bidders and one uninformed bidder is similar to that of an auction with one informed and one uninformed bidder. We restrict attention to two-period auctions in this section.[5] For simplicity, we assume $m = 0$. The results can be extended to any opening price.

Although sniping does not occur in the game with just two uninformed bidders, it has some interesting equilibrium 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. In this game, bidders 1 and 2 are uninformed and they get signals in every period. There is no equilibrium 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 follow mixed strategies. To allow for mixed strategies, we redefine strategy $\sigma = \{\sigma_t\}_{t=1}^{T-1}$ where $\sigma_t : H_t \times [0, 1] \to \Delta\,([0, 1])$ and $\Delta(X)$ denotes the set of all probability measures over X.

Let strategy σ^G be such that bidder *i* chooses his period 1 bid from $[\underline{b}_1, \overline{b}_1)$ according to the distribution G for some positive $\underline{b}_1$. He bids $\mathbf{E}\,[y \mid [p_2, 1]]$ in period 2 if $v_i \geq p_2$. If (σ^G, σ^G) 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 $[\underline{b}_1, \overline{b}_1)$. In any equilibrium, both players choose their bids from the same support and we also can show that they would choose their bids using the same distribution. Thus, all equilibria of this game are symmetric. Theorem 2 and Proposition 5 can be extended to a game with n_U uninformed bidders for any $n_U \geq 2$.

Theorem 2. There exists a distribution G such that (σ^G, σ^G) is an equilibrium.

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

[5] Not all results will go through in an auction with $T > 2$.

Proposition 5. If G is an equilibrium distribution, then so is $\overline{G}$.

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 $y \in [\underline{b}_1 + \psi, \overline{b}_1)$ 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 goes 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.

In the presence of both uninformed and informed bidders, there will be both sniping and nibbling. Suppose there are $n_I + n_U$ bidders. Bidders 1 to n_I are informed and bidders $n_I + 1$ to $n_I + n_U$ are uninformed with $n_I \geq 1$ and $n_U \geq 2$.[6] Uninformed bidder u learns whether $v_u \geq p_t$ in both periods. There is a symmetric equilibrium where an informed bidder i follows strategy σ^v where she bids v_i in period 1 if $v_i \leq v$ and bids nothing in period 1 and v_i in period 2 if $v_i > v$; that is, she snipes if $v_i > v$. Strategy σ^G for uninformed bidder j is such that his period 1 bid is from $[\underline{b}_1, \overline{b}_1]$ according to the distribution G . In period 2, he bids $\mathbf{E}\,[y \mid [p_2, 1]]$ if $v_j \geq p_2$. The pair (v, G) is an *equilibrium cutoff-distribution pair* if informed bidders following σ^v and uninformed bidders following σ^G is an equilibrium.

Theorem 3. There exists an equilibrium cutoff-distribution pair (v, G) and for any equilibrium cutoff-distribution pair, $v < 1$.

Although the price may change even if informed bidder i does not bid in period 1, p_2 is likely to be higher if she bids in the first period. She snipes when her valuation is high enough, so that giving uninformed bidders less information is optimal. The cutoff value is strictly below 1. The distribution G makes both uninformed bidders indifferent between any bid in $[\underline{b}_1, \overline{b}_1]$. Informed bidders play pure strategy and uninformed bidders play mixed strategies in this equilibrium. There is no equilibrium where an uninformed bidder follows a pure strategy in period 1.

5. Secret reserve price auctions

■ In online auctions, the opening price serves as the public reserve price. In eBay auctions, the seller can set a secret reserve price in addition to the opening price by paying an extra fee. The seller announces that there is a secret reserve 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.

In a standard second-price auction in the IPV setting, all bidders with valuation above the public reserve bid their valuations. A bid below the secret reserve is ignored in determining the auction outcome. Thus, a secret reserve price auction has the same outcome as that of a public reserve price auction with the same reserve. Hence, paying a fee for a secret reserve price is suboptimal for the seller. With uninformed bidders, however, secret reserve auctions may lead to a different outcome than a public reserve auction.

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, r_s, is received, that bid becomes the current price. If $\max_i \beta^i_{t-1} \leq r_s$ then $p_t = \max_i \beta^i_{t-1}$ and otherwise p_t equals the second highest of β^1_{t-1} and β^2_{t-1}. Thus, p_t can change even when only one bidder places bids. The object stays unsold if $\max_i \beta^i_T < r_s$. In a public reserve auction, there is only an opening price m and in a secret reserve auction there is an opening price m_s and a secret reserve r_s that is unknown to both bidders. We assume that both bidders believe that r_s is drawn according to density $\frac{f}{F(R)}$ on support

[6] When $n_U = 1$, then the uninformed bidder plays a pure strategy in equilibrium and the analysis is similar to the game with one informed and one uninformed bidder.

$[m_s, R]$ for some $R \leq 1$. The seller chooses the opening price and the secret reserve optimally given this belief. We thus assume that the buyers do not correctly anticipate that the seller chooses the optimal secret reserve. This assumption seems quite reasonable from the real-world point of view where bidders usually cannot surmise the secret reserve exactly. Our results do not depend on the exact distribution that the bidders believe that the secret reserve is drawn from.

As in Section 3, there is one informed and one uninformed bidder and we assume $T = 2$ for illustrative purposes. First we show that for a given opening price, the informed bidder's probability of sniping will be lower and the uninformed bidder's first-period bid will be higher if there is a secret reserve price in addition to the opening price. Suppose the equilibrium cutoff-action pair of an auction without a secret reserve price, as in Section 3, is (v, x). If there is an additional secret reserve price, suppose the corresponding equilibrium cutoff-action pair is (v^s, x^s). Proposition 6 shows that a secret reserve price reduces the probability of sniping and raises bidder 2's bid in period 1.

Proposition 6. The probability of sniping is lower in the secret reserve price auction.

With a secret reserve, the uninformed bidder knows that the price can change even when the informed bidder snipes and the probability of him becoming the high bidder with any given first-period bid is also lower. As a result, he bids more aggressively in period 1. This reduces the benefit from sniping for the informed bidder and the cutoff value for sniping for bidder 1 increases. Thus, a secret reserve price leads to a lower probability of sniping and higher bids by uninformed bidders, supporting the last stylized fact from Section 2.

We show that, in our model, the optimal revenue from a secret reserve price auction is strictly higher than the optimal revenue from a public reserve price auction. Specifically, we show that there is a secret reserve auction that generates higher revenue than the optimal public reserve auction. The model thus provides a rationale behind the frequent use of secret reserve prices on eBay, as evidenced in the seventh stylized fact.

Proposition 7. The optimal revenue is higher from a secret reserve price auction.

An uninformed bidder may win the object even when his valuation is below the reserve price when it is secret. That will not happen if the reserve is public. A secret reserve also leads to a higher probability of bidding by the uninformed bidder compared to a public reserve. Moreover, for a given opening price, a secret reserve increases the uninformed bidder's first-period bid and reduces the probability of sniping by the informed bidder. As a result, the optimal secret reserve price may lead to a higher expected revenue net of the fee for keeping the reserve secret. This also implies that the revenue of auctions depends on both the effective reserve and whether the reserve is public or secret.

6. Conclusion

■ This article 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 some bidders know their private valuations and the other bidders can only learn whether their valuations are above the current price in 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.

The main idea proposed in this article 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 article suggests a simple model where players use the current price as a benchmark for learning about their own preferences. This model can provide a stepping stone for more complicated models of bounded rationality of this nature. 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 negotiation of price can also be used in bargaining or principal-agent problems.

Appendix

■ This Appendix provides the proofs of all the results in the text.

Proof of Theorem 1. In any weakly undominated strategy, bidder 1 will never bid above v_1 and will bid exactly v_1 in one of the periods; that is, $\beta_T^1 = v_1$. If $v_2 \geq p_T$, then bidder 2 chooses b_T^2 to maximize $\int_{p_T}^1 \int_a^{b_T^2} (v - y)\, dG(y)\, dF(v)$ bids and hence bids $b_T^2 = \mathbf{E}\,[y \mid [p_T, 1]]$.[7] Henceforth we use $e\,(p)$ to represent $\mathbf{E}\,[y \mid [p, 1]]$.

First we show that (σ^v, σ^x) is indeed an equilibrium. In this equilibrium, bidder 1 either bids her valuation only in the first period or only in the last period and does not bid in other periods. Expressing m and 0 by x_0 and x_{-1}, respectively, bidder 1's expected utility from bidding in period 1 equals

$$\sum_{t=0}^{T-1} (v_1 - x_t)(F(x_t) - F(x_{t-1}))\,\mathbf{1}_{\{v_1 > x_t\}} + (1 - F(x_{T-1}))(v_1 - e(x_{T-1}))\,\mathbf{1}_{\{v_1 > e(x_{T-1})\}}.$$

Her expected utility from bidding in period T (if $v_1 > e(m)$) equals

$$(v_1 - m)\,F(m) + (1 - F(m))(v_1 - e(m)) = v_1 - mF(m) + \int_m^1 y\, dF(y).$$

If $v_1 > e(x_{T-1})$, bidder 1's expected utility from bidding in period 1 is

$$v_1 - \sum_{t=0}^{T-1} x_t\,(F(x_t) - F(x_{t-1})) + \int_{x_{T-1}}^1 y\, dF(y)$$

and she will be better off by bidding in period T. If $v_1 \leq e\,(m)$ then she is clearly better off by bidding her valuation in the first period. In fact, she is indifferent between bids in periods 1 and T when $v_1 = v \in (x_{T-1}, e(x_{T-1}))$ where

$$\sum_{t=0}^{T-1} (v - x_t)(F(x_t) - F(x_{t-1})) = v - mF(m) + \int_m^1 y\, dF(y)$$

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

When $v_1 \in (x_l, x_{l+1})$ for some $l < T - 1$ and $v_1 \geq e(m)$, the expected payoff from bidding v_1 in the first period is $\sum_{t=0}^{l}(v_1 - x_t)(F(x_t) - F(x_{t-1}))$. By construction of v, for any $v_1 < x_{l+1} < v$,

$$\sum_{t=0}^{l} (v_1 - x_t)(F(x_t) - F(x_{t-1})) > \sum_{t=0}^{T-1} (v_1 - x_t)(F(x_t) - F(x_{t-1}))$$

$$> v_1 - mF(m) + \int_m^1 y\, dF(y).$$

Therefore, if $v_1 < v$ then bidder 1 prefers bidding her valuation in period 1. She strictly prefers sniping if $v_1 > v$. Given σ^x, no other strategy gives bidder 1 a higher payoff. If bidder 1 follows another strategy where if $v_1 \leq v$, $\beta_t^1 \geq \min\{x_t, v_1\}$ for all t and she snipes if $v_1 > v$, then the history after any period t would be the same as it would have been if $b_1^1 = v_1$. Bidder 2 will place the same bids in both cases and bidder 1 will be equally well off by breaking her bids into many separate bids or by bidding once, leading to the same outcome in both cases. If $\beta_t^1 \in (x_{t-1}, x_t)$ where $v_1 \geq x_t$ and $p_t > p_{t-1}$ (that is, bidder 2 placed a bid in period $t - 1$) then, given σ^x, bidder 1 will be strictly better off by choosing $b_1^1 = v_1$ if $v_1 \leq e\,(\beta_t^1)$ and by sniping if $v_1 > e\,(\beta_t^1)$. Thus, σ^v is a best response to σ^x.

Now suppose in period $t < T$, bidder 1 is the high bidder and $v_2 \geq p_t = x_{t-1}$. Expressing v by x_T, bidder 2 will choose x_T to maximize

$$\frac{\sum_{l=t-1}^{T-1} \int_{x_l}^1 \int_{x_l}^{x_{l+1}} (z - y)\, dF(y)\, dF(z)}{1 - F(x_{t-1})}.$$

[7] Suppose β_T^1 is drawn from the continuous distribution G with support $[a, c]$ where $p_T \leq a$ and $c \leq 1$.

Therefore, the first-order conditions require

$$\int_{x_{t-1}}^{x_{t+1}} \left(y - x_t^*\right) dF(y) = 0$$

$$\Longrightarrow x_t^* = \frac{\int_{x_{t-1}}^{x_{t+1}} y\, dF(y)}{F(x_{t+1}) - F(x_{t-1})}$$

$$\therefore x_t^* = \mathbf{E}[y|[x_{t-1}, x_{t+1}]].$$

The second-order condition is satisfied by x_t^* and boundary points are not optimal. Using the Envelope Theorem, we can see that a second-order condition on a single variable around x_t^* maximizes bidder 2's expected payoffs. Also, his *ex ante* expected utility-maximizing bid choice is dynamically consistent. By construction, σ^x is a best response to σ^v. Any belief system consistent with the player strategies works for this to be an equilibrium. Therefore, if a pair (v, x) satisfying the optimality conditions exists, then (σ^v, σ^x) is an equilibrium. Moreover, we can see from the first-order conditions that $x_t - x_{t-1} > 0$ and $v < e\,(x_{T-1}) < 1$.

We can use Brouwer's fixed point theorem on $M : [m, 1]^T \longrightarrow [m, 1]^T$ to show the existence of an equilibrium cutoff-action pair. The first $T - 1$ elements of M are given by $\mathbf{E}\,[y \mid [x_{t-1}, x_{t+1}]]$. The last element is the function

$$\max\left[\frac{\int_m^1 y\, dF(y) - \sum_{t=1}^{T-1} x_t\,(F(x_t) - F(x_{t-1}))}{1 - F(x_{T-1})}, m\right].$$

The domain is nonempty, convex, and compact and M is continuous. By Brouwer's fixed point theorem, there exists a fixed point of M. By construction, $x_t \leq x_{t+1}$ and $x_{T-1} \leq v$. Furthermore, $v > m$ 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. There is no $v \leq 1$ such that (v, x_a) and (v, x_b) satisfy the equilibrium criteria. If $x_{a,1} < x_{b,1}$ then that will lead to $x_{a,2} < x_{b,2}$ 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 and thus bids in the two periods are strategic complements. Strict single-crossing property defined by Milgrom and Shannon (1994) is also satisfied. If $x_{a,1} < x_{b,1}$ then $x_{a,t} < x_{b,t}$ for all t, then the cutoff corresponding to x_a is smaller than the cutoff corresponding to x_b.

Now suppose (v_a, x_a) and (v_b, x_b) are two equilibrium cutoff-action pairs with $v_a < v_b \leq 1$. Hence, $x_{a,t} < x_{b,t}$ 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 $v_1 - y$, $\int_m^{v_1} (v_1 - y)\, dF(y)$ using x for the intervals. Therefore, if $v_1 = v_a$, bidder 1's expected payoff from bidding in period 1 is higher when bidder 2 follows x_a instead of x_b. She is indifferent between bidding in period 1 and T when bidder 2 follows x_a. She gets the same utility from bidding late whether bidder 2 follows x_a or x_b. Therefore, bidder 1 would strictly prefer bidding late if bidder 2 follows x_b implying that $v_b < v_a$. Therefore, a unique equilibrium cutoff-action pair (v, x) exists.

Now we show that this is essentially the unique equilibrium when both bidders place a bid in the first period with positive probabilities. Define $\chi\,(\sigma_1, \sigma_2)$ to be a T-element vector of bidder 2's bids in the T periods of the auction along a history where $w_t \neq 2$ and $v_2 \geq p_t$ when the two bidders follow the strategy profile (σ_1, σ_2). Lemma A1 shows that if (σ_1, σ_2) is an equilibrium, then $b_\tau^1 > m$ implies that $\beta_t^1 \geq \min\{\chi_t(\sigma_1, \sigma_2), v_1\}$ for all $t \geq \tau$. Suppose (σ_1, σ_2) is an equilibrium. Then the public history and, therefore, the information bidder 2 receives about v_2 at any point will be the same if we change σ_1 in a way that if $b_\tau^1 > m$ then $b_\tau^1 = v_1$. Given that bidder 1 bids in period 1 for some valuations, bidder 2 will place a bid in the first period if $v_2 \geq m$. Suppose bidder 1 prefers placing her first bid in period $\tau \in \{2, \ldots, T - 1\}$ to bidding in the first period if $v_1 \in [\underline{v}, \overline{v}]$. This implies that $\underline{v} > \chi_1(\sigma_1, \sigma_2)$. Then, if bidder 1 does not place a bid by period $\tau - 1$, bidder 2 will bid at least $\min\{\underline{v}, e(m)\}$ in period τ as a best response to σ_1. This implies that in any equilibrium, bidder 1 places her first bid in either the first or the last period. This implies that there is a value ς such that (σ_1, σ_2) leads to the same outcome as $(\sigma^\varsigma, \sigma_2)$. Given the construction of the equilibrium at the beginning of a theorem, σ_2 can be characterized by σ^χ for some vector χ. As there is a unique cutoff-action pair, any equilibrium (σ_1, σ_2) leads to the same winner and final price as in equilibrium (σ^v, σ^x) and the equilibrium outcome is unique. *Q.E.D.*

Lemma A1. If (σ_1, σ_2) is an equilibrium then $b_\tau^1 > m$ implies that if $p_t > p_{t-1}$ then $\beta_t^1 \geq \min\{\chi_t(\sigma_1, \sigma_2), v_1\}$ for all $t \geq \tau$.

Proof. We will show that $\beta_t^1 \in (\chi_{t-1}(\sigma_1, \sigma_2), \chi_t(\sigma_1, \sigma_2))$ while $v_1 \geq \chi_t(\sigma_1, \sigma_2)$ and $p_t > p_{t-1}$ cannot happen in an equilibrium. If $\beta_t^1 \in (\chi_{t-1}(\sigma_1, \sigma_2), \chi_t(\sigma_1, \sigma_2))$ then $w_{t+1} = 2$ if $v_2 \geq p_t$ (otherwise this is irrelevant). On the equilibrium path, after observing $w_{t+1} = 2$, bidder 2 correctly anticipates that either $v_1 = p_{t+1}$ implying that bidder 2 will be the winner of the auction or v_1 is greater than or equal to some $\underline{v}$ so that $\beta_t^1 \in (\chi_{t-1}(\sigma_1, \sigma_2), \chi_t(\sigma_1, \sigma_2))$ is profitable for bidder 1. Note that by choosing $\beta_t^1 < \chi_t(\sigma_1, \sigma_2)$, compared to $\beta_t^1 \geq \chi_t(\sigma_1, \sigma_2)$, bidder 1 reduces p_{t+1} thus reducing the probability of bidder 2 stopping bidding (because of learning $v_2 < p_{t+1}$) whereas his payment if bidder 2 gets a negative signal stays unchanged. Therefore, $\underline{v} > \chi_t(\sigma_1, \sigma_2)$ if (σ_1, σ_2) is an equilibrium. Therefore, if $v_2 \geq p_{t+1}$ then bidder 2's best response is to bid at least $\min\{\underline{v}, e\,(p_{t+1})\}$. If $p_\tau = p_{t+1}$ for all $\tau \in \{t + 2, \ldots, T\}$ then bidder 2 will bid $e\,(p_{t+1})$ in at least one of the periods. If bidder 1 bids in any period between period $t + 1$ and $T - 1$, bidder 2 will want his bid to be at least $\underline{v}$ as he knows that $v_1 \geq \underline{v}$ in that case. However, if $b^2_{t+1} \geq \min\{\underline{v}, e(p_{t+1})\}$ then σ_1 is clearly suboptimal for bidder 1 when $v_1 = \underline{v}$. Note that if $\underline{v} > e(p_{t+1})$ and bidder 1 does not bid until the last period, she wins for sure. Using the calculations in Theorem 1, we can easily show that bidder 1 would have been better off by just sniping in that case. Thus, $\beta^1_t \in (\chi_{t-1}(\sigma_1, \sigma_2), \chi_t(\sigma_1, \sigma_2))$ whereas $v_1 \geq \chi_t(\sigma_1, \sigma_2)$ cannot happen for any t in an equilibrium and $b^1_\tau > m$ implies that, for all $t \geq \tau$, $\beta^1_t \geq \min\{\chi_t(\sigma_1, \sigma_2), v_1\}$. *Q.E.D.*

Proof of Proposition 1. 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 to each other. Because F is atomless and strictly increasing, $F(x_t) - F(x_{t-1}) \to 0$. By using Theorem 1, $x_t = \mathbf{E}[y \mid [x_{t-1}, x_{t+1}]]$. 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 $x_{T-1} \to \mathbf{E}[y \mid [m, e(m)]] > m$. 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, $\lim_{T\to\infty} x(x_t - x_{t-1}) = 0$ for all t. Now suppose, $\lim_{T\to\infty} x_{T-1} = 1 - \eta$ for some $\eta > 0$. Because $\lim_{T\to\infty}(x_t - x_{t-1}) = 0$ for all t, $x_{T-1} = \mathbf{E}[y \mid [x_{T-2}, v]]$ implies $v \to 1 - \eta$. However, equation 2 implies that $v \to e(1\text{-}\eta)$. 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 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 the same as the standard model's outcome as $T \to \infty$. *Q.E.D.*

Proof of Proposition 2. When bidder 2 is uninformed, β^2_T is either above v_2 or equals his conditional valuation given $v_2 \geq p_T$. That is, bidder 2 bids above v_2 in expectation. As $\beta^1_T = v_1$, expected revenue with one informed bidder and one uninformed bidder,

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

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

$$\pi_I = 2mF(m)(1 - F(m)) + \int_v^1\left(\int_m^z y\,dF(y) + z(1 - F(z))\right)dF(z) + \int_m^v\left(\int_m^z y\,dF(y) + z(1 - F(z))\right)dF(z).$$

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. *Q.E.D.*

Proof of Proposition 3. When $v_1 = v$, bidder 1's expected payoff from bidding v_1 in period 1 is

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

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

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

Because (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, the seller's expected revenue equals $\alpha\,\pi_U + (1 - \alpha)\,\pi_I$ for all $\alpha \in (0, 1]$. As π_U is greater than π_I for any given T, the expected revenue is increasing in α. *Q.E.D.*

Proof of Proposition 4. Bidder 1's payoff from bidding late and early are the same if $v_1 = v$. That is,

$$v - mF(m) - \int_m^1 y\,dF(y) = vF(x_h) - mF(m) - x_l(F(x_l) - F(m)) - \int_{x_l}^b y\,dF(y) - x_h(F(x_h) - F(b))$$

$$\Longrightarrow v = \frac{\int_m^1 y\,dF(y) - x_l(F(x_l) - F(m)) - \int_{x_l}^b y\,dF(y) - x_h(F(x_h) - F(b))}{1 - F(x_h)}.$$

Bidder 2 chooses b to maximize

$$\int_0^b \int_0^{x_l} (z - y)\,dF(y)\,dF(z) + \int_{x_l}^b \int_{x_l}^{\mathbf{E}[y|[x_l, b]]} (z - y)\,dF(y)\,dF(z)$$
$$+ \int_b^1 \int_0^{x_h} (z - y)\,dF(y)\,dF(z) + \int_{x_h}^1 \int_{x_h}^{v} (z - y)\,dF(y)\,dF(z).$$

Bidder 2's payoffs are optimized when

$$x_l = \mathbf{E}[y|[m, \mathbf{E}[y|[x_l, b]]]],\; x_h = \mathbf{E}[y|[b, v]],\; \text{and } b = \mathbf{E}[y|[\mathbf{E}[y|[x_l, b]], x_h]].$$

Following Theorem 1, we can easily show that this is indeed an equilibrium. *Q.E.D.*

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

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

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

$$f(y) \int_y^{\overline{b}_1} (u - y)\,dG(u) + \frac{f(y)}{F(y)} \int_y^1 \int_y^{\overline{b}_1} (z - u)\,dG(u)\,dF(z) + g(y) \int_0^y (z - y)\,dF(z) = 0, \tag{A1}$$

with $G(\underline{b}_1) = 0$ and $G(\overline{b}_1) = 1$. To show existence, define $\Gamma(y) = \int_y^{\overline{b}_1} G(u)\,du$. Then, $G(y) = -\Gamma'(y)$ and $g(y) = -\Gamma''(y)$. Equation A1 becomes

$$f(y)\left(\int_y^1 z\,dF(z) - (1 - 2F(y))\overline{b}_1 - yF(y)\right) - F(y) \int_0^y (z - y)\,dF(z)\Gamma''(y)$$
$$+ f(y)\left(\int_y^1 (z - y)\,dF(z)\Gamma'(y) + (1 - 2F(y))\Gamma(y)\right) = 0. \tag{A2}$$

In addition, $\Gamma(\overline{b}_1) = 0$ and $\Gamma'(\overline{b}_1) = -1$. Equation A2 is a linear second-order differential equation with two initial conditions. Therefore, there exists a unique solution for $G(.)$ for given $\underline{b}_1$ and $\overline{b}_1$. We can find $\underline{b}_1$ using $G(\underline{b}_1) = 0$. To find $\overline{b}_1$, notice that when $y \to \overline{b}_1$,

$$\int_{\underline{b}_1}^{\overline{b}_1} \int_0^1 F(u)(z - u)\,dF(z)\,dG(u) = K_1. \tag{A3}$$

From equation A3, 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 $\overline{b}_1 > \mathbf{E}[y|[0, 1]]$. For a given support, a unique G satisfies the indifference condition given by equation A1. 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 is closed on the left and open on the right. If it is closed on the right then $g(\overline{b}_1) \int_0^{\overline{b}_1} (z - \overline{b}_1)\,dF(z) = 0$. That implies that $g(\overline{b}_1) = 0$. 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 $\eta \to 0^+$, 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 $[\underline{b}_1, \overline{b}_1)$ where $\underline{b}_1 > 0$. Hence, G is an equilibrium distribution and (σ^G, σ^G) is an equilibrium. *Q.E.D.*

Proof of Proposition 5. Suppose G solves equation A1 for $y \in [\underline{b}_1, \overline{b}_1)$ and $\psi \in (0, \overline{b}_1 - \underline{b}_1)$. Define $\overline{G}$ to be truncated distribution G on $[\underline{b}_1 + \psi, \overline{b}_1)$. Then $\overline{G}$ solves equation A1, implying that $\overline{G}$ is an equilibrium distribution. Thus, there is a continuum of equilibria. *Q.E.D.*

Proof of Theorem 3. First we will show the existence of an equilibrium cutoff-distribution pair when there are one informed bidder and two uninformed bidders. Then we will argue that this result carries through when there are one or more informed bidders and two or more uninformed bidders. Suppose bidder 1 is informed and bidders 2 and 3 are uninformed. Bidder 1 bids v_1 in period 1 if $v_1 \le v$ and bids nothing in period 1 and v_1 in period 2 otherwise. Bidder $i \in \{2, 3\}$ bids $y \in [\underline{b}_1, \overline{b}_1]$ using the distribution G in period 1. In period 2, he bids $\mathbf{E}\,[y \mid [p_2, 1]]$ if $v_i \ge p_2$. In equilibrium, bidder 2 is indifferent between bidding any $y \in [\underline{b}_1, \overline{b}_1]$ in period 1. This occurs if

$$\begin{aligned}
&F^2(y)\int_0^y (z-y)\,dF(z)\,g(y)\\
&+ f(y)F(y)\left(\int_y^1 \left(z-\overline{b}_1\right)\,dF(z) + \left(\overline{b}_1 - y\right)F(y) + (1-2F(y))\int_y^{\overline{b}_1} G(u)\,du\right)\\
&+ f(y)G(y)\left(\begin{array}{c} F(y)\left(\displaystyle\int_0^{\min[e(y),v]}(z-y)\,dF(z) - \int_y^1 (z-y)\,dF(z)\right)\\ + \displaystyle\int_y^1\int_y^{\min[e(y),v]} (z-z_1)\,dF(z_1)\,dF(z)\end{array}\right) = 0.
\end{aligned} \tag{A4}$$

Distribution G will have no mass at 0 as the left-hand side is strictly positive when $y = 0$ in that case. Hence, an uninformed bidder will place a bid in the first period with probability 1. To maximize the expected payoff at the supremum of the support, we need

$$F\left(\overline{b}_1\right)\int_0^v \left(z-\overline{b}_1\right)\,dF(z) + \int_{\overline{b}_1}^1\int_{\overline{b}_1}^v (z-u)\,dF(u)\,dF(z) = 0. \tag{A5}$$

For a given v, there is a unique $\overline{b}_1$ that satisfies equation A5. Using arguments similar to those in the proof of Theorem 2, we can show that there is a unique solution to G. Equation A4 implies that there is a unique $\underline{b}_1$ that solves $G(\underline{b}_1) = 0$. That is, the support of G cannot be truncated like in the two-uninformed bidder case. The support will be closed given equations A4 and A5.

Now, $W(1, v_1) = 1$ and $P(1, v_1) > P(2, v_1)$ if $v_1 > e(\overline{b}_1)$ where $P(t, v_1)$ and $W(t, v_1)$, respectively, denote expected payment and probability of winning of bidder 1 when she bids v_1 in period t and nothing in the other period. There is a $v < 1$ such that bidder 1 is indifferent between bidding in period 1 and 2 if $v_1 = v$. For a given G, there is at most one equilibrium cutoff v. Then, $vW(1, v) - P(1, v) = vW(2, v) - P(2, v)$ and (v, G) is an equilibrium cutoff-distribution pair.

For a given v, G and $[\underline{b}_1, \overline{b}_1]$ are known. By equalizing payoffs from bidding early and late, we get an expression for v that depends on F and G and we label that function $M_1(v)$. Define M:$[\mathbf{E}\,[y \mid [0, 1]], 1] \to [\mathbf{E}\,[y \mid [0, 1]], 1]$ such that $M(v) = \max\,[M_1(v), \mathbf{E}\,[y \mid [0, 1]]]$. Because M is continuous and the domain is nonempty, compact, and convex, we can use Brouwer's fixed point theorem. At the fixed point, $v > \mathbf{E}\,[y \mid [0, 1]]$, as otherwise bidder 1 gets zero profit from bidding late. Hence, an equilibrium cutoff-distribution pair exists.

Suppose there are n_I informed bidders and n_U uninformed bidders where $n_I \ge 1$ and $n_U \ge 2$. Then, the uninformed bidders will play mixed strategy in period 1 in any equilibrium. To find the distribution from which uninformed bidders draw their first-period bids in a symmetric equilibrium and show its existence, we can follow the method used for the two uninformed bidders, and the two uninformed bidders and one informed bidder case. The differential equation expressing the distribution is similar to the case with three bidders but is quite cumbersome and hence is not presented here. Informed bidders with high valuation will be better off by sniping. Notice that, when the second-highest bidder is an uninformed bidder, conditional on the probability of winning being close to one, an informed bidder's expected payment is lower when she snipes because the uninformed second-highest bidder will overbid in expectation. Therefore, even with more than one informed bidder, $P(1, 1) > P(2, 1)$ and $W(1, 1) \le W(2, 1)$. Hence, sniping is better for an informed bidder i if $v_1 = 1$. This implies that when there are a finite number of bidders there will exist a cutoff $v < 1$ such that informed bidders with valuation above v will snipe. Knowing uninformed bidders' equilibrium distribution, existence of informed bidders' equilibrium cutoff can be shown using Brouwer's fixed point theorem. There is a symmetric equilibrium where informed bidders use the same cutoff value v and uninformed bidders draw first-period bids from the same distribution G on interval $[\underline{b}_1, \overline{b}_1]$. *Q.E.D.*

Proof of Proposition 6. Suppose given the bidders' belief on the distribution of r_s, the uninformed bidder believes that the highest and second highest of v_1 and r are distributed according to distributions H and S on $[m_s, 1]$ and

$[m_s, R]$, respectively. Thus, H first-order stochastically dominates F. Now, the expected payoff for bidder 2 from bidding $x^s{}_1$ is

$$\int_{m_s}^{x_1^s}\int_{m_s}^{1}\int_{w}^{x_1^s}(z-y)\,dH(y)\,dF(z)\,dS(w)$$
$$+\int_{m_s}^{R}\int_{\max[x_1^s,w]}^{1}\int_{\max[x_1^s,w]}^{\min[v^s,e(\max[x_1^s,w])]}(z-y)\,dH(y)\,dF(z)\,dS(w).$$

This leads to the first-order condition,

$$h\left(x_1^s\right)\int_{m_s}^{x_1^s}\left(z-x_1^s\right)\,dF(z)+f\left(x_1^s\right)\int_{x_1^s}^{\min[v^s,e(x_1^s)]}\left(z-x_1^s\right)\,dH(z)=0.$$

The cutoff v^s equates bidder 1's expected payoff from bidding v_1 in period 1 and sniping. Recall that, in a public reserve auction, x_1 satisfies

$$f\left(x_1\right)\int_{m_s}^{x_1}\left(z-x_1\right)\,dF(z)+f\left(x_1\right)\int_{x_1}^{\min[v,e(x_1)]}\left(z-x_1\right)\,dF(z)=0.$$

Given the structure of the distribution H, these imply that $x^s{}_1 > x_1$ and, in turn, $v^s > v$. *Q.E.D.*

Proof of Proposition 7. The optimal revenue from secret reserve price auctions will be at least as high as that from a public reserve price auction because the seller can always choose a secret reserve equaling the opening price. Suppose the optimal reserve in a public reserve auction is m^* and bidder 2's corresponding first-period bid and bidder 1's cutoff value are x_1^* and v^*, respectively. Now suppose the seller chooses secret reserve $r_s = m^*$ and opening price $m_s = \widetilde{m}$ such that corresponding bidder 2's first-period bid x_1^s equals x_1^*. Using Proposition 6, we can easily show that $\widetilde{m} < m^*$. This implies that the uninformed bidder is more likely to place a first-period bid where the first-period bid will be the same as in the optimal auction without a secret reserve price. As a result, the corresponding cutoff value v^s in the secret reserve price auction will be greater than v^*. Thus bidder 1 will snipe less frequently and sniping reduces expected revenue. Given the secret reserve of m^*, the minimum price conditional on a sale is the same and the probability of a sale is higher, as the uninformed bidder may bid above m^* even when $v_2 < m^*$ and bidder 2's average bid is higher. If bidder 1 does not snipe then expected revenue also rises, as bidder 2 bids at least once with higher probability. Hence, the expected revenue will be higher from this secret reserve price auction than from the optimal public reserve auction. Therefore, the optimal secret reserve price auction must generate strictly higher revenue than the optimal public reserve price auction. *Q.E.D.*

References

Ariely, D., Ockenfels, A., and Roth, A.E. "An Experimental Analysis of Ending Rules in Internet Auctions." *RAND Journal of Economics,* Vol. 36 (2005), pp. 891–908.

Bajari, P. and Hortacsu, A. "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions." *RAND Journal of Economics,* Vol. 34 (2003), pp. 329–355.

Compte, O. and Jehiel, P. "The Wait and See Option in Ascending Price Auctions." *Journal of the European Economic Association*, Vol. 2 (2004), pp. 494–503.

Ely, J.C. and Hossain, T. "Sniping and Squatting in Auction Markets." Working Paper, Northwestern University, 2006.

Ergin, H. "Costly Contemplation." Working Paper, MIT, 2003.

Hasker, K., Gonzalez, R., and Sickles, R.C. "An Analysis of Strategic Behavior in eBay Auctions." Working Paper, Rice University, 2004.

Hossain, T. "Essays on Bounded Rationality in Auctions and Bargaining." Chapter 1, Ph.D. Dissertation, Princeton University, 2004.

——— and Morgan, J. ". . . Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay." *Advances in Economic Analysis and Policy,* Vol. 62 (2006), Article 3.

Milgrom, P. and Shannon, C. "Monotone Comparative Statics." *Econometrica,* Vol. 62 (1994), pp. 157–180.

Rasmusen, E.B. "Strategic Implications of Uncertainty over One's Own Private Value in Auctions." *Advances in Theoretical Economics,* Vol. 61 (2006), Article 7.

Roth, A.E. and Ockenfels, A. "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Theory and Evidence from a Natural Experiment on the Internet." NBER Working Paper no. 7729, 2001.

Simon, H.A. "Rational Choice and the Structure of the Environment." *Psychological Review,* Vol. 63 (1956), pp. 129–138.

———. *Models of Man: Social and Rational*. New York: Wiley, 1957.

Sloman, S.A. "Two Systems of Reasoning." In T. Gilovich, D. Griffin, and D. Kahneman, eds., *Heuristics and Biases*. Cambridge, UK: Cambridge University Press, 2002.

Song, U. "Nonparametric Estimation of an eBay Auction Model with an Unknown Number of Bidders." Working Paper, University of British Columbia, 2004.