Dynamic Bidding in Second Price Auctions

Hugo Hopenhayn* **Maryam Saeedi†**

UCLA The Ohio State University

December 21, 2015

Abstract

We develop a dynamic model of biddings in second price auctions where agents' bidding opportunities and values follow a joint Markov process. We prove that equilibrium exists and is unique, providing a recursive representation and algorithm to solve for bids as a function of time and values. The equilibrium bid equals the expected final value conditional on being the bidder's final one: either there is no further rebidding opportunity or the bidder chooses not to increase this bid if given the option. This results in adverse selection with respect to a bidder's own future values, and as a result bids are shaded. This is true in spite of values being independent across bidders. Under mild conditions, desired bids increase as time increases and the close of the auction is approached. Our results are consistent with repeated bidding and sniping, two puzzling observations in eBay auctions. We estimate the model by matching moments from eBay auctions and consider a series of counterfactuals.

*hopen@econ.ucla.edu

†saeedi.2@osu.edu

...if something went over my limit early on, I might re-evaluate my budget and make a higher bid because I want the product. If it was the last few minutes I wouldn't have the time to consider if I can afford it.

...work + college = does not give you the right time you need to baby an auction.

The wife wanted to go to lunch right at the same time as the end of the auction, so I decided to drop an early bid an hour before the close.

1 Introduction

Various auction mechanisms happen in a dynamic setting while most of the theoretical models of auctions are static. For example, on eBay, auction listings usually last seven days. Bidders can bid at anytime during the active time of the auction and also increase their bids at any time prior to the end. As the first quote suggests, information might arrive during the auction that can change a bidder's value.[1] As the second and third quotes suggest, bidders are often inattentive and not have full control of the timing of future bids.[2] In this paper, we develop a general model of a dynamic second price auction that captures these two features.

In our model, agents' bidding opportunities and values follow a joint Markov process. At each time of arrival bidders observe a new signal and choose whether to place a bid or increase a previous one. At the end of the auction, the winner is the highest bidder and the price equal to the second highest bid. We characterize the equilibrium of this dynamic auction.

While our theory applies generally to dynamic second price auctions, it is motivated by bidding behavior in eBay auctions, which are close to second price auctions, where bidders can place bids at any point in time up until the end of the auction. Several authors have emphasized what seem to be anomalies in bidding behavior, such as submission of multiple bids throughout the auction and sniping, i.e. a higher concentration of bids towards the end

[1] The first two quotes were obtained from a blog http://www.neowin.net/forum/topic/587154-why-do-people-bid-so-early-on-ebay/. The second one is from http://www.kenrockwell.com/tech/ebay/early-bidding.htm.

[2] Sniping programs are an incomplete solution as they do not condition bids on possible changes of valuation that could arise from new information.

of the auction, as we document below. This type of behavior cannot be rationalized through the lens of a static model, where bidders should bid their valuation only once and at the time of arrival to the auction. Our model provides a way to rationalize this behavior combining two ingredients: new information can change the optimal bid throughout the auction and the possible lack of a future opportunity to bid provides a rationale for early bidding.

In most of the paper, we consider the case of independent private values, where signals and final values are independent across bidders, and where bidding times are exogenous.[3] We model this through a–otherwise unrestricted–joint Markov process of signals and bidding times that are also sufficient statistics for the final expected value. Our process is very general, allowing for instance that attentiveness to vary with value and nests the case where a bidder bids at the end of the auction almost surely. We prove equilibrium is unique, providing also a recursive representation and algorithm to solve for bids as a function of time and values.

Bidding behavior has an intuitive analog to that of a standard second price auction. Consider first the case where bidders have a single opportunity to bid prior to the end of the auction. In that case, they bid their expected final value as in a standard second price auction. In contrast, when bidders can return to the auction prior to its end with positive probability and rebid, the expected value is modified taking into account that the current bid will apply if the bidder chooses not to exercise that option. The equilibrium bid equals the expected final value conditional on being the bidder's final one: either there is no further rebidding opportunity or the bidder chooses not to increase this bid if given the option.

The possibility of rebidding is a source of *adverse selection* against future self as the current bid applies when the bidder chooses not to exercise future options of rebidding and this is correlated with lower future value.[4] As a consequence, bidders shade bids below their unconditional expected final value. In particular, as the probability of bidding at the very end of the auction converges to one, any prior bid goes to zero while at the same time the maximum bid converges to the value at the end of the auction, consistent with sniping

[3]Section 11.1 analyzes a simplified model with endogenous bidding and Section 11.2 analyzes a simplified model with correlated values among bidders.

[4]An analogue result is found in Harris and Holmstrom [1982], where initially worker's wages are shaded below marginal products, as the wage is effective in the future only if it is less or equal than the realized marginal product of the worker.

behavior.

There is an important difference to the case of adverse selection with correlated values. While bids are shaded relative to expected values, the *expected final bid* still equals the expected final value and is thus unbiased. These two results are consistent since bids, while shaded, are still unbiased in the set where they *apply.* As a comparison, an auction that admits bidders to retract on bids will have no shading, since a bid only applies when no rebidding takes place or when the final expected value is the same at a new bidding time.

The problem of solving for the optimal bidding function is not a simple one, as current bids depend on all future bidding opportunities and strategies. We provide a recursive formulation for bidders optimal strategy. We can prove that the solution to the dynamic programming exist, is unique and under mild conditions, desired bids increase as the end of the auction is approached. This provides a rationale why bidders might increase their bids over time independently of competitive pressures. The intuition for this result is straightforward and goes back to the incentives for shading: as the end of the auction approaches and the option of rebidding becomes less likely, the adverse selection problem mentioned above and the incentives for shading tend to disappear.

To derive further properties, we consider the specialized case where values follow a Brownian motion and bidding opportunities a Poisson process,.We derive a partial differential equation that is used to solve for the equilibrium bidding function. Matching a series of moments from eBay data, we estimate parameter values for the Brownian motion and Poisson process.We then consider as a counterfactual, a similar bidding mechanism to our baseline, where instead bidders can retract bids at any of their random bidding times. The possibility of retraction eliminates the source of adverse selection mentioned earlier and thus the incentives for bid shading. As a consequence, at any bidding time the bidder will choose a bid equal to the expected final value. We prove that the distribution of final bids for a bidder is a mean preserving spread of those obtained with no retraction. In addition, as a reference point, we compute the efficient allocation and prices, i.e. that corresponding to a second price auction taking place at the end of period. The correlation between final values and final bids increases moderately (from 0.85 to 0.88) when comparing the baseline scenario to the case of bid-retraction. As a result the expected final value increases, closing 1/3 of a fairly moderate gap of 3.2% between the baseline model and the efficient allocation. While the gap in values is moderate, the gap in prices between baseline and the efficient level is more than twice as high. This is precisely the consequence of bid shading in our baseline scenario. As a comparison, the price gap is in the order of only 2% for the bid retraction case. As a consequence, allowing for bid-retraction would increase revenues of sellers by approximately 10.2% while decreasing the expected utility of buyers by 33%.

Related Literature Several papers have tried to rationalize sniping behavior and incremental bidding. Bajari and Hortacsu [2003] incorporate a model of common value auction to explain this phenomenon. Having informed and uninformed participants, the informed bidders do not bid before the last period since that would reveal their private information to other potential bidders. Hence in the equilibrium, all bidders only place a bid at the very last period of the auction, leading to sniping by everyone. In another set of papers that focus on sniping, Ockenfels and Roth [2002] and Ockenfels and Roth [2006] compare auctions mechanisms at eBay and Amazon which have hard ending and soft ending, respectively. Hard ending refers to a fixed ending time for an auction which cannot be extended by the seller or the marketplace; whereas soft ending refers to a tentative ending time: placing a bid within the last few minutes of the auction extends the duration of the auction for another ten minutes. They argue that along with the hard-ending rule at eBay, there are incremental bidders on eBay whom increase their maximum bid as they get outbid by other bidders. Therefore, a strategic bidder facing incremental bidders places bids in the last possible moment not to give incremental bidders time to react and place higher bids, giving rise to sniping.

Gray and Reiley [2004] and Ely and Hossain [2009] independently run experiments evaluating the benefit of sniping. The former does not find a statistically significant value to sniping although they reconfirm results in Ockenfels and Roth [2002] and Ockenfels and Roth [2006] regarding the prevalence of sniping; in their dataset 50 percent of bids are placed within the last minute of an auction. On the other hand, the latter find a significant value to sniping, about $1 per auction for various DVD listings. Backus et al. [2015] in a more recent paper consider the effect of sniping on the return rate of new buyers to the marketplace. In all of these papers, the concentration is mainly on rationalizing sniping, and neither early bidding nor bidding multiple time is rationalized alongside sniping.

In this paper, we concentrate mainly on dynamics within the span of an auction. However, we are borrowing the result from Zeithammer [2006], Said [2011], Hendricks and Sorensen [2014], Backus and Lewis [2012], and Coey et al. [2015]implicitly. These papers model the dynamic option value of not winning the current auction and the opportunity cost of participating in the next available auction. On eBay, many closely substitutable items are simultaneously available and if bidders do not win a particular auction they can participate in the next closing one. This results in a reservation price below their values. Changes in the available alternative items can change the reservation price for bidders over time, and as they get closer to the end of the auction their valuation of the item at that instant becomes closer to their valuation at the end of the auction. While our paper misses some of the interesting features arising from the explicitly modeling the link between these auctions, it provides a very tractable reduced form.

In a related paper, Ambrus et al. [2013] model gradual bidding on eBay-like auctions. They model bidders as having the same common value for the item and with random bidding opportunities. They show the equilibrium bidding strategy is to increase the price with the minimum increment and the person who showed up last will be the winner. Groeger and Miller [2015] , in a similar set up, study the optimal bidding strategy of first price auction when the bidders have random bidding windows.

The paper is organized as follows. Section 2 discusses the evidence on eBay bidding behavior and reviews related literature. Section 3 provides a simple example that conveys the main intuition and results in the paper. Section 4 describes the general model and defines an equilibrium. Section 5proves existence and uniqueness of equilibrium, characterizes bidding behavior and provides the dynamic programming algorithm. Section 6 gives properties for the case where values are independent of Poisson arrivals for bidding times and solves the two special cases described above. Section 8 shows that the bidding function derived before is still applicable under arbitrary assumptions about the information a bidder observes on the past bids of other bidders. Section 9 provides the estimates and counterfactuals. Sections 11.1 and 11.2 consider, respectively, the cases of endogenous bidding and correlated signals.

2 Evidence

The model presented in this paper closely mimics eBay auctions. eBay is an online auction and shipping website launched in 2005. Sellers can sell their items either through an auction or by setting a fixed price for their item, an option called "Buy it Now." The auction mechanism is similar to a second price or Vickery auction. A seller sets the starting bid of an auction and bidders can bid repeatedly for the item until the end of the auction. Each bidder observes all previous bids, except for the current highest bid. A bidder should bid an amount higher than the current second highest bid, plus some minimum increment.[5] If this value is higher than the current highest bid, the bidder becomes the new highest bidder. Otherwise, he becomes the second highest bidder. The winner has to pay the second highest bid, plus the increment or his/her own bid, whichever is smaller. Auctions last for one to ten days and they have a pre-determined and fixed ending time that cannot be changed once the auction is active. As a general rule, bidders cannot retract or cancel a bid.

On eBay, bidders' bid-placement, as noted in the literature, does not follow the prediction of a static model of auction, for instance a disproportional share of bids are placed in the last few seconds of an auction. As a recent paper by Backus et al. [2015] shows, about third of winning bids are placed in the last ten seconds of an auction. Hayne et al. [2003] report that about 15 percent of bids placed are within the last 60 minutes of an auction and also about 61 percent of bids placed by bidders who submitted more than one bid. Gonzalez et al. [2009] similarly report 11 percent of bids placed within the last minute of auction and 77 percent of auctions have a bidder placing more than one bid. Hayne et al. [2003] further note that on average there are 6.78 bids placed and about 3.98 unique bidders per auction. Moreover, they show that late bidding has a much higher success rate, about 75 percent, much more than early bidding, 7.3 percent, or bidding in between, 40.5 percent. Using a static model of auctions, strategic late bidding or bidding multiple times is not rational.

In order to get a deeper understanding of bidding behavior, we use eBay data on successful auction sales during the first week of June, 2014. There are an average of 2.93 bidders per auction and an average of of 2 bids per bidder. We focus on three aspects of bidding behavior

[5]The increment is a function of the second highest bid, fixed for all auctions, and is set by eBay.

that are more closely related to our theory: 1) distribution of bidding time; 2) frequency of bidding 3) the process for the increase in bids of repeat bidders.

Figure 1 gives the distribution of all bidding. While there is some bias towards later bidding, its not excessively strong: 55% of the bids are placed in the last quintile of time and about 70% after the midpoint of the auction. When looking at the distribution of winning bids, the concentration towards the end is considerably larger, as should be expected. As can be seen in Figure 2, over 75% of winning bids take place in the last quintile and a very large concentration of these bids occur in the last few hours of the auction.

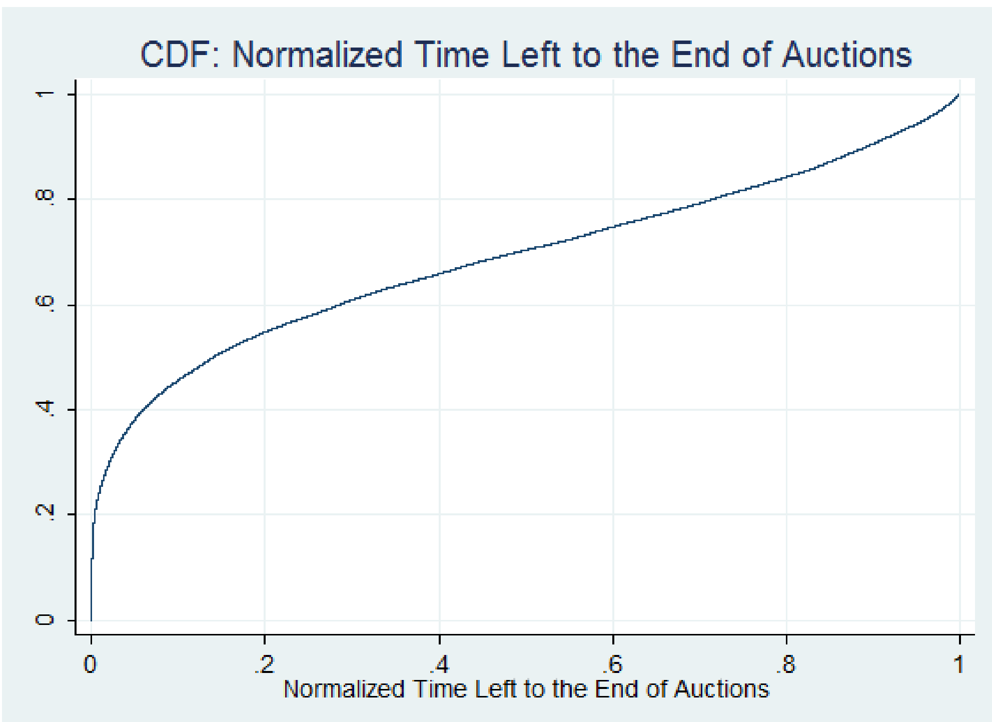

Figure 1: Distribution of bidding times

Figure 3 gives the cumulative distribution of number of bids each bidder submitted during any auction where he participated. Let $i = 1, ..., I$ be bidders and $j = 1, ..., J$ denote auctions and N_{ij} the number of original bids (not including proxy bidding) submitted by bidder i in auction j. This graph is the cdf of all positive N_{ij} values.[6] It can be seen that more than 30% of bidders place more than one bid while 15% over two. The average number of bids per bidder in our data is slightly less than two.

[6]For display purposes we truncate this distribution at 10 bids per bidder, while the maximum we observe in our data is 80.

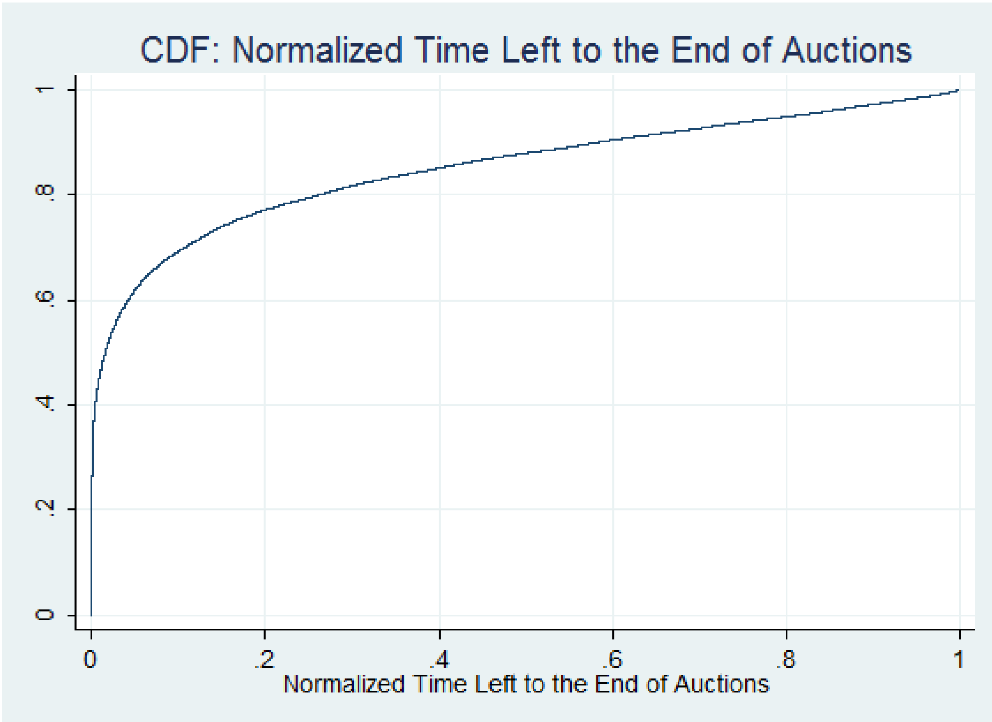

Figure 2: Distribution of time of winning bids

Our theory focuses precisely on the incentives for rebidding and the shading of bids. It will become useful to identify the parameters in our model to document more fully these bidding patterns. Figure 4 gives the ratio of second to first bids for those bidders that place two bids in an auction, as a function of the time difference between these two bids. Each point is an average across auctions for that particular time interval of one minute. As seen, the ratios are quite high, and increase with the time difference. The fanning out is a result of the increase in variance of the ratio of rebids over time. In order to get a more transparent picture, identical data is aggregated into 100-minute in Figures 5 and 6. Mean bid ratio double when going to the smallest to longest time interval between bids, while the variance quadruples.

3 A simple example

This section considers a simple example and derives some key properties of the equilibrium bidding function that extend to the general case. There are two periods $t = \{0, 1\}$. Bidders

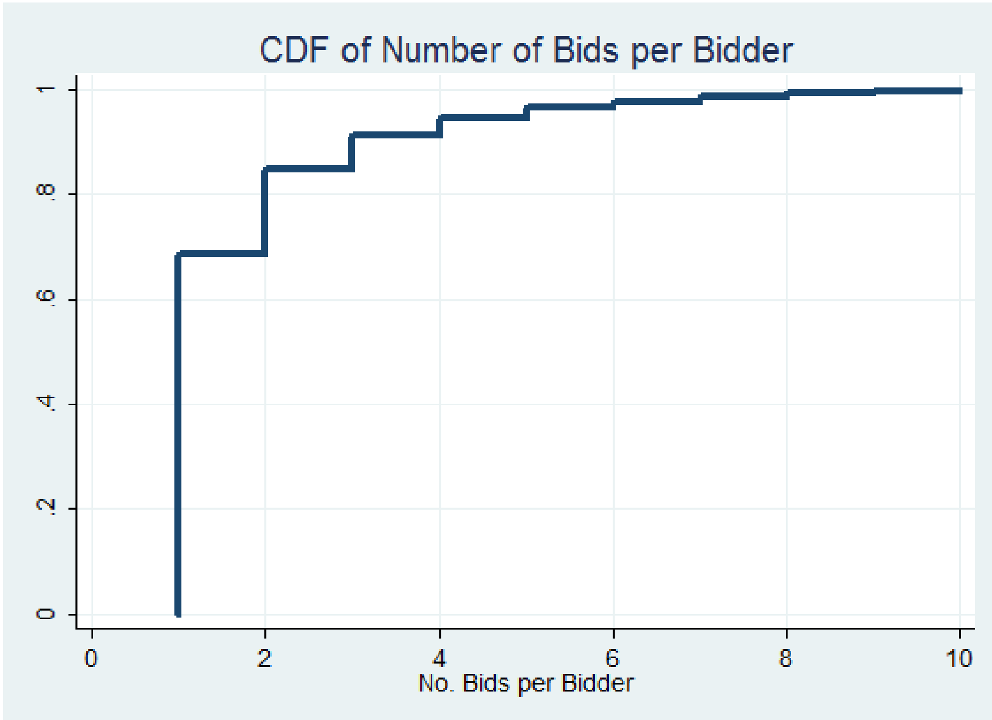

Figure 3: Number of bids per bidder

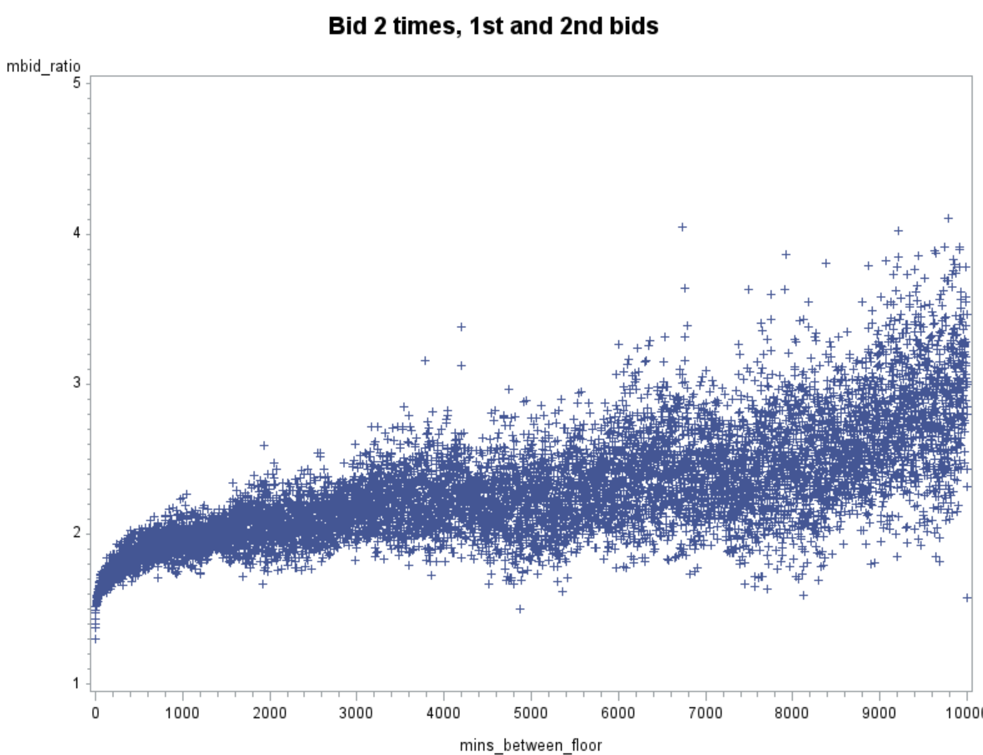

Figure 4: Average re-bid ratio by time elapsed (minutes)

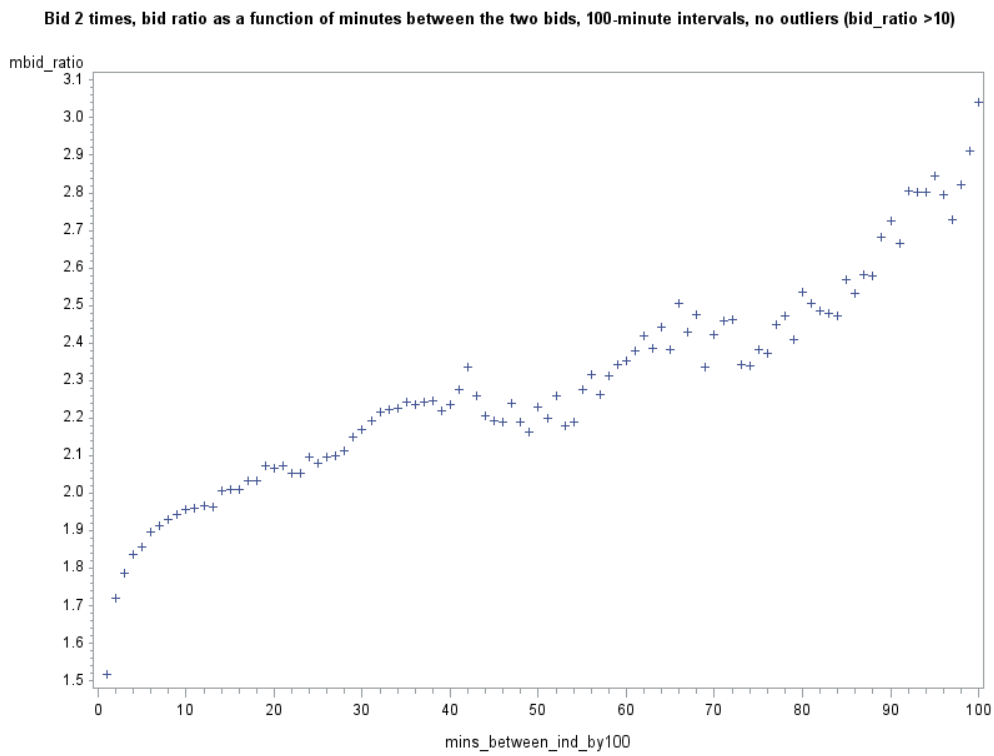

Figure 5: Average re-bid ratio by time elapsed (100-min intervals)

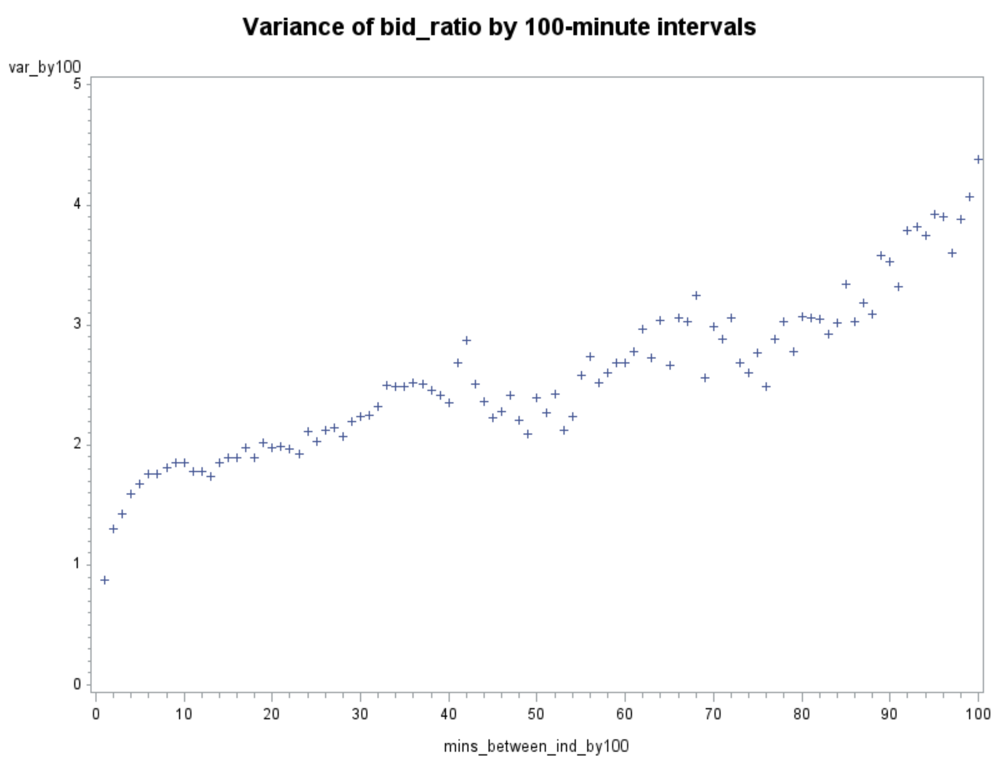

Figure 6: Variance of re-bid ratio

can submit a bid for sure in the first period and with probability p in the second period. Bidders have no information in the first period and draw values $v\epsilon\left[0,1\right]$ independently from distribution F in the second period, prior to bidding time. Since the auction is second price, it follows that the second period bid will equal v_i for all bidders. Let b_{-i} denote the maximum final bid among all bidders, excluding i and let G denote its cdf. The first period, expected value for bidder i is:

$$\begin{aligned}
U_i &= (1-p)\int_0^{b_0}\left(Ev-b_{-i}\right)dG\left(b_{-i}\right)\\
&+ p\int_0^{b_0}\int_0^{b_0}\left(v-b_{-i}\right)dF\left(v\right)dG\left(b_{-i}\right)+p\int_{b_0}\int_0^{v}\left(v-b_{-i}\right)dF\left(v\right)dG\left(b_{-i}\right)
\end{aligned}$$

Taking derivate with respect to b and equating to zero:

$$i\left(1-p\right)\left(Ev-b_0\right)dG\left(b_0\right)+p\int_0^{b_0}\left(v-b_0\right)dG\left(b_0\right)dF\left(v\right)=0 \tag{1}$$

It can be easily verified that as this equation is strictly decreasing in b_0 so there is a unique solution that satisfies:

$$b_0=\frac{\left(1-p\right)Ev+p\int_0^{b_0}vdF\left(v\right)}{\left(1-p\right)+pF\left(b_0\right)} \tag{2}$$

This expression has a natural interpretation: the optimal initial bid equals the expected value conditional on no rebidding. Given a current bid b_0, the final bid prevailing at the end of the auction is then b_0 with probability $\left(1-p\right)$ and $\max\left(b_0,v\right)$ with probability p. The initial bid b_0 binds in two cases: 1) there is no opportunity to rebid and 2) there is an opportunity to rebid but $v\leq b_0$ therefore the bidder decides not to bid. The second term in the above equation represents an adverse selection effect. It follows that $b_0<E\left(v\right)$ as $\int_0^{b_0}vdF\left(v\right)/F\left(b_0\right)<E\left(v\right)$. Indeed, as $p\rightarrow 1$ it is easy to see that b_0 decreases monotonically to zero.

The probability distribution $\tilde{F}$ for the final bids of a player is given by:

$$\tilde{F}\left(b\right)=\begin{cases}0 & \text{if } b<b_0\\ \left(1-p\right)+pF\left(b\right) & \text{if } b\geq b_0\end{cases} \tag{3}$$

with mean

$$E(b) = b_0 [p + (1-p) F(b_0)] + p \int_{b_0}^{1} v dF(v)$$

Substituting the first term using (2) it follows that:

$$E(b) = (1-p) Ev + p \int v dF(v) = Ev$$

so the expected final bid is unbiased, equaling the expected final value.

It is also interesting to compare the above results to a case where bidders can retract their bids. It is easy to se that with retraction there is no incentive to shade, so b_0 is equal to Ev; the final bid equals Ev with probability $(1-p)$ and equals the final value with probability p. The distribution for final bids is thus:

$$\tilde{F}(b) = \begin{cases} pF(b) & \text{if } b < Ev \\ (1-p) + pF(b) & \text{if } b \geq Ev \end{cases}$$

Compared to (3) it is easy to see that this is a mean preserving spread of bidding with no retraction. The higher dispersion is the simple result of a higher correlation with final values.

4 The Model

There are i=1,...,N potential bidders in an auction. The auction is sealed bid second price and takes place in time interval $[0, T]$ where bids are submitted. As shown below in Section 8 equilibrium bidding functions are also applicable to sequential auctions -such as in eBay- with publicly available information on past bids.

Each bidder has the option of submitting bids only at random times $\tau_1, \tau_2, ...$ Bids can only be increased at any of these random bidding times and cannot be retracted. The valuation of the bidder is modeled as a stochastic process $v_i(t)$ of signals where without loss of generality $v(T)$ corresponds to the final value. Since these signals are only relevant at bidding times τ_n we restrict attention to the corresponding signals v_n at these dates.[7] Assume

[7]The process for the value $v_i(t)$ can be considered a continuous Markov process sampled at random stopping times τ_n that are Markov with respect to last stopping time and valuation at that time.

$\{v_{in}, \tau_{in}\}$ follow a joint Markov process with transition function which is independent across bidders[8], and inscribed in a common probability space (Ω, σ) with typical element ω. Note that this specification allows for random entry and rebidding, as well as a random number of bidders. It also allows for the degree of inattention to depend on value and accommodates the case where a bidder is sure to come at the end of the auction, the latter by letting $P_i(\tau' \leq T|v, \tau) = 1$ for all v, τ. From now on for notational convenience we drop the i index.

We consider here the case of a sealed bid auction, where bidders information sets consist only of their past bids and valuations. In section 8 we show that the equilibrium bidding functions obtained here are also equilibrium bids for general information structures and in particular for that of eBay auctions. Together with the assumption that values and bidding times are Markov, this implies that we can identify information sets with pairs (v, τ) corresponding to the time and value in the last bidding window. A bidding function for bidder i specifies at each possible bidding time τ_n and given signal v_n a desired bid $B_i(\tau_n, v_n)$. Given that bids can only be increased, $b_i(t) = \max\{B_i(\tau_m, v_m)|\tau_m \leq t\}$ is the bid that prevails at time t and in particular $b_i(T)$ is the final bid. Let $b_{-i}(T)$ denote the maximum bid over the remaining bidders at time T and $F_{-i}(b)$ its distribution. Utility for bidder i is given by:

$$U(v(T), b(T)) = \int^{b(T)} (v(T) - u)\, dF_{-i}(u).$$

An optimal bidding function B_i for bidder i is the one that maximizes expected utility at the time the bidder enters the auction, i.e.,

$$\max_{B(.)} E(U(v(T), b(T))|\tau_1, v_1) \tag{4}$$

where τ_1 is the time at which the bidder enters the auction and v_1 the initial value.

Definition. An *equilibrium* for the auction is a vector of bidding functions B_i and final distributions F_{-i} for each player i, such that for every bidder i the bidding function B_i is the best responses to F_{-i} and bidding functions are consistent with the final distribution of

[8]As shown in Section 11.2, results extend to the case of common values

bids.

5 Equilibrium Bidding

In this section, we characterize the equilibrium bidding functions and provide a dynamic programming problem that can be used to solve them. We show that the optimal bid $B(v_n, \tau_n) = E[v(T)|v_n, \tau_n, b(T) = B(v_n, \tau_n)]$, namely the expected final value for bidder i conditional on the current state and the event that no higher bids are placed by the same bidder at a later instance so that the final bid $b(T)$ equals the current bid. Intuitively, this mimics the notion that in a second price auction it is weakly dominant strategy to bid the valuation, or when it is random, the expected valuation.

Proposition 1. *The optimal bid* $B(v_n, \tau_n) = E[v(T)|v_n, \tau_n, b(T) = B(v_n, \tau_n)]$.

Proof. See Appendix. □

We provide an intuitive argument here, while the formal proof is in the Appendix. Take a candidate optimal bidding function B_i for bidder i and consider bid b in state v_n, τ_n. Let $H(b)$ denote all paths $\omega = \{\tau_{m,} v_m\}_{m>n}$ starting from the current history where $B(\tau_m, v_m) \leq b$ including $\tau_{n+1} > T$ (i.e. no rebidding opportunity.) The expected value of bidding b equals:

$$V(v_n, \tau_n, b) = \int_{H(b)} \int^{b} (v(T) - u)\, dF_{-i}(u)\, dP(\omega) + \int_{H(b)^c} \int^{b(\omega,T)} (v(T) - u)\, dF_{-i}(u)\, dP(\omega)$$

We claim that the optimal bid is $b = B(v_n, \tau_n) = E_{H(b)} v(T)$. First, note that the boundary of the set $H(b)$ consists of all those paths starting from (v_n, τ_n) for which the final bid $b(T)$ is equal to b. So when considering the derivative of the above we can ignore the effect of the change in the supports of the two integrals. The first order condition is then $\partial V / \partial b = \int_{H(b)} (v(T) - b)\, dF_{-i}(b)\, dP(\omega) = 0$ which is equivalent to the statement that $b = E_{H(b)} v(T)$.

While Proposition 1 characterizes bidding at a given history, it also shows that current bidding behavior depends on the whole strategy for future bidding, making the problem of calculating equilibrium bids potentially very complicated. However, there is a natural recursive structure to this problem which we exploit to define a dynamic programming problem that will help derive the optimal bidding function and establish uniqueness.

Define recursively the following function[9]:

$$W(b,v,\tau) = \int_{\tau}^{T} \min\left(W\left(b,v',\tau'\right),0\right) dP\left(v',\tau'|v,\tau\right) + P\left(\tau' > T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau' > T\right] - b\right) \tag{5}$$

where τ' denotes the following bidding time.

Assumption 1. *Assume $P(\tau' > T|v,\tau) > \delta > 0$ for all (v,τ). In addition, assume all conditional probabilities and expectations are continuous in the state.*

By Assumption 1 and using standard dynamic programming arguments, it follows that there is a unique function satisfying functional equation (5), that it is strictly decreasing in b and continuous. Moreover, it is greater or equal to zero when $b = 0$ and negative for large b. It follows, by the intermediate value theorem, that there is a unique value $B(v,\tau)$ such that $W(B(v,\tau),v,\tau) = 0$. Given a new bidding time τ', value $v(\tau')$ and outstanding bid, this function also defines the rebidding region $\{(b, v(\tau'), \tau')\,|W(b, v(\tau'), \tau') > 0\}$ We will next show that this bidding function maximizes the agents expected utility.

Proposition 2. *Given Assumption 1, the function $B(v,\tau)$ defined implicitly by $W(B(v,\tau),v,\tau) = 0$ satisfies $E[v(T)\,|v,\tau,b(T) = B(v,\tau)]$ and is the unique optimal bidding function.*

Proof. See Appendix. □

6 Properties

This section derives general properties on equilibrium bidding behavior and considers some special cases.

6.1 Bid Shading

Our leading example suggests that bidders will shade bids as a consequence of adverse selection given by the option of future rebidding. In this section we prove this is true for an

[9]When conditioning with respect to τ and $v(\tau) = v$ we will write for short $P[.|v,\tau]$.

arbitrary process. We first establish the following result:

Lemma 1. *$W(b,v,\tau) \leq E[v_T|v,\tau] - b$ with strict inequality if $P(\tau' \leq T|v,\tau) > 0$.*

Proof. See Appendix. □

Since $W(b,v,\tau)$ is decreasing in b, it follows that:

Proposition 3. *$B(v,\tau) \leq E[v_T|v,\tau]$ with strict inequality if and only if rebidding occurs with positive probability.*

6.2 Unbiased expected final bids

Bids are shaded as the result of conditioning on the adverse selected set of no-rebidding. However, in the complement of this set the current bid is replaced by a higher one. What is the overall effect on the final expected bid? The answer is given in the following Proposition.

Proposition 4. *Suppose v follows a Martingale. Consider a bidding time τ with value v and bid $b = B(v,\tau)$. Then the expected final bid equals v, i.e. $E(B(T)|v,\tau,b = B(v,\tau)) = v$.*

Proof. Let H_0 consist of all histories ω starting from (v,τ,b) where $B(T) = b$. By Proposition 2 $E(V(T)|H_0) = b$. Now consider all histories ω on the complement H_0^c where rebidding occurs with probability one. Let τ be the stopping time corresponding to the first time where rebidding occurs. Assume, by way of induction, that $E[B(T)|v(\tau),\tau,b = B(v(\tau),\tau)] = v(\tau)$. By the optional stopping theorem it follows that $E(B(T)|H_0^c) = E(v(\tau)|H_0^c) = E(v(T)|H_0^c)$. □

This Proposition implies, in particular, that when a bidder places the initial bid, the expected final bid is unbiased and equals the expected final value. This result is consistent with bid shading since bids, while shaded, are still unbiased in the set where they *apply.*

We now specialize the model here to a case where the process for rebidding times τ_n is Poisson with intensity $\rho(t)$ that is independent of the signals and value $v(T)$. The main result in this section is that the bidding function $B(v,\tau)$ is increasing in t.

Proposition 5. *Assume $E[v_T|v,\tau]$ is weakly increasing in τ. Then $\partial W(b,v,\tau)/\partial\tau \geq 0$ and bid $B(v,\tau)$ increases with t.*

Proof. See Appendix. □

The above result implies that bids increase over time even in the absence of competitive pressure.

6.3 Independent increments

In this section we consider two cases where bidding behavior is simplified: 1) increments independent of the current value v and 2) increments proportional to v. In particular, these apply to the cases where v follows an arithmetic (geometric) Brownian motion (respectively). In both cases we assume that bidding arrival times are given by a homogeneous Poisson process.

Proposition 6. *Assume $P(v' = v + \varepsilon | v, \tau)$ is independent of v for all ε and all τ and arrival rates for bidding are independent of v. Then $W(b + \delta, v + \delta, \tau) = W(b, v, \tau)$.*

Proof. Follows by standard induction argument on the Bellman equation. So assume the function W has this property. Then evaluate:

$$\begin{aligned}
W(b+\delta, v+\delta, \tau) &= \int_t^T \min\left(W(b+\delta, v'+\delta, \tau'), 0\right) dP(v'+\delta, \tau' | v+\delta, \tau) \\
&\quad + P(\tau' > T | v, \tau)\left(E\left[v_T | v+\delta, \tau, \tau' > T\right] - (b+\delta)\right) \\
&= \int_t^T \min\left(W(b, v', \tau'), 0\right) dP(v', \tau' | v, \tau) \\
&\quad + P(\tau' > T | v, \tau)\left(E\left[v_T | v, \tau, \tau' > T\right] + \delta - (b+\delta)\right) \\
&= W(b, v, \tau)
\end{aligned}$$

□

The following Corollary simplifies bidding behavior to a simple one dimensional shading function $s(t) = v - B(v, \tau)$ that is independent of v.

Corollary 1. $B(v + \delta, \tau) = B(v, \tau) + \delta$.

Consider now the case where $P(\gamma v', \tau' | \gamma v, \tau) = P(v', \tau' | v, \tau)$. As the next Proposition shows, this implies that $W(\gamma b, \gamma v, \tau) = \gamma W(b, v, \tau)$.

Proposition 7. *Assume* $P(\gamma v', \tau'|\gamma v, \tau) = P(v', \tau'|v, \tau)$. *Then* $W(\gamma b, \gamma v, \tau) = \gamma W(b, v, \tau)$ *and consequently* $B(\gamma v, \tau) = \gamma B(v, \tau)$.

Proof. Follows by standard induction argument on the Bellman equation. So assume the function W has this property. Then evaluate:

$$\begin{aligned} W(\gamma b, \gamma v, \tau) &= \int_t^T \min(W(\gamma b, \gamma v', \tau'), 0)\, dP(\gamma v', \tau'|\gamma v, \tau) + P(\tau' > T|v, \tau)(E[v_T|\gamma v, \tau, \tau' > T] - \gamma b) \\ &= \int_t^T \min(\gamma W(b, v', \tau'), 0)\, dP(v, \tau'|v, \tau) + P(\tau' > T|v, \tau)\gamma(E[v_T|v, \tau, \tau' > T] - b) \\ &= \gamma W(b, v, \tau). \end{aligned}$$

The second property follows immediately from the definition of the bidding function. □

7 Brownian motion and Poisson arrivals

In this section, we introduce a special case which we later take into data for estimation and to run counterfactual studies. Suppose that $v(t)$ follows a Brownian motion with drift μ and volatility σ; and the process for rebidding is Poisson with arrival ρ, as in the previous case. The functional equation (5) is now:

$$\begin{aligned} W(b, v, \tau) &= \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W\left(b, v + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma\varepsilon, \tau'\right)\right) d\Phi(\varepsilon)\, d\tau' \\ &\quad + \exp(-\rho(T - t))(v + \mu(T - t) - b) \end{aligned} \tag{6}$$

Note that for this case, $E[v(T)|v, \tau] = v + \mu(T - t)$ so if the drift $\mu \leq 0$ then $E[v(T)|v, \tau]$ is increasing with t, so by Proposition 5, $B(v, \tau)$ is increasing in t.

Our specification satisfies the condition in 6; therefore, the proposition applies. Specifically, $W(b, v, \tau) = W(0, v - b, t)$. In consequence , we can write the value function $W_\mu(x, t)$ where $x = v - b$,

$$\begin{aligned} W_\mu(x, t) &= \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W_\mu\left(x + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma\varepsilon, \tau'\right)\right) d\Phi(\varepsilon)\, d\tau' \\ &\quad + \exp(-\rho(T - t))(x + \mu(T - t)) \end{aligned} \tag{7}$$

Scaling is another property worth mentioning which is used in estimation section as well. If we scale σ and μ of the Brownian motion by a factor λ, the corresponding W function also scaled proportionally and so are the equilibrium bids.

Lemma 2. $W(\lambda x, t; \lambda\mu, \lambda\sigma) = \lambda W(x, t, \mu, \sigma)$ *and further* $B(\lambda v, \tau; \lambda\mu, \lambda\sigma) = \lambda B(v, \tau; \mu, \sigma)$ *.*

Proof. The proof is by induction, assume the right hand side of (7) satisfies this property. It verifies immediately that $W(\lambda x, t; \lambda\mu, \lambda\sigma) = \lambda W(x, t, \mu, \sigma)$. This also implies that the shading function $x(t; \lambda\mu, \lambda\sigma) = \lambda x(t; \mu, \sigma)$ or equivalently $B(\lambda v, \tau; \lambda\mu, \lambda\sigma) = \lambda B(v, \tau; \mu, \sigma)$. □

The effect of drift on the value function can also be easily determined as shown by the following Lemma.

Lemma 3. *Let* W_μ *denote the value function with drift* μ *and* W *the value function with drift zero. Then* $W_\mu(x, t) = W(x + \mu(T - t), t)$ *.*

Proof. This is proved inductively using:

$$
\begin{aligned}
W(x + \mu(T - t), t) &= \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W\left(x + \mu(T - \tau') + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma\varepsilon, \tau'\right)\right) d\Phi(\varepsilon)\, d\tau' \\
&\quad + \exp(-\rho(T - t))(x + \mu(T - t)) \\
&= \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W_\mu\left(x + \mu(\tau' - t) + \sqrt{\tau' - t}\sigma\varepsilon, \tau'\right)\right) d\Phi(\varepsilon)\, d\tau' \\
&\quad + \exp(-\rho(T - t))(x + \mu(T - t)) \\
&= W_\mu(x, t)
\end{aligned}
$$

□

It also follows that the bid shading function is of the form $x(t) + \mu(T - t)$ where $x(t)$ is the shading function when $\mu = 0$. In what follows we assume $\mu = 0$.

7.1 Partial Differential Equation

In the estimation section, we need to solve the functional equation **??** to find the bidding function. We do that by numerically solving its corresponding PDE. To find the PDE for the Bellman equation (the Hamilton-Jacobi equation), first subtract $W(x, t)$

$$
\begin{aligned}
0 &= \rho \int_t^{t+\Delta} \exp(-\rho(\tau' - t)) \int \left[\min\left(0, W\left(x + \sqrt{\tau' - t}\sigma\varepsilon, \tau'\right)\right) - W(x, t)\right] d\Phi(\varepsilon)\, d\tau' \\
&\quad + \exp(-\rho\Delta)\left[\int W\left(x + \sqrt{\Delta}\sigma\varepsilon, t + \Delta\right) - W(x, t)\right] d\Phi(\varepsilon)
\end{aligned} \tag{8}
$$

Taking derivative with respect to Δ and evaluating at $\Delta = 0$

$$\begin{aligned} 0 &= \rho\left[\min\left(0, W\left(x,t\right)\right) - W\left(x,,t\right)\right] + \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial x^2} W\left(x,t\right) + \frac{\partial}{\partial t} W\left(x,t\right) \\ \rho \max\left(W\left(x,t\right), 0\right) &= \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial x^2} W\left(x,t\right) + \frac{\partial}{\partial t} W\left(x,t\right) \end{aligned}$$

8 Bidding with Partially Observed Competing Bids

The above assumed that the auction is sealed bid, so there is no information on competing bids. However, if bidders can observe other bidders' previous bids, they may use that information strategically. In this section, we show that the previous derived equilibrium will remain valid given any information structure. However, we cannot rule out that there might be other equilibria. Throughout this section, we maintain our assumption of independent values across bidders. We first provide details of the information structure and define an equilibrium. We next show that the equilibrium derived before remains an equilibrium of the extended game.

At any point in time each agent's information set includes a common public history $h(t)$ in addition to the private history $h_i(t)$ of all past realizations $\{\tau_{in}, v_{in}\}$ and bids for $\tau_{in} \leq t$. A strategy s_i for player i specifies at every bidding time τ_{in} and corresponding information set $(h(\tau_{in}), h_i(\tau_{in}))$ a bid b_{in} with the restriction that $b_{in} \geq b_{i,n-1}$ for all $n \geq 2$. Let $v_i(\omega, T)$ denote bidder i's value for the object at the end of the auction, $b_i(T, s, \omega)$ denote the final bid of player i and $b_{-i}(T, s, \omega)$ the highest bid among the remaining bidders, where s is the vector of strategies $s = (s_1, ..., s_N)$. For simplicity of notation, we drop ω from the bidding function. Finally let $u(v_i, b_i, b_{-i}) = \chi_{\{b_i \geq b_{-i}\}}(v - b_{-i})$ denote the final payoff for player i. We can define the final expected payoff for player i under strategy s, after history $h(\tau_{in}), h_i(\tau_{in})$, by:

$$E_{\tau'_{in}}\left(u\left(v_i\left(T\right), b_i\left(T,s\right), b_{-i}\left(T,s\right)\right) | h\left(\tau_{in}\right), h_i\left(\tau_{in}\right)\right) \tag{9}$$

where beliefs are derived from information and other player's strategies. An equilibrium is a vector of strategies s such that for all players and at all information sets strategy s_i maximizes (9) given s_{-i}.

We now show that the equilibrium derived previously remains an equilibrium in this game by showing that if all bidders use strategies that are functions of their private history only, then it is a best response to do so.

Proposition 8. *Let $B_i(v_n, \tau_n)$ denote the equilibrium bidding functions derived in section 5. Then the strategies $s_i(v_n, \tau_n; h_{in}, h_n) = \max_{m \leq n} B_i(v_m, \tau_m)$ are an equilibrium of the game.*

Proof. Suppose for all remaining players s_{-i} is defined as above. We now show the following strategy is a best response for player i. It is sufficient to check that the bidder will not deviate from the bidding strategy at any arbitrary information set, $(v_n, \tau_n; h)$. Let player i follow this prescribed strategy in all future information sets. Without loss of generality assume the optimal bid chosen at the information set considered is greater than the outstanding bid of the player. Let $F_{-i}(u|h)$ denote the distribution of highest bid of b_{-i} (h_{-i} is enough) . The optimal bid at (v_n, τ_n) solves:

$$V(v_n, \tau_n, h) = \max_b \int_{H(b)} \int^b (v(T) - u)\, dF_{-i}(u|h)\, dP(\omega) + \int_{H(b)^c} \int^{b(\omega,T)} (v(T) - u)\, dF_{-i}(u|h)\, dP(\omega)$$

where as before, $H(b)$ denotes the sets of histories where given future bidding behavior, bid b is he final bid. So the public history h is relevant only when considering its impact on the expected distribution for the highest bid of the remaining players in that set. Given the optimality of future bidding and the definition of $H(b)$, the effect on the integration set can be ignored, so it immediately follows that the optimal bid is independent of the distribution F_{-i} and consequently independent of h. □

8.1 Censoring and Observed Behavior

Let $B_i(v, \tau)$ denote the maximum bid that bidder i is willing to place at time t. When taking the model to the data one must take into account that when a bidder observes an outstanding bid that is greater than $B(v, \tau)$, he will choose not to bid. Hence, the process for observed bids is censored and this censoring depends on the available information. In eBay auctions, bidders can see the outstanding second highest value thus defining the threshold below which bids are censored. The model has implications for the observed timing of bids. If at time t a bidder has a high probability of returning to the auction right before its end, its desired bid $B(v, \tau)$ will be very low and thus it is likely to be censored by an existing higher second bid. As a consequence, it is unlikely to see any bids from this bidder until the end of the auction, consistent with the observed sniping behavior.

There is likely to be asymmetry in the frequency of bidding times for different bidders. The observation of sniping–defined as bids that are overtaken in the last few minutes of the auction–is usually interpreted as indication that many bidders follow this kind of strategy. As the following example shows, this might not be true. Suppose there are n bidders. One of these bidders can bid with probability one at the end of the auction, while the other $n-1$ bidders only bid at the beginning of the auction. To make things extreme, suppose the final value is uniform between $[0, 1]$ and all the $n-1$ initial bidders have no information and thus approximately the same expectation equal to $1/2$ and so the initial winning bid is also around $1/2$. It follows that the probability that the remaining "sniping" bidder wins the auction is $1/2$. So in this auction, $1/2$ of the times the auction will be *sniped* while the share of snipers is only $1/n$. Note also that given the information structure sniping is still efficient as in absence of correlated information the expected value of the sniped bidder is $1/2$ and thus lower than the value for the sniper. This is in line with the distribution of observed bids versus the distribution of winning bids. Winning bids tend to be more skewed toward the end when compared to all bids placed.

9 Estimation

In this section, we provide estimates and some counterfactual exercises for the Brownian motion setting described in section 7. In our basic setting, bidders are symmetric and have independent draws for their arrival times, values and bidding times. Arrival times are given by a Poisson process with parameter η. This implies that both the total number of bidders and their arrival times are random. We normalize the total time of the auction $T = 1$, so that both, the expected number of bidders participating in an auction and the variance are equal to η. It is interesting to note that conditional on the number of bidders participating in an auction, arrival times are uniform. Bidding times are given by a Poisson process with arrival rate ρ.

The process for values is modeled in the following way. For an agent arriving at time $t\epsilon [0,1]$, the initial value is drawn from a Normal distribution with mean μ and standard deviation $\sigma_0+\sigma t$. After this arrival it follows a Brownian motion with zero drift and volatility σ. This process implies that the unconditional distribution of values for a bidder bidding at time τ is also Normal with mean zero and standard deviation $\sigma_0 + \sigma\tau$ *independently of the time of arrival to the auction.* Our assumption captures the idea that the amount of information an agent has in an auction only depends on calendar time and not on the time of entry. In particular this would be true if all relevant information to determine valuation is publicly available to all bidders.

We consider three sub categories of eBay marketplace: tablets and iPads, tickets, arts. We chose these categories because they are more consistent with private value assumption in our setup. In this simple estimation method we identify four parameters. The parameters to estimate for each case are $(\eta, \mu, \sigma_0, \sigma, \rho)$. To identify these parameters we use the simulation method of moments and a dynamic grid search on the parameter space. We choose moments that we believe are informative of the above parameters.

We consider the following four moments for the estimation:

- Average number of bidders-defined as those that placed either a highest or second highest bid at some point during the auction. We refer to these as *recorded* bids.
- Average number of bids per bidder.
- Variance of bids recorded within the first 50% of duration of auction normalized by final price.
- For bidders who bid both in the first 50% of the time and also in the last 10% of the time the increment between the two bids normalized by the final price.

It is important to note that while we have a total of five parameters, $(\eta, \mu, \sigma_0, \sigma, \rho)$, three of them (μ, σ_0, σ) can be identified, with our moments, only up to a scalar. (There is no problem identifying the other two.) This is a consequence of Lemma 2 and the normalization

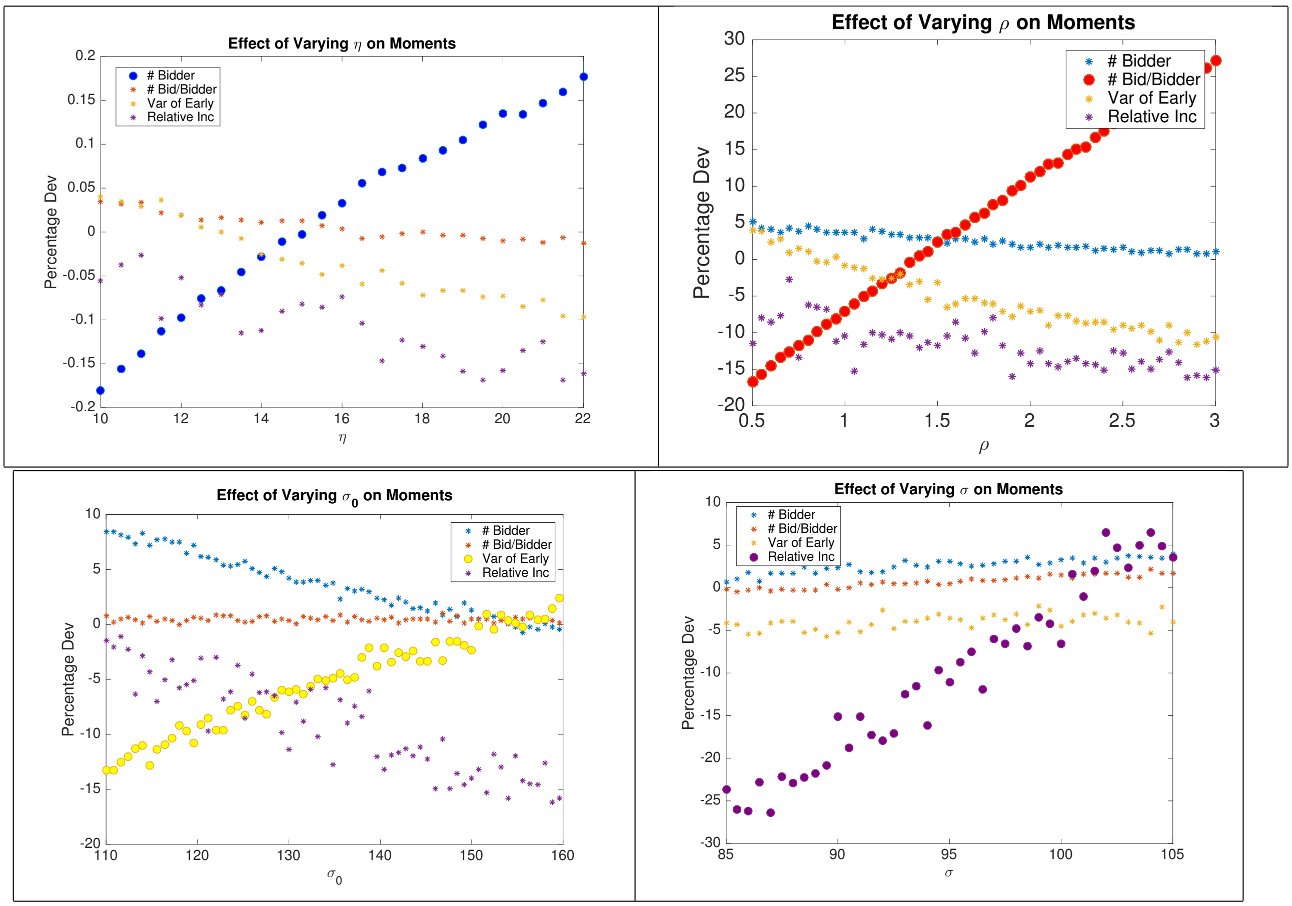

Figure 7: Identification, Effect of Varying Parameters on Moments

used to calculate our two last two moments. Intuitively, scaling μ, σ_0 and σ by a factor λ is equivalent to an equal change in scale of bidders' values. This in turn implies that for any realization of values and bidding times in an auction, all bids will be scaled by λ and in particular so will the final price. Given that in computing the last two moments bids are normalized by this final price, the corresponding ratios are unaffected by the scale change. Hence our moments allow us to identify only the ratios σ_0/μ and σ/μ. This is not a problem as these ratios are also sufficient for our counterfactual exercises.

Figure 7 illustrate the identification strategy. In each of these sub-figures we vary one of the parameters while keeping the rest constant at their estimated values. Then we calculate the value for the four targeted moments at each point and report the percentage deviation from the corresponding target values. Most of the parameters affect all the moments to some extent, but one of them the largest Increasing η increases the number of bidders; increasing ρ increases the number of bids per bidder; and increasing σ increases the increment between two bids submitted by a bidder. The only exception is σ_0; however, as this is the only moment that affects the variance of early bids, we can use this moment to identify σ_0.

The estimation is based on simulated method of moments. We start with a coarse grid for parameters and make the grid finer in each round. The steps within each round is as follows:

1. Solve the PDE for the shading function as described in section 7 for a fixed set of parameters (σ, ρ) .

2. Simulate all bids taking into account censoring, i.e., only consider a bid if it was either first or second highest bid at the time it was placed.

3. Construct moments from $n = 500$ simulations.

4. Define a loss function as squared of percentage deviations from above moments.

The estimated parameters are the the ones that minimize the above loss function. Table 1 gives the moments for three different categories of items on eBay: Tablets, Tickets, and Arts. These products have arguably more of a private value rather than common value. Given that gradual bidding cannot be generated by our model, we remove all such bids by only

Table 1: Moments

	#of bidders	# bids per bidder	Var of early bids	Average Increment
Tablets	6.28	1.38	0.077	36%
Tickets	2.97	1.53	0.088	27%
Arts	2.03	1.27	0.114	43%

Table 2: Estimated Parameters

	η	$\frac{\sigma_0}{\mu}$	ρ	$\frac{\sigma_0}{\mu}$
Tablets	16	0.46	1.4	0.31
Tickets	3.9	0.61	1.6	0.29
Arts	2.6	0.5	0.7	0.57

considering the last bid a bidder placed within two minutes of one another. This reduces the average number of bids per bidder.

Table 2 shows the estimated parameters, $\left(\eta, \frac{\sigma_0}{\mu}, \rho, \frac{\sigma}{\mu}\right)$.

10 Counterfactuals

In this section, we first consider bid retraction as it may apply to all or a subset of bidders. This exercise has a series of motivations. First, in the absence of restrictions on bid retraction, bidders have no incentives to shade bids and this allows us to quantify the effect of shading. Second, the question might be of interest as a matter of auction design. Third, in practice some bidders are able to work around the restrictions on bid retraction by the use of a sniping program that submits bids automatically in the last few seconds of the auction. Bidders using these programs are able to change their bids at any point in time and thus can retract their intended bids if desired so. We compare the results with both our baseline and an alternative benchmark given by an ex-post efficient allocation, as would be given by a second price auction that takes place at the end of time

10.1 Bid Retraction

For the bid retraction case, we assume that bidders can choose to retract/lower bids during any one of their bidding times but not outside these bidding times. This is a natural restriction, since when bidders are allowed to retract at the end, bidding would become meaningless. The possibility of retraction eliminates the source of adverse selection mentioned earlier and thus the incentives for bid shading. Given that values follow a Martingale, at every bidding time τ a bidder will choose a bid equal to the current value v_τ.

The impact of bid retraction on the distribution of final bids for a given bidder can be further characterized.

Proposition 9. *The final bids under bid retraction are a mean preserving spread of those without retraction.*

Proof. Consider an information set (v_n, τ_n) where bid b_n is placed. From Proposition 4 it follows that $Eb(T|v_n, \tau_n) = v_n$ and the same holds true under bid retraction. Let $H_0(b)$ denote the set of histories starting at this node where bid b is the final bid. On $H_0(b)^c$ the expected final bid must be the same with and without retraction. This is because the first time there is rebidding in that set, final bids are unbiased and equal to the corresponding value, for both cases. It must then follow that the expected final bid in $H_0(b)$ under bid retraction is the same that under no retraction, which by definition is b, implying that bid retraction gives a mean preserving spread of final bids in this set. Applying this argument forward on $H_0(b)^c$ proves that final bids under bid retraction are a mean preserving spread of those without retraction. □

The mean preserving spread under bid retraction has two implications. First, it represents the fact that bids are more closely connected to values. Secondly, it is likely to lead to an increase the second order statistic and therefore price, particularly in the presence of many bidders. The first effect can be seen in Figure 8 that compares the correlation between values and bids for the baseline and no-retraction case. In the left hand we consider simulated final values and final bids in both scenarios. There is substantial dispersion in both cases, as bids are placed some random time before the end of the auction giving rise to some residual uncertainty as of final value. The correlation between bids and value is 0.85 for the baseline case and 0.88 for the no-retraction case. The right hand side figure plots the correlation between bids and the expected value at the time of the last bid/retraction opportunity. Here by definition bids equal to the value with no-retraction, so the figure illustrates more

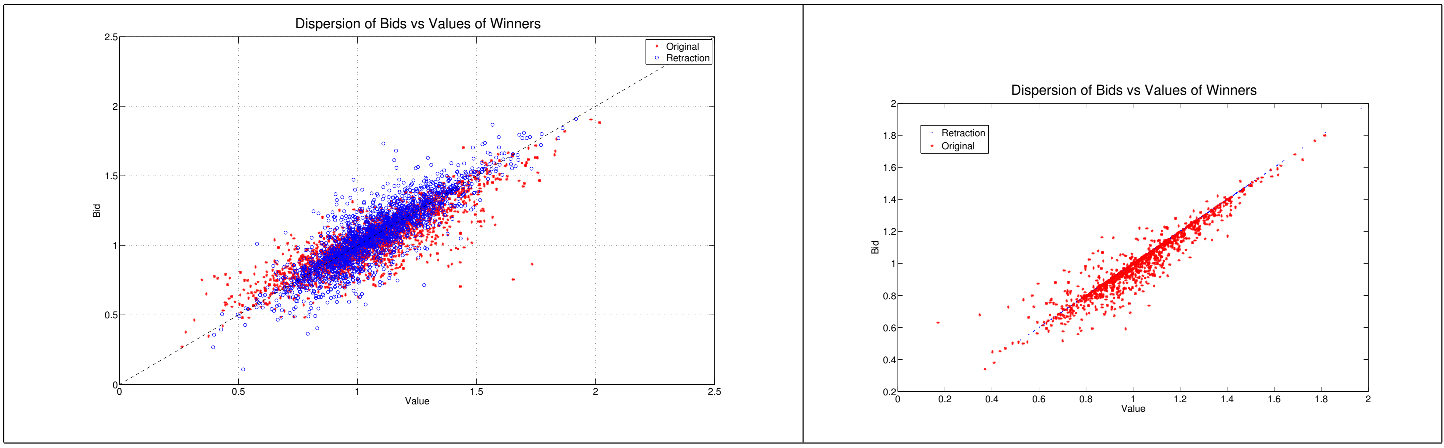

Figure 8: Correlation between bids and values

Table 3: Value and Revenues Compared To Baseline, Tablets

	Bid-retraction	Efficient
Expected value	+1.3%	+3.2%
Expected price	+10.2%	+12.2%
Expected payoff	-33.0%	-31.4%

clearly the added dispersion in the baseline scenario, which is consistent with the lower correlation indicated above.

Table 3 gives the average performance of these bidding scenarios compared to the baseline case. The expected value in the efficient final allocation is 3.2% higher than that in the baseline. Bid retraction reduces this gap by almost half, which is consistent with the higher correlation between bids and values. In turn, average prices are 10.2% higher under bid retraction, a consequence of the mean preserving increase in spread discussed above. The increase in prices exceeds the increase in value, so bid-retraction decreases the expected payoffs for bidders by approximately 33% while increasing the sellers' revenue by 10.2%.

10.2 Incentives For Bid Retraction

Our results above show that bidders collectively lose if all have the ability of retracting bids. The question we ask here is if a single bidder might gain when given the possibility of retracting bids as could be done with a sniping program. To be consistent, in our analysis we maintain the assumption of inattention: that bidders can make changes to the desired

Table 4: One Sniper, Percentage Deviations from Baseline Case

	ΔExpected Payoff for Snipers	ΔExpected Payoff for non-Snipers	ΔProb of Winning for Snipers
Tablets	12.4%	-2.4%	55.0%
Tickets	4.8%	-2.7%	24.6%
Arts	32.6%	-2.4%	62.8%

Table 5: Two Types Bidders, Percentage Deviations from Baseline Case

	Δ Expected Payoff For Snipers	ΔExpected Payoff For non-snipers	ΔProb of Winning for Snipers
Tablets	-13.2%	-29.2%	22.8%
Tickets	-6.8%	-13.3%	11.0%
Arts	1.1%	-30.8%	55.5%

bids only at their randomly given bidding times.

Table 4 shows that if a bidder can singly retract his bid, e.g., by using a sniping program, his expected payoff increases substantially. This increase comes from an increase in the probability of winning compared to the baseline case. However, the increase in expected payoff is less than the increase in this probability because bidders payoff conditional on winning goes down as they win more marginal cases. The remaining bidders are worse off as their probability of winning decreases and their payoffs conditional on winning are also slightly lower.

10.3 Two Bidder Types

Previous sections shows an incentives for bidders to use sniping programs. In this section, we identify the model by assuming a fraction of bidders use these programs while others follow the strategy in the baseline case. To identify the new parameter, share of bidders using sniping program, we add a new moment: share of winning bidders placed in the last 10% of auction. The identification strategy follows the same steps as the baseline model; only by adding the two different type of bidders. We assume that both groups are exactly the same and the only difference for them is that one group uses the sniping programs and therefore does not shade their bid and only their last bid is recorded.

As shown in Table 5, when a bigger fraction of bidders use sniping programs their probability of winning compared to the baseline does not increase as much as when only one bidder uses sniping. The final result is a lower payoff for snipers compared to the baseline (except for Arts), as bidders' payoff conditional on winning is lower when they do not shade their bid. This change in payoffs resembles a prisoners' dilemma game. If no one is using a sniping program, each bidder has incentives to unilaterally deviate and use a sniping program, thus increasing the winning probability substantially. However, as more bidders use such programs, then their probability of winning increases by less. At the same time, according to Table 5, snipers' payoff is still higher than that of non-snipers. If the remaining bidders were to use also a sniping program, the result is the one obtained in our first counterfactual where all bidders are able to retract bids, and where bidders are worse off compared to the baseline. In this extreme case, sniping does not increase the probability of winning compared to the baseline but makes, however prices in equilibrium are higher.

11 Extensions

This section considers two extensions to our model. The first one is endogenous rebidding, where there is a cost to rebid and bidders exercise this option optimally. The environment is considerably more difficult as the value of rebidding is highly dependent of the distribution of other players' bids, in contrast to the simple bidding derived above. The second extension is to allow for correlation between the signals received by bidders. While we still assume that conditional on these signals values are independent across bidders, the adverse selection problem becomes more severe: winning the auction when placing an early bid is correlated with negative signals observed by later bidders. As we show below, this results in a larger incentive to shade bids.

11.1 Endogenous Rebidding

We consider here a very simple model. There are $N = 2$ bidders and 2 periods. Let v_i denote the value for bidder i in the first period drawn from distribution $G(v)$ and v_i' the value in the second period, drawn from conditional distribution $F(v'|v)$. Both bidders can bid freely the first period but must pay a cost $c > 0$ to rebid in the second period. We assume bids are sealed.

Strategies for the bidders can be defined as follows: a bidding function $B_i(v)$ for the first period and a rebidding set $R_i(v)$ for the second period. Let $N_i(v)$denote the complement of $R_i(v)$. Given the strategy for the other player, player $i's$ expected utility is given by:

$$\begin{aligned} U(v) &= \int_{N_i(v)} Q_i(B_1(v))(v' - P_i(B_i(v)))\, dF(v'|v) \\ &+ \int_{R_i(v)} -c + Q_i(v')(v' - P_i(v'))\, dF(v'|v) \end{aligned}$$

where Q_i is the probability of winning function and P_i the expected payment conditional on winning. A symmetric equilibrium $(Q, R)$is a Nash equilibrium in these strategies.

Example 1. Both bidders draw independently their initial value $v\epsilon[0, 1]$ from distribution G and with probability $1 - \rho$ get zero value next period and with probability ρ the value remains equal to v. Conjecture a threshold v^* so that:

$B_1(v) = \rho v$ if $v < v^*$and zero otherwise

$R(v) = \{\}$ for $v < v^*$and $R(v) = \{v\}$ otherwise

Consider the player with $v = v^*$. Bidding first or second period doesn't change his probability of winning since for $\rho v^* \leq b \leq v^*$, $Q(b) = G(v^*) + (1 - \rho)(1 - G(v^*))$. The expected payment is also the same in both cases equal to:

$$\frac{\rho \int_0^{v^*} v dG(v)}{G(v^*) + (1 - \rho)(1 - G(v^*))}.$$

The difference is that if he chooses not to rebid, he pays this expected value for sure while if he chooses to rebid he pays it with probability ρ. The difference in expected payment is then:

$$(1 - \rho) \frac{\rho \int_0^{v^*} v dG(v)}{G(v^*) + (1 - \rho)(1 - G(v^*))}$$

to which we need to add that he pays an expected cost ρc. So v^* must be such that:

$$c = \frac{(1 - \rho) \int_0^{v^*} v dG(v)}{\rho G(v^*) + (1 - \rho)}. \tag{10}$$

The derivative of the right hand side with respect to v^* equals in sign to

$$v^* \left[\rho G\left(v^*\right)+(1-\rho)\right]-\rho \int_0^{v^*} v dG\left(v\right)>0.$$

The last step is to show that for all $v>v^*$it is optimal to rebid and conversely for those not in this set. The difference between rebidding and not is

$$\begin{aligned} & \rho \textstyle\int_0^v (v-x)\, dF_{-i}(x)-\int_0^{\rho v}(\rho v-x)\, dF_{-i}(x) \\ = \quad & \textstyle\int_0^v (\rho v-\rho x)\, dF_{-i}(x)-\int_0^{\rho v}(\rho v-x)\, dF_{-i} \end{aligned}$$

Using the envelope condition and taking derivatives with respect to v (keeping bids fixed)

$$\rho F_{-i}(v)-\rho F_{-i}(\rho v)>0.$$

This establishes that the gains from rebidding are increasing in v so the thresholdv^*defined by (10) is an equilibrium and it is unique.

One might expect that similar results will hold with more bidders. Moreover, it is natural to conjecture that the threshold for rebidding increases with the number of bidders, as expected payoffs decrease. Following this conjecture, with endogenous bidding one might expect that the number of bids per player decreases too.

11.2 Correlated Information

In many cases, it is likely that information or signals observed are correlated across bidders. For example, the arrival of a competing auction with a similar product is an event that creates an opportunity cost and is likely to affect in a correlated way the value of all bidders that keep track of that information. Our results extend easily to the case of pure common values, where all agents have the same values but different bidding windows. For more general cases we do not have a general result so we restrict to a simplified scenario.

11.2.1 Pure common values

The setting is identical to that described in our general model, but where all bidders values v_i are identical, bidders observe the same signals but have independent bidding windows. As before, it follows that the optimal bid at state (v, τ) satisfies $b_i = E(v(T)|v, \tau, b(T) = b_i)$ where now $b(T) = \max_{ij} b_i(v_j, t_j)$ is the maximum over all bids, including those of other bidders. The recursive representation given before still holds,

$$W(b, v, \tau) = \int_{\tau}^{T} \min(W(b, v', \tau'), 0)\, dP(v', \tau'|v, \tau) + P(\tau' > T|v, \tau)(E[v_T|v, \tau, \tau' > T] - b)$$

where the interpretation of arrivals now is for the arrival to any bidder. As an example, if opportunities for bidding are Poisson with arrival ρ then the total arrival rate that would be used in this dynamic programming equation is $N\rho$, where N is the number of bidder. The effect on bidding behavior is similar to an increase in ρ that intuitively should result in greater shading of bids. This follows naturally the interpretation of increased adverse selection as now the current bid will win not only when that bidder decides not to bid any higher but when any other bidder chooses not to do so.

11.2.2 Imperfectly correlated information

Consider the following simplified scenario. Suppose time is discrete and there are two periods $\{1, 2\}$ where 2 corresponds to the end of the auction. Each bidder receives an independent initial value v_i, drawn from distribution $F(v)$, that we interpret as the unconditional expected final value in absence of further information. They simultaneously bid in the first period. With probability ρ, they have a chance to bid in the second period. Second period valuations are drawn from distribution $G(v'|v, \theta)$ where θ is a common observable shock that is independent of the initial values. Here v (resp. v') denotes the initial (resp. final) value for bidder one and v_2 (v_2') the corresponding values for bidder two.

An agent that bids in the second period will choose $b_2(v') = v'$. Let $B_1(v)$ denote the bid in the first period. This bid should equal the expected value conditional on the union of the following events: (1) none of the two agents get to rebid and $B_1(v) > B_1(v_2)$; (2) agent one gets to rebid but $v' \leq B_1(v)$and $v_2 \leq v$; (3) other agent rebids but $v_2' < B_1(v)$

and $B_1(v_2) < B_1(v)$, and (4) both get to rebid but $v', v_2' < B_1(v)$ and $B_1(v_2) < B_1(v)$. In a symmetric equilibrium with monotone bidding functions, the condition $B_1(v) \geq B_1(v_2)$ can be substituted by $v \geq v_2$.

Consider a symmetric equilibrium Given a bid b for the first player, we claim the following:

Proposition 10. *$B_1(v)$ is an equilibrium if and only if $E(v'|v, B_1(v)) = B_1(v)$, where*

$$E(v'|v,b) = \frac{F(v)(1-\rho)^2 v_1(b,v) + F(v)\rho(1-\rho)v_2(b,v) + \rho(1-\rho)v_3(b,v) + F(v)\rho^2 v_4(b,v)}{\pi(v,b)} \tag{11}$$

where:

$$\begin{aligned}\pi(v,b) &= F(v)(1-\rho)^2 + F(v)\rho(1-\rho)G(b|v) + \rho(1-\rho)P(v_2 \leq v, v_2' \leq b)\\ &\quad + F(v)\rho^2 \int \chi_{v' \leq b, v_2' \leq b} dP(v', v_2'|v, v_2 \leq v)\end{aligned}$$

$$v_1(b,v) = E[v'|v]$$

$$v_2(b,v) = E[v'|v' \leq b, v]$$

$$v_3(b,v) = \int v' dG(v'|v,\theta)\, dP(\theta|v_2' \leq b \& B_1(v_2) < b)$$

$$v_4(b,v) = \int \chi_{v' \leq b, v_2' \leq b} v' dP$$

and $H(v_2', v_2)$ is the joint distribution that is assumed to be independent of v.

Proof. Note that the denominator is also $\pi(v,b)$ the probability of winning with a bid equal to b. Suppose the bidding function satisfies this condition for all v. By symmetry and monotonicity of the bidding function, $B_1(v_2) \leq B_1(v)$ if and only if $v_2 \leq v$. Hence conditioning on $v_2 \leq v$ is equivalent to conditioning on $B_1(v_2) \leq b$ for $b = B_1(v)$. It is now easy to verify the expression given in equation (11) for $b = B_1(v)$ is precisely the conditional expectation described above. □

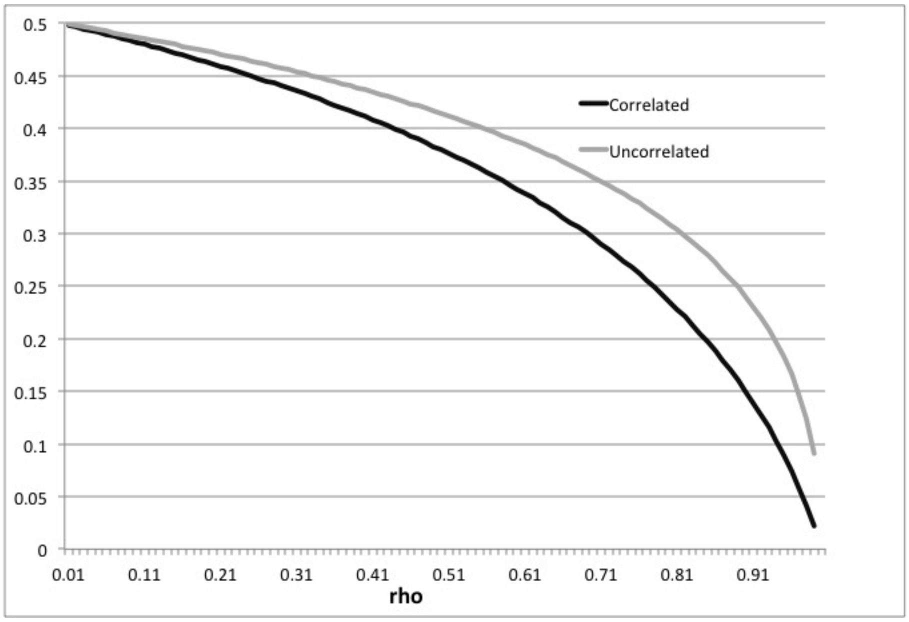

Figure 9: Bidding functions: effect of correlation

Suppose $v_i' = \theta v_i$ where θ is the common component. Rewriting equation (11),

$$E\left(v'|v,b\right) = \frac{F\left(v\right)\left(1-\rho\right)^2 v \int \theta dG\left(\theta\right) + F\left(v\right)\rho\left(1-\rho\right) v \int^{b/v} \theta dG\left(\theta\right) + \rho\left(1-\rho\right) v \int \chi_{\{v_2 \leq v, \theta \leq b/v_2\}} \theta dG\left(\theta\right) dF\left(v_2\right)}{F\left(v\right)\left(1-\rho\right)^2 + F\left(v\right)\rho\left(1-\rho\right) G\left(b/v\right) + \rho\left(1-\rho\right) \int \chi_{\{v_2 \leq v, \theta \leq b/v_2\}} dG\left(\theta\right) dF\left(v_2\right)} \tag{12}$$

Suppose in addition that v and θ are both uniform $[0,1]$. Closed form expressions can be found when θ is uniform $[0,1]$ and v_2 is also uniform.

$$E\left(v'|v,b\right) = \frac{v^2\left(1-\rho\right)^2/2 + \rho\left(1-\rho\right)b^2/2 + \rho\left(1-\rho\right) v \int^v \min\left(\frac{1}{2}, \frac{b^2}{2v_2^2}\right) dv_2 + \frac{\rho^2 v b^2}{2}}{v\left(1-\rho\right)^2 + \rho\left(1-\rho\right) b + \rho\left(1-\rho\right) \int^v \min\left(1, b/v_2\right) dv_2 + \rho^2 b} \tag{13}$$

$$= \frac{v^2\left(1-\rho\right)^2/2 + \rho\left(1-\rho\right)b^2/2 + \rho\left(1-\rho\right) v \left(b - \frac{b^2}{2v}\right) + \frac{\rho^2 v b^2}{2}}{v\left(1-\rho\right)^2 + \rho\left(1-\rho\right) b + \rho\left(1-\rho\right)\left(b + b\ln\left(v/b\right)\right) + \rho^2 b} \tag{14}$$

$$= \frac{v^2\left(1-\rho\right)/2 + \rho v b + \frac{\rho^2 v b^2}{2(1-\rho)}}{v\left(1-\rho\right) + \rho\left(2b + b\ln\left(v/b\right)\right) + \frac{\rho^2 b}{(1-\rho)}} \tag{15}$$

We solve for the fixed point above numerically and compare the bidding function with the one obtained for the uncorrelated case, where θ is independently drawn for the two players from a uniform distribution. Figure 9 plots bidding functions in both scenarios for an initial value $v = 1$. The x axis shows different probabilities of rebidding $\rho$and the y axis the corresponding bids. Consistent with our findings, as $\rho \to 1$ bids go to zero in both cases and at the other extreme, when $\rho = 0$ bids equal the unconditional mean of $\theta = 1/2$. More importantly, when information is correlated bidders shade their bids even more.

References

Attila Ambrus, Yuhta Ishii, and James Burns. Gradual bidding in ebay-like auctions. *Economic Research Initiatives at Duke (ERID) Working Paper*, (129), 2013. 1

Matt Backus, Thomas Blake, Dimitriy V Masterov, and Steven Tadelis. Is sniping a problem for online auction markets? In *Proceedings of the 24th International Conference on World Wide Web*, pages 88–96. International World Wide Web Conferences Steering Committee, 2015. 1, 2

Matthew Backus and Gregory Lewis. A demand system for a dynamic auction market with directed search. *Harvard University, October*, 2012. 1

P. Bajari and A. Hortacsu. The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions. *RAND Journal of Economics*, pages 329–355, 2003. 1

Dominic Coey, Bradley Larsen, and Brennan C Platt. A theory of bidding dynamics and deadlines in online retail. 2015. 1

Jeffrey C Ely and Tanjim Hossain. Sniping and squatting in auction markets. *American Economic Journal: Microeconomics*, pages 68–94, 2009. 1

Raul Gonzalez, Kevin Hasker, and Robin C Sickles. An analysis of strategic behavior in ebay auctions. *The Singapore Economic Review*, 54(03):441–472, 2009. 2

Sean Gray and David Reiley. Measuring the benefits to sniping on ebay: Evidence from a field experiment. *Preliminary and Incomplete Draft*, pages 1–18, 2004. 1

Joachim R Groeger and Robert A Miller. Bidding frictions in ascending auctions. 2015. 1

Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *The Review of Economic Studies*, 49(3):315–333, 1982. 4

Stephen C Hayne, CAP Smith, and Leo R Vijayasarathy. Who wins on ebay: An analysis of bidders and their bid behaviours. *Electronic Markets*, 13(4):282–293, 2003. 2

Kenneth Hendricks and Alan Sorensen. The role of intermediaries in dynamic auction markets. Technical report, Working Paper, 2014. 1

Axel Ockenfels and Alvin E Roth. 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), 2002. 1

Axel Ockenfels and Alvin E Roth. Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic behavior*, 55(2):297–320, 2006. 1

Maher Said. Sequential auctions with randomly arriving buyers. *Games and Economic Behavior*, 73(1):236–243, 2011. 1

Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476, 2006. 1

12 Proofs

Proof of Proposition 1

Consider the optimal policy starting at (v, τ). Let $V(v, \tau, b)$ denote the expected utility bidding b at this state and following the optimal bidding policy for the future. A convenient expression can be obtained as follows:

$$\begin{aligned} V(v,\tau,b) &= \int_{H(b)} \int^{b} (v(T) - x)\, dG_{-i}(x)\, dP(\omega) \\ &+ \int_{H(b)^c} \int^{b(\omega,T)} (v(T) - x)\, dG_{-i}(x)\, dP(\omega) \end{aligned}$$

where $b(\omega, T)$ is the final bid for path ω when following the optimal bidding function and $H(b)$ denotes the set of paths starting from this node such that b is the final bid, i.e. $H(b) = \{\omega \epsilon \Omega(v,\tau) : b(\omega,T) = b(\omega,\tau) = b\}$. Abusing notation we will let $H(b)^c = \{\omega \epsilon \Omega(v,\tau) : b(\omega,T) > b(\omega,\tau) = b\}$ and in the following we $b = B(v,\tau)$ is the optimal bid at this state. We want to show that this optimal bid

$$b = E_{H(b)} \int^{b} v(\omega, T)\, dP(\omega).$$

To prove this, we establish as usual the two corresponding inequalities. First we show that $b \geq E_{H(b)} \int^{b} v(\omega,T)\, dP(\omega)$. Take $V(b+\varepsilon, v, \tau)$, the optimal value if the current bid is $b+\varepsilon$ followed by future optimal bids. A particular feasible policy after bidding currently $b+\varepsilon$ is to keep this bid in the original set $H(b)$, and set $b' = \max(B(v', \tau'), b+\varepsilon)$ in the complement. Letting $b(\omega,T)$ denote the final bid after path ω induced by the optimal bidding policy (starting at b), this alternative policy from $b+\varepsilon$ will give final bid equal to $\max(b(\omega,T), b+\varepsilon)$ and an expected utility as given from the right hand side of the equation below. Since this policy is feasible but not necessarily optimal starting from $b+\varepsilon$, the following inequality holds:

$$
\begin{aligned}
V(b+\varepsilon, v, \tau) &\geq \int_{H(b)} \int^{b+\varepsilon} (v(T)-x)\, dG_{-i}(x)\, dP(\omega) \\
&+ \int_{b(\omega,T)>b+\varepsilon} \int^{b(\omega,T)} (v(T)-x)\, dG_{-i}(x)\, dP(\omega) \\
&+ \int_{b<b(\omega,T)<b+\varepsilon} \int^{b+\varepsilon} (v(T)-x)\, dG_{-i}(x)\, dP(\omega). \qquad (16)
\end{aligned}
$$

Lemma 4. *Let $g^{+}_{-i}(b)$ denote the right derivative of cdf G_{-i} at b, which is defined almost everywhere. Then:*

$$
\begin{aligned}
\lim_{\varepsilon\downarrow 0} \frac{V(b+\varepsilon, v, \tau) - V(b, v, \tau)}{\varepsilon} &\geq \lim_{\varepsilon\downarrow 0} \frac{\int_{H(b)} \int_b^{b+\varepsilon} (v(\omega,T)-x)\, dG_{-i}(x)\, dP(\omega)}{\varepsilon} \\
&= g^{+}_{-i}(b) \int_{H(b)} (v(\omega,T)-b)\, dP(\omega).
\end{aligned}
$$

It also follows that $b \geq E_{H(b)} v(\omega, T)$.

Proof. Using (16) it follows that

$$
\begin{aligned}
\lim_{\varepsilon\downarrow 0} \frac{V(b+\varepsilon, v, \tau) - V(b, v, \tau)}{\varepsilon} &\geq \lim_{\varepsilon\downarrow 0} \frac{\int_{H(b)} \int_b^{b+\varepsilon} (v(\omega,T)-x)\, dG_{-i}(x)\, dP(\omega)}{\varepsilon} \\
&+ \lim_{\varepsilon\downarrow 0} \frac{\int_{b<b(\omega,T)<b+\varepsilon} \int_{b(\omega,T)}^{b+\varepsilon} (v(\omega,T)-x)\, dG_{-i}(x)\, dP(\omega)}{\varepsilon}
\end{aligned}
$$

The inner integral in the second term is bounded below by

$$
\left[-(b+\varepsilon) + \bar{\min}(v(\omega,T), 0)\right] (G_{-i}(b+\varepsilon) - G_{-i}(b))
$$

and thus

$$
\begin{aligned}
&\lim_{\varepsilon\downarrow 0} \frac{\int_{b<b(\omega,T)<b+\varepsilon} \int_{b(\omega,T)}^{b+\varepsilon} (v(\omega,T)-x)\, dG_{-i}(x)\, dP(\omega)}{\varepsilon} \\
\geq\ &\lim_{\varepsilon\downarrow 0} \frac{[G(b+\varepsilon) - G(b)]}{\varepsilon} \int \chi(\omega : b < b(\omega,T) < b+\varepsilon) \left[-(b+\varepsilon) + \bar{\min}(v(\omega,T), 0)\right] dP(\omega).
\end{aligned}
$$

As $\varepsilon \downarrow 0$ the function under the integral converges pointwise (and monotonically) to zero so by the monotone convergence theorem the limit is zero, completing the first inequality in the proof. The second equality follows from the definition of $g_{-i}^{+}(b)$ as the right limit of the cdf. Finally that $b \geq E_{H(b)} v(\omega, T)$ follows from the fact that $V(b+\varepsilon, v, \tau) - V(b, v, \tau) \leq 0$ since b is the optimal choice at (v, τ) and thus $g_{-i}^{+}(b) \int_{H(b)} (v(\omega, T) - b)\, dP(\omega) \leq 0$. □

We next show that the reverse inequality $\int_{H(b)} (v(\omega, T) - b)\, dP(\omega) \geq 0$ also holds which together with the above imply that $b = E_{H(b)} v(\omega, T)$. To prove this, consider bidding $b - \varepsilon$. Starting at $b-\varepsilon$ the following is a feasible policy: 1) maintain the bid at $b-\varepsilon$ in the original set $H(b)$ and also follow the original bidding function in its complement. Since this is feasible but not necessarily optimal, it follows that:

$$\begin{aligned} V(b, v, \tau - \varepsilon) \;\; \geq & \;\; \int_{H(b)} \int^{b-\varepsilon} (v(\omega, T) - x)\, dG_{-i}(x)\, dP(\omega) + \\ & \;\; \int_{H(b)^c} \int^{B(\omega,T)} (v(\omega, T) - x)\, dG_{-i}(x)\, dP(\omega). \end{aligned}$$

But since $V(b, v, \tau - \varepsilon) \leq V(b, v, \tau)$ it follows that

$$\int_{H(b)} \int^{b-\varepsilon} (v(\omega, T) - x)\, dG_{-i}(x)\, dP(\omega) \leq \int_{H(b)} \int^{b} (v(\omega, T) - x)\, dG_{-i}(x)\, dP(\omega)$$

This in turn implies that

$$\int_{H(b)} \int_{b-\varepsilon}^{b} (v(\omega, T) - x)\, dG_{-i}(x)\, dP(\omega) \geq 0$$

or changing order of integration

$$\int_{b-\varepsilon}^{b} \int_{H(b)} (v(\omega, T) - x)\, dP(\omega)\, dG_{-i}(x) \geq 0$$

for all ε. Suppose now towards a contradiction that $\int_{H(b)} (v(\omega, T) - b)\, dP(\omega) < 0$. By continuity it follows that for some small ε the above inequality will be violated.

Proof or Proposition 2

Any bidding strategy $B(v,\tau)$ defines a joint stochastic process for final bids and values. As before, we will denote $b(T)$ and $v(T)$ the corresponding final bids and values, respectively. Let $B(v,\tau)$ be a bidding strategy with the property that

$$E\left[v(T)\,|v,\tau \text{ and } b(T)=B(v,\tau)\right]=B(v,\tau) \tag{17}$$

for all states (v,τ), i.e., the bid $B(v,\tau)$ equals the expected final value in the set where this is the final bid. As seen in Proposition 1, this is a necessary condition for the optimal bidding strategy.

We will prove that there is a unique bidding function satisfying property (17) and that it solves $W(B(v,\tau),v,\tau)=0$, where $W(b,v,\tau)$ is the unique solution to the Bellman equation:

$$\begin{aligned} W(b,v,\tau) &= \int_{\tau}^{T} \min\left(W\left(b,v',\tau'\right),0\right) dP\left(v',\tau'|v,\tau\right) \\ &+ P\left(\tau'>T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau'>T\right]-b\right) \end{aligned} \tag{18}$$

The first step is to establish *necessity*, i.e., any bidding function $B(v,\tau)$ with property (17) corresponds to a function $W(b,v,\tau)$ satisfying this functional equation and $W(B(v,\tau),v,\tau)=0$. The second step is to show that the Bellman equation is a contraction mapping. Therefore, the function $W(b,v,\tau)$ is unique and strictly decreasing in b, and consequently there is a unique bidding function $B(v,\tau)$ for which $W(B(v,\tau),v,\tau)=0$.

Step 1. Necessity

Take a candidate bidding function $B(v,\tau)$ that satisfies property (17). For any state (v,τ) and bid b consider the process for future bids obtained by setting $b(\tau)=b$ and applying the given bidding function in the future. Using the above notation, this bidding behavior defines a joint stochastic process for final bids and values that we denote by $b(T)$ and $v(T)$ respectively. We will be interested in histories, starting from τ, where $b(T)=b$, namely those where all interim bidding opportunities (v_n, τ_n) have the property that $b \geq B(v_n, \tau_n)$. Let $Q(b, v, \tau)$ denote the resulting probability that $b(T) = b$. Now define

$$W(b, v, \tau) = \{E[v(T)|v, \tau, b(T) = b] - b\} Q(b, v, \tau), \tag{19}$$

where again the expectation is taken under the bidding process described above. We now show that $W(b, v, \tau)$ is a solution to functional equation (18). Substituting in the right hand side of (18)

$$\begin{aligned} W(b, v, \tau) &= \int \min(\{E[v(T)|v', \tau', b(T) = b] - b\} Q(b, v', \tau'), 0) \, dP(v', \tau'|v, \tau) \\ &+ P(\tau' > T|v, \tau)(E[v(T)|v, \tau, \tau' > T] - b) \end{aligned}$$

As indicated bid b will only remain valid in the next arrival state v', τ' when $B(v', \tau') \leq b$, and by Lemma 5 this will occur if and only if $E[v(T)|v', \tau', b(T) = b] \leq b$. It follows that

$$Q(b, v, \tau) = \int_{\tau}^{T} \chi_{E[v(T)|v', \tau', b(T)=b] - b \leq 0} Q(b, v', \tau') \, dP(v', \tau'|v, \tau) + P(\tau' > T|v, \tau)$$

and using the law of iterated expectations it follows that:

$$W(b, v, \tau) = \{E[v(T)|v, \tau, b(T) = b] - b\} Q(b, v, \tau).$$

To complete the first step, we need to show that $W(B(v, \tau), v, \tau) = 0$; this follows immediately from the definition of W and the property of the bidding function given by equation (17).

Step 2. Sufficiency and Uniqueness We first show that there is a unique function $W(b, v, \tau)$ satisfying (18) by establishing it is a contraction mapping. Check Blackwell sufficient conditions: Monotonicity is trivially satisfied. To check discounting, consider the function $W(b,v,\tau)+a$ for $a \geq 0$ on the right hand side of Bellman equation (18).

$$\begin{aligned}
& \int_{\tau}^{T} \min\left(W\left(b,v',\tau'\right)+a,0\right) dP\left(v',\tau'|v,\tau\right) \\
+ \;& P\left(\tau'>T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau'>T\right]-b\right) \\
\leq \;& \int_{\tau}^{T} \min\left(W\left(b,v',\tau'\right),0\right) dP\left(v',\tau'|v,\tau\right)+a\left(1-P\left(\tau'>T|v,\tau\right)\right) \\
+ \;& P\left(\tau'>T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau'>T\right]-b\right) \\
= \;& W(b,v,\tau)+a\left(1-P\left(\tau'>T|v,\tau\right)\right)
\end{aligned} \tag{20}$$

By Assumption 1 $P(\tau'>T|v,\tau)>(1-\delta)$ for some $0<\delta<1$, proving the second Blackwell sufficient condition.

To complete the proof, we need to show that there is a unique bidding function $B(v,\tau)$ satisfying $W(B(v,\tau),v,\tau)=0$ for all (v,τ). We show this recursively by establishing that W is strictly decreasing in the first argument, for any (v,τ). The proof is the by induction. So assume that the W function in the right hand side of Bellman equation (18) is weakly decreasing. Letting $b'>b$,

$$\begin{aligned}
TW\left(b',v,\tau\right) &= \int_{\tau}^{T} \min\left(W\left(b',v',\tau'\right),0\right) dP\left(v',\tau'|v,\tau\right) \\
&+ P\left(\tau'>T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau'>T\right]-b'\right) \\
&\leq \int_{\tau}^{T} \min\left(W\left(b,v',\tau'\right),0\right) dP\left(v',\tau'|v,\tau\right) \\
&+ P\left(\tau'>T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau'>T\right]-b'\right)
\end{aligned}$$

which is strictly less than $TW(b,v,\tau)$ since by Assumption 1 $P(\tau'>T|v,\tau)>0$ thus proving sufficiency and uniqueness.

Supporting Lemmas

Lemma 5. *Consider a bidding function $B(v,\tau)$ such that for all $\tau>\tau_0$ $E[v(T)|v,\tau$ and $b(T)=B(v,\tau)]=$ $B(v,\tau)$. Starting with an arbitrary bid b at state (τ_0,v_0) define $b_{\tau'}=\max(b,B(\tau',v(\tau')))$ for all $\tau'>\tau_0$. Then $E[v(T)|\tau_0,v_0$ and $b(T)=b]-b$ is decreasing in b.*

Proof. Take an initial bid b at some point (τ_0,v_0) and $\varepsilon>0$. For any path ω following (τ_0,v_0), i.e. $\omega\epsilon\Omega(\tau_0,v_0)$ and $\tau>\tau_0$ define $B_0(b,\omega,\tau)$ to the bid obtained when starting with bid b at (τ_0, v_0) and following the above bidding rule. Let $H(b)$ denote the set of paths where $B_0(b,\omega,T) = b$ and $H(b+\varepsilon)$ the set where $B_0(b+\varepsilon,\omega,T) = b+\varepsilon$ when applying same bidding strategy with initial bid $b+\varepsilon$. It follows immediately that $B_0(b+\varepsilon,\omega,T) = \max\{B_0(b,\omega,T), b+\varepsilon\}$. Then

$$\begin{aligned}
\int_{H(b+\varepsilon)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega) &= \int_{H(b)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega) \\
&+ \int_{\{b<B_0(b,\omega,T)\le b+\varepsilon\}} (v(\omega,T) - (b+\varepsilon))\, dP(\omega) \\
&\le \int_{H(b)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega) \\
&+ \int_{b<B_0(b,\omega,T)\le b+\varepsilon} (v(\omega,T) - (B_0(b,\omega,T)))\, dP(\omega) \\
&= \int_{H(b)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega)
\end{aligned}$$

where the last equality follows from Lemma 6. It follows that

$$\begin{aligned}
&\textstyle\int_{H(b+\varepsilon)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega) - \int_{H(b)} (v(\omega,T) - b)\, dP(\omega) \\
\le\; &\textstyle\int_{H(b)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega) - \int_{H(b)} (v(\omega,T) - b)\, dP(\omega) < 0 \quad .
\end{aligned}$$

Now

$$\begin{aligned}
E_{H(b+\varepsilon)} v(\omega,T) - (b+\varepsilon) &= \frac{\int_{H(b+\varepsilon)} (v(\omega,T) - (b+\varepsilon))\, dP(\omega)}{P(H(b+\varepsilon))} \\
&\le \frac{\int_{H(b)} (v(\omega,T) - (b))\, dP(\omega)}{P(H(b+\varepsilon))} \\
&\le \frac{\int_{H(b)} (v(\omega,T) - (b))\, dP(\omega)}{P(H(b))} \\
&= E_{H(b)} v(\omega,T) - b
\end{aligned}$$

□

Lemma 6. $\int_{B(\omega,T)\epsilon B} (v(\omega,T) - (B(\omega,T)))\, dP(\omega) = 0$ *for any (Borel set)* B

The proof of the Lemma follows from the definition of conditional expectation. We use the following theorem in Ash (add reference):

Theorem. *Let Y be an extended random variable on (Ω, F, P), and $X : (\Omega, \mathcal{F}) \to (\Omega', F')$ a random object. If $E(Y)$ exists, there is a function $g : (\Omega', \mathcal{F}') \to (\bar{R}, \mathscr{B})$ such that for each $A \epsilon \mathcal{F}'$,*

$$\int_{\{X \epsilon A\}} Y dP = \int_A g(x)\, dP_x(x) \tag{21}$$

where $P_x(A) = P(\omega | X(\omega) \epsilon A)$. The function $g(x)$ is interpreted as $E(Y|X = x)$.

Proof of Lemma Define a series of random variables which are a function of ω :

$b(\omega)$: final bid for that ω

$\tau'(\omega)$: time at which that final bid was placed

$v(\omega)$: final value for that ω

Let $X(\omega) = (\tau(\omega), b(\omega))$. The random variable that we will consider is $Y(\omega) = v(\omega) - b(\omega)$. Let $A = \{\tau_0 < \tau \leq T\} \times B$, where B is a Borel set (of bids) in $\mathbb{R}$. What we want to show is that:

$$\int_{X \epsilon A} Y(\omega)\, dP(\omega) = 0.$$

We know that for all $\tau > \tau_0$ and b, $E(Y|\tau, b) = 0$ where the pair (τ, b) are interpreted as the last bid and the time it was placed. Substituting $x = (\tau, b)$ and using (21),

$$\int_{\{\tau(\omega) > \tau_0, b(\omega) \epsilon B\}} (v(\omega) - b(\omega))\, dP(\omega) = \int_{\{\tau > \tau_0, b \epsilon B\}} E(Y|\tau, b)\, dP_x(\tau, b) = 0$$

Proof or Lemma 1

First note that by the law of iterated expectation, $E\left(E\left([v_T|v',\tau'] - b|v',\tau'\right)|v,\tau\right) = E\left[v_T|v,\tau\right] - b$ and that

$$\begin{aligned}
E\left(E\left([v_T|v',\tau'] - b|v',\tau'\right)|v,\tau\right) &= \int_\tau E\left([v_T|v',\tau'] - b|v',\tau'\right) dP\left(v',\tau'|v,\tau\right) \\
&= \int_\tau^T E\left([v_T|v',\tau'] - b|v',\tau'\right) dP\left(v',\tau'|v,\tau\right) \\
&\quad + \int_T E\left([v_T|v',\tau'] - b|v',\tau'\right) dP\left(v',\tau'|v,\tau\right) \\
&= \int_\tau^T E\left([v_T|v',\tau'] - b|v',\tau'\right) dP\left(v',\tau'|v,\tau\right) \\
&\quad + P\left(\tau' > T|v,\tau\right) E\left([v_T|v',\tau'] - b|v,\tau,\tau' > T\right).
\end{aligned}$$

We now prove the Lemma by induction. So suppose that $W\left(b,v',\tau'\right) \leq E\left(v\left(T\right)|v',\tau'\right) - b$. Then

$$\begin{aligned}
W\left(b,v,\tau\right) &= \int_\tau^T \min\left(W\left(b,v',\tau'\right),0\right) dP\left(v',\tau'|v,\tau\right) + P\left(\tau' > T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau' > T\right] - b\right) \\
&\leq \int_\tau^T W\left(b,v',\tau'\right) dP\left(v',\tau'|v,\tau\right) + P\left(\tau' > T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau' > T\right] - b\right) \\
&\leq \int_\tau^T E\left([v_T|v',\tau'] - b|v',\tau'\right) dP\left(v',\tau'|v,\tau\right) + P\left(\tau' > T|v,\tau\right)\left(E\left[v_T|v,\tau,\tau' > T\right] - b\right). \\
&= E\left[v_T|v,\tau\right] - b
\end{aligned}$$

where the first inequality is strict if $P\left(\tau' \leq T|v,\tau\right) > 0$, thus completing the proof.

Proof of Proposition 5

Let $S\left(\tau,\tau'\right) = \exp\left(-\int_\tau^{\tau'} \rho\left(s\right) ds\right)$ be the probability of no arrival of bidding time in the interval $[\tau,\tau']$. For the specific process considered here, the value function (5) specializes to:

$$\begin{aligned}
W\left(b,v,\tau\right) &= \int_\tau^T \rho\left(\tau'\right) S\left(\tau,\tau'\right)\left[\int \min\left(W\left(b,v',\tau'\right),0\right) dP\left(v'|v\right)\right] d\tau' \\
&\quad + S\left(\tau,T\right)\left[E\left(v_T|v,\tau\right) - b\right]
\end{aligned}$$

We prove inductively that $W(b,v,\tau)$ is weakly increasing in τ.

$$\begin{aligned}
\partial W(b,v,\tau)/\partial\tau &= -\rho(\tau)\int \min\left(W(b,v',\tau),0\right)dP(v'|v) \quad (22)\\
&\quad +\int_{\tau}^{T}\rho(\tau')\rho(\tau)\left[\int \min\left(W(b,v',\tau'),0\right)dP(v'|v)\right]d\tau'\\
&\quad +\rho(\tau)\left[E(v_T|v,\tau)-b\right]+S(\tau,T)\frac{\partial}{\partial\tau}E(v_T|v,\tau) \quad (23)\\
&> -\rho(\tau)\int \min\left(W(b,v',\tau),0\right)dP(v'|v)\\
&\quad +\rho(\tau)\left[E(v_T|v,\tau)-b\right]+S(\tau,T)\frac{\partial}{\partial\tau}E(v_T|v,\tau) \quad (24)\\
&\geq S(\tau,T)\frac{\partial}{\partial\tau}E(v_T|v,\tau)\geq 0
\end{aligned}$$

where the last inequality follows from Lemma 1. This completes the inductive proof. The second claim of the Proposition follows from the first one, the fact that $W(B(v,\tau),v,\tau)=0$ and that W is decreasing in b.