Equilibria in eBay Auctions

July 2017

Unjy Song

Seoul National University

1 Introduction

This paper characterizes the set of symmetric Bayesian-Nash equilibria in eBay auctions within the independent private values model. There are multiple equilibria with respect to the bidding time but top two bidders bid their valuations by the end of auction and hence the expected revenue is the same in every equilibrium.

EBay auctions are a mixture of traditional ascending and second-price sealed-bid auctions. The eBay auction takes an ascending-bid format but the highest bid is not revealed. A bidder can know every bid except the highest and whether or not he is the highest bidder. The eBay auction ends at the fixed ending time, hence a bidder may not have time to response to his competitior's bid at the very end of the auction. Therefore, the eBay auction works more like the second-price sealed-bid auction rather than the ascending auction at the end of the auction. In eBay, there are auctions which sell multiple items and/or have a fixed-price component such as "buy-it-now" option. In this paper I focus only on auctions which sell a single item without using a buy-it-now option.

While there has beeen a huge literature on empirical analysis of eBay auction data including field experiments, research on equilibria in eBay auctions has been limited. This paper sets up a bidding model for eBay auctions and characterizes all equilibria; this would improve understanding of eBay auction data for empirical work. The rest of the paper is organized as follows: The next section explains the eBay auction format, and Section 3 presents an eBay bidding model and its equilibria with discussion.

2 Background: eBay Auctions

This section briefly explains the eBay auction mechanism. Bajari and Hortaçsu (2002a,b) and Bryan et al. (2000) offer richer descriptions. I consider only auctions in which a single item is sold. I exclude secret reserve price auctions as well as auctions ended by a bidder's use of a "buy it now" option. An eBay auction starts as soon as a seller registers it. An eBay seller has several options when she lists her item. She can set a starting price and also choose the time length of her auctions: one, three, five, seven or ten days. Potential buyers can find auctions of interest by browsing the categorized auction listings or by using a search engine. No advance announcement of an auction exists, therefore there is no reason to expect all potential bidders

to become aware of the auction at the same time. In the next section, I will explicitly model the stochastic arrival times of bidders during an auction.

All auctions proceed according to the rules pre-announced by the eBay. All eBay auctions use an open, ascending-bid format that is different from a more traditional ascending auction's bid format in two respects. First, there is a fixed ending time instead of a "going-going-gone" ending rule. Second, eBay uses the proxy bidding system. A new bidder is asked to submit a *cutoff-price*[1], a maximum bid, instead of his instant bid amount. If the new bidder outibids the current winner, the proxy bidding system then will issue a proxy bid equal only to the minimum increment over the current winner's bid which is now the next highest bid. If the new bidder's cutoff price is less than the current winner's cutoff price, the proxy bidding system will not issue a bid for the new bidder. Instead, the proxy bidding system issue a proxy bid for the current winner as much as the minimum increment over the new bidder's cutoff price. The maximum proxy bid is posted as the *standing price* next to a current winner's identity. For example, consider an auction in which a seller starts the bidding at $5, and the *first-arrived* bidder submits $25 as his cutoff-price. The proxy server issues a proxy bid of $5 on the first-arrived bidder's behalf and posts $5 as the standing price. Suppose another bidder arrives and submits a cutoff-price of $20. The proxy server then bids $20 plus the minimum increment for the first-arrived bidder and displays it as the standing price. As a result, under the eBay's proxy bidding system, the standing price is the second-highest existing cutoff-price plus the minimum increment. During the course of the auctions, whenever the cutoff-price submitted by a new bidder is not high enough to lead the auction, or the current auction leader is outbid, eBay notifies the bidder via e-mail so that he may revise his cutoff-price if he so desires. A bidder may keep or increase his previous cutoff-price at any time, but may not decrease it. A bidder may retract his bid upon the seller' permission. But the bid retraction is extremely rare. If there are more than one bidder who submitted the same bid, an early bidder will win the auction and pay the same bid.

Once an auction has concluded, the winner is notified by e-mail and pays the standing price posted at the closing time. Thus, a winner pays the second-highest bid plus the minimum increment. During the auction eBay reveals bidders' cutoff-prices and bidding times except

---

[1] I introduced a new term "cutoff price" in order to distinguish the intermediate bid and the final bid in the bidding model in the next section. EBay does not use a terminology of "cutoff price".

the highest bid. If a bidder submitted cutoff-prices multiple times, every cutoff-price and its corresponding bidding time is shown. All auction listings and their results remain publicly available on eBay for at least one month after the auction closes.

3   Setup and Characterization of Equilibria

Consider an eBay auction of a single object. The number of potential bidders, N, is a random variable, with $p_n = \Pr(N = n)$ where $\Pr(N \geq 2) > 0$. A potential bidder i's valuation V^i is an independent draw from the continuous distribution $F(\cdot)$, having support on $[\underline{v}, \overline{v}]$. Each bidder knows only his valuation, the distribution $F(\cdot)$, and the probabilities p_n. The analysis would be identical were bidders to know the realization of N, as is usually assumed in the literature. For the sake of simplicity, I ignore the minimum increment.[2]

The auction is conducted over an interval of time $[0, \tau]$. An eBay seller can choose τ among 1, 3, 5, 7, and 10 days. Since the eBay auction is held too long to monitor every minute, I assume that each bidder monitors the auction only at finite monitoring times. A set of monitoring times of bidder i is exogenously given by a finite set: $T^i = \{t_0^i, t_1^i..., \tau^i\}$. The number of monitoring times and each monitoring time are known to bidder i at time t_0^i but are unknown to i's competitors. At each monitoring time in T^i, bidder i sees the standing price and can submit a new cutoff-price if bidder i wants. There are no limitations on the total number of submissions. Especially I call the final cutoff-price the *bid* which is represented by a random variable B^i.

The standing price as a function of time is denoted by $P(t)$ ($t \in [0, \tau]$). The standing price at time 0, $P(0)$ is initialized at the starting price set by the seller. As the auction proceeds, $P(t)$ is raised to the value of the second-highest cutoff-price transmitted prior to t. If the number of existing bidders is less than two, $P(t)$ stays at the starting price. The auction ends at time τ, with the highest bidder declared the winner at a price of $P(\tau)$ which is equal to the second-highest cutoff-price. If there are more than one bidder who submitted the same cutoff price, the bidder who submitted the winning bid earlist wins.

I consider a strategy for bidder i that specifies the cutoff price he will submit at each monitoring time $t \in T^i$ as a function of his valuation, the history of the standing prices that he has

---

[2]The amount of the minimum increment is predetermined and posted on eBay website. The minimum increment is, for instance, $0.50 when the standing price is $5.00 - $24.99. Such a small, minimum increment seems unlikely to affect bidders' bidding behaviors significantly.

monitored.[3] The history of the standing prices at bidder i's monitoring time is denoted by a set $H^i(t^i) = \{P(t) | t \in T^i, t \leq t^i\}$. For example, $H^i(t_2^i) = \{P(t_0^i), P(t_1^i), P(t_2^i)\}$. A random variable $C^i(t)$ $(t \in [0, \tau])$ represents bidder i's most recent cutoff-price at time t. If $C^i(t) = 0$, it indicates that bidder i has not transmitted a cutoff-price by time t. Obviously $C^i(t) = 0$ for $t < t_0^i$; and $B^i = C(\tau^i) = C(\tau)$. Those who ever submitted a cutoff-price are called *actual bidders.* Potential bidder i's strategy can be described by functions $s_t(v^i | H^i(t))$ $(t \in T^i)$ which specify his cutoff-price at each monitoring time given the history of the standing prices.

Proposition 1 provides a characterization of all symmetric Bayesian-Nash equilibria of this game and show that all equilibria satisfy two key conditions: (a) No bidder ever submits a cutoff-price greater than his valuation; and (b) At his final monitoring time (τ^i), bidder i submits a cutoff-price equal to his valuation if he has not yet done so on the condition that his valuation is greater than the standing price, $P(\tau^i)$. Note that many patterns of bidding behavior are possible in the equilibrium. For example, bidder i may submit a cutoff-price equal to his valuation v^i as soon as he finds the auction; he may postpone his submission until time τ^i; or he may submit a cutoff-price lower than v^i and update his cutoff-price over time.

Proposition 1 *The strategies $S^o = (\ s_{t_0^i}^o\ (v^i | H^i(t_0^i)), ...s_{\tau^i}^o(v^i | H(\tau^i))\ )$ constitute a symmetric Bayesian-Nash equilibrium if and only if they induce:*

(a) $C^i(t) \leq V^i$, $\forall\ t\ \in [0, \tau]$ and (b) $B^i = V^i$ if $V^i > P(\tau^i)$.

Proof. Let $M = \underset{k \neq i}{\arg \max} B^k$. Namely, M is the identity of bidder with the highest bid other than bidder i.

(Sufficiency) Note that the auction price is B^M whenever bidder i wins the auction. Given B^M, bidder i's payoff is:

$$
\begin{cases}
V^i - B^M, & \text{if } B^i > B^M; \text{ or} \\
& \text{if } B^i = B^M \text{ and bidder } i \text{ bid earlier than bidder } M \\
0, & \text{if } B^i < B^M; \text{ or} \\
& \text{if } B^i = B^M \text{ and bidder } M \text{ bid earier than bidder } i
\end{cases}
$$

.

---

[3] I do not include in the bidding function the previous bidders' identities and cutoff prices revealed during the auction. First, among cutoff prices only the standing price matters to bidders' payoffs. Second, the previous bidders' identities played the same role as the number of previous bidders in symmetric equilibria, on which I focus. Inclusion of the number of previous bidders in bidding functions would have no effect on Proposition 1.

Thus the highest payoff bidder i can achieve is $(V^i - B^M)$ if $V^i > B^M$; and zero if $V^i \leq B^M$. Accordingly, any strategy which induces bidder i (i) to win the auction when $V^i > B^M$; and (ii) to obtain zero payoff when $V^i \leq B^M$, is a best response. First, condition (a) guarantees (ii), because a bidder cannot obtain negative payoff if he bids no more than his valuation. Next, condition (b) guarantees (i). By construction, $B^M \geq P(t)$ for all t. Hence if $V^i > B^M$, then $V^i > B^M \geq P(\tau^i)$ and so $B^i = V^i > B^M$ according to condition (b). Accordingly if bidder i bids according to condition (b), bidder i wins the auciton whenever $V^i > B^M$.

(Necessity)

A. The necessity of that (b) $B^i = V^i$ if $V^i > P(\tau^i)$

Below I first show that B^M has a full support over $[\underline{v}, \overline{v}]$ and then show that any strategy profile on the equilibrium should induce $B^i = V^i$ if $V^i > P(\tau^i)$.

(1) $\Pr(b_1 < B^M < b_2) > 0, \quad \forall\, \underline{v} < b_1 < b_2 < \overline{v}$.

Since we are considering symmetric equilibria with symmetric bidders, the support of B^i is the same for all bidder i and the support of B^M is also the same as the common support of B^i. Suppose there exist b_1 and b_2 such that $\Pr(b_1 < B^M < b_2) = 0$ where $\underline{v} < b_1 < b_2 < \overline{v}$. Let

$$\underline{b} = \inf_b \{b| \Pr(b < B^M < b_2) = 0\}; \text{ and } \overline{b} = \sup_b \{b| \Pr(b_1 < B^M < b) = 0\}.$$

By construction for all $\varepsilon > 0$,

$$\Pr(\underline{b} < B^M < \overline{b}) = 0; \ \Pr(\underline{b} - \varepsilon < B^M \leq \underline{b}) > 0; \text{ and } \Pr(\overline{b} \leq B^M < \overline{b} + \varepsilon) > 0.$$

Consider a bidder j with a valuation V^j where $\underline{b} < V^j < \overline{b}$. Since we are assuming that $\Pr(\underline{b} < B^M < \overline{b}) = 0$, $\Pr(\underline{b} < B^j < \overline{b}) = 0$ as well. Hence either $B^j \leq \underline{b}$ or $B^j \geq \overline{b}$. In addition, neither that $B^j < \underline{b}$ nor $B^j > \overline{b}$ can happen on the equilibrium. To see this:

(i) Compare expected payoffs of a strategy inducing $B^j = V^j$ with a strategy inducing $B^j = \underline{b} - \varepsilon$, $\forall \varepsilon > 0$. If $\underline{b} - \varepsilon < B^M < V^j$, which happens with a positive probability, $B^j = V^j$ results in the positive payoffs, $V^j - B^M$, but $B^j = \underline{b} - \varepsilon$ results in the zero payoff. If $B^M = \underline{b} - \varepsilon$, $B^j = V^j$ always results in the positive payoffs, $V^j - B^M$, but $B^j = \underline{b} - \varepsilon$ results in the positive payoffs, $V^j - B^M$, or the zero payoff. For other range of B^M, both strategies yield the same payoffs. Accordingly, a strategy inducing $B^j = V^j$ yields higher expected payoffs than any strategy inducing $B^j = \underline{b} - \varepsilon$, $\forall\, \varepsilon > 0$.

(ii) Compare expected payoffs of a strategy inducing $B^j = V^j$ with a strategy inducing $B^j = \bar{b} + \varepsilon$, $\forall \varepsilon > 0$. If $V^j \leq B^M < \bar{b} + \varepsilon$, which happens with a positive probability, $B^j = V^j$ results in the zero payoff, but $B^j = \bar{b} + \varepsilon$ results in the negative payoff, $V^j - B^M$. If $B^M = \bar{b} + \varepsilon$, $B^j = V^j$ results in the zero payoff, but $B^j = \bar{b} + \varepsilon$ results in the negative payoff, $V^j - B^M$, or the zero payoff. For other range of B^M, both strategies yield the same payoff. Accordingly, a strategy inducing $B^j = V^j$ yields higher expected payoff than any strategies inducing $B^j = \bar{b} + \varepsilon$, $\forall \varepsilon > 0$.

Now we know that B^j should be either $\underline{b}$ or $\bar{b}$ on the equilibrium if $\underline{b} < V^j < \bar{b}$. That means B^j has a point mass at $\underline{b}$, $\bar{b}$, or both, and so should B^M. However if B^M has a point mass at $\underline{b}$, $B^j = \underline{b}$ cannot happen on the equilibrium; and if B^M has a point mass at $\bar{b}$, $B^j = \bar{b}$ cannot happen on the equilibrium. To see this:

(i) Suppose B^M has a point mass at $\underline{b}$. Compare expected payoffs of a strategy inducing $B^j = V^j$ with a strategy inducing $B^j = \underline{b}$. If $B^M = \underline{b}$, which happens with a positive probability by assumption, bidder j always wins if $B^j = V^j$, but bidder j wins with a probability less than 1 if $B^j = \underline{b}$. For other values of B^M, both strategies yield the same payoff. Hence a strategy inducing $B^j = V^j$ yields higher expected payoff than any strategy inducing $B^j = \underline{b}$.

(ii) Suppose B^M has a point mass at $\bar{b}$. Compare expected payoffs of a strategy inducing $B^j = V^j$ with a strategy inducing $B^j = \bar{b}$. If $B^M = \bar{b}$, which happens with a positive probability by assumption, bidder j always loses if $B^j = V^j$, but bidder j wins with a positive probability to obtain a negative payoff if $B^j = \bar{b}$. For other values of B^M, both strategies yield the same payoff. Hence a strategy inducing $B^j = V^j$ yields higher expected payoff than any strategy inducing $B^j = \bar{b}$. Therefore B^j doesn't have a point mass at $\underline{b}$ nor $\bar{b}$; this implies that $\Pr(B^j \leq \underline{b}) = 0$ and $\Pr(B^j \geq \bar{b}) = 0$.

In summary, if we assume that B^M does not have a full support over $[\underline{v}, \bar{v}]$, there exist $\underline{b}$ and $\bar{b}$ such that $\Pr(\underline{b} < B^j < \bar{b}) = 0$, $\Pr(B^j \leq \underline{b}) = 0$, and $\Pr(B^j \geq \bar{b}) = 0$; this is a contradiction. Hence we can conclude that there is no b_1 and b_2 such that $\Pr(b_1 < B^M < b_2) = 0$ where $\underline{v} < b_1 < b_2 < \bar{v}$.

(2) Here I show that a strategy inducing $B^i = V^i$ yields higher expected payoff than a strategy inducing $B^i \neq V^i$. Then the necessity of (b) follows. Below I compare payoffs among

strategies inducing (i) $B^i = V^i - \varepsilon$; (ii) $B^i = V^i$; and (iii) $B^i = V^i + \varepsilon$ for all $\varepsilon > 0$:

Event	(i) $B^i = V^i - \varepsilon$	(ii) $B^i = V^i$	(iii) $B^i = V^i + \varepsilon$
$B^M < V^i - \varepsilon$	$V^i - B^M$	$V^i - B^M$	$V^i - B^M$
$B^M = V^i - \varepsilon$	0 or $V^i - B^M$	$V^i - B^M$	$V^i - B^M$
$V^i - \varepsilon < B^M < V^i$	0	$V^i - B^M$	$V^i - B^M$
$B^M = V^i$	0	0	0
$V^i < B^M < V^i + \varepsilon$	0	0	$V^i - B^M$
$B^M = V^i + \varepsilon$	0	0	0 or $V^i - B^M$
$B^M > V^i + \varepsilon$	0	0	0

Since B^M has a full support over $[\underline{v}, \overline{v}]$, $\Pr(V^i - \varepsilon < B^M < V^i) > 0$ where a strategy inducing $B^i = V^i$ yields higher payoff than a strategy inducing $B^i = V^i - \varepsilon$. Similarly, $\Pr(V^i < B^M < V^i + \varepsilon) > 0$ where a strategy inducing $B^i = V^i$ yields higher payoff than a strategy inducing $B^i = V^i + \varepsilon$. Accordingly a strategy inducing $B^i = V^i$ yields higher expected payoff than a strategy inducing either $B^i = V^i - \varepsilon$ or $B^i = V^i + \varepsilon$.

B. The necessity of (a) $C^i(t) \leq v^i, \quad \forall\, t \in [0, \tau]$

It is straightforward from the necessity of (b). Since an eBay bidder may not decrease his previous cutoff-price, any strategy inducing $c^i_t > v^i$ cannot equalize B^i with v^i. ∎

The idea of necessity proof of Proposion 1 is similar to that of Proposition 1 of Blume and Heidhues (2004). It looks more complicated mainly because we incorporate bidding times in the model. Proposition 1 shows that, although there are multiple equilibria with respect to the bidding time, in every equilibrium, all bidders submit their true valuations before the auction ends as long as the standing price has not raised over their valuations when they place a bid. Note that eBay's tie breaking rule favorable to early bidder cannot give an incentive to bid early. As long as everyone bid his valuation, an eBay bidder will pay as much as his valuation whenever the tie breaking rule applies. As a result, the two highest-valued potential bidders, whose valuations cannot be lower than the standing price at any time, always bid their valuations before the auction ends. On the other hand, lower-than-second-highest-valued bidders will bid their valuations only if they choose to do so before the standing price rises above their

valuations. Therefore, some potential bidders may not make a bid at all or may not update their early cutoff prices, even though these were lower than their valuations. This could be critical to interpretation of eBay bidding data because a set of observed bids may not be the same as a set of potential bidders' bids.

The proof of Proposition 1 demonstrates that every equilibrium is an *ex-post* equilibrium:[4] even if the actual number of potential bidders n and all potential bidders' private information $\{v^i, T^i\}_{i=1}^n$ were known to a particular bidder i, his equilibrium strategy would still be optimal. This suggests robustness to changes in assumptions that a bidder knows distribution $F(\cdot)$ and the probabilities p_n.

I will briefly discuss the relationship of this model to some stylized facts about eBay auctions, and to other existing models of eBay auctions. Previous research concerning eBay auctions such as Bajari and Hortaçsu (2002a), Ockenfels and Roth (2002), and Roth and Ockenfels (2002) has pointed out that late-bidding is prevalent. The above result does not contradict late-bidding. For example, in one equilibrium, all bidders wait until their own, last monitoring times to submit any cutoff prices. However, this model does not explain why late-bidding is observed more frequently than early-bidding: here, bidders have no reason to bid late, but also no reason not to.

Bajari and Hortasçu (2002a) study eBay auctions within the common value paradigm. They show that on an equilibrium, bidders will bid at the end of the auction in order not to reveal their private information to other bidders in a common value environment. Ockenfels and Roth (2002) construct a model which, like mine, has multiple equilibria, including one involving last-minute bidding in private value environments. However, in their model, on an equilibrium path in which the last-minute bidding happens, every bid in an auction should be submitted in the auction's last seconds; that is hardly ever observed in practice. The strength of Proposition 1 is that it characterizes all symmetric Bayesian-Nash equilibria. In Bajari and Hortasçu (2002a) and Ockenfels and Roth (2002), they show only that their equilibria is an equilibrium.

---

[4] For a definition of ex post equilibrium, see Appendix F of Krishna (2002) or references therein.

References

Bajari, Patrick and Ali Hortaçsu (2002a), "Winner's curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auction, " *RAND Journal of Economics*, Summer 2003, 329-355.

Bajari, Patrick and Ali Hortaçsu (2002b), "Cyberspace Auctions and Pricing Issues: A Revies of Empirical Findings," *New Economic Handbook*, Edited by Derek C. Jones.

Blumea, Andreas and Paul Heidhuesb (2004), "All equilibria of the Vickrey Auctions," *Journal of Economic Theory*, 114, 170-177.

Krishna, Vijay (2002), *Auction Theory*, San Diego: Academic Press.

Ockenfels, Axel and Alvin E. Roth (2002), "Late Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction," Working Paper, Harvard University.

Roth, Alvin E. and Axel Ockenfels (2002), "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92 (4), 1093-1103.