Shills and Snipes

Subir Bose
University of Leicester
sb345@leicester.ac.uk

Arup Daripa
Birkbeck College
University of London
a.daripa@bbk.ac.uk

June 2011

Abstract

Many online auctions with a fixed end-time experience a sharp increase in bidding towards the end despite using a proxy-bidding format. We provide a novel explanation of this phenomenon under private values. We identify a simple environment in which the seller bids in her own auction (shill bidding) and show that the targeted bidders optimally bid late to snipe the shill bid. A technical contribution of our work involves modeling continuous bid times and bid arrival process. We show that in all equilibria bidding is delayed to the latest instance involving no sacrifice of probability of bid arrival.

JEL CLASSIFICATION: D44

KEYWORDS: Online auctions, private values, last-minute bidding, sniping, shill bidding, random bidder arrival, continuous bid time, continuous bid arrival process

1 Introduction

Online auctions on eBay as well as many other platforms have a pre-announced end time, and in many such auctions there is a noticeable spike in bidding activity right at the end, a phenomenon often called "last minute bidding." In an English auction in which bidding is meant to be done incrementally, such behavior clearly makes sense: by bidding just before the auction closes, a bidder might be able to foreclose further bids–a practice known as sniping–and win at a low price. To prevent such behavior, eBay allows bidders to use a proxy bidding system. A bidder submits a maximum price, and the proxy bid system then bids incrementally on behalf of the bidder up to the maximum price. The advantage of the system is that the proxy-bot cannot be sniped: so long as the highest bid of others is lower then the maximum price that a bidder has submitted to the proxy bid system, he wins.

Of course, in common value environments, e.g. coin auctions, bidders might have an incentive to delay their bids even in a proxy bidding auction format in order to optimally hide the information content of their bids from other bidders (see Bajari and Hortaçsu (2003), Ockenfels and Roth (2006)).

However, a large fraction of auctions on online platforms such as eBay fit the private values paradigm well, and experience significant last minute bidding.[1] What explains such bidder behavior in a private values setting? This is the question we address in this paper, and suggest a novel solution. In contrast with the literature, our analysis establishes the equilibrium bid time using a framework in which bid times and bid arrival times are chosen in a continuous manner.

Our analysis starts by considering another phenomenon that occurs in online auctions. Sellers often put in bids assuming different identities and/or get others to bid on their behalf. While the practice–known as "shill bidding"–is illegal, and frowned upon by the online auction community, prevention requires verification, which is obviously problem-

[1]See, for example, Roth and Ockenfels (2002) and Wintr (2008) for evidence of late bidding in eBay auctions for items such as computers, PC components, laptops, monitors etc. These items are fairly standardized products and would seem to fit the private values framework much better than a common values one. Wintr reports that on eBay, around 50% of laptop auctions and 45% of auctions for monitors receive their last bid in the last 1 minute, while around 25% of laptop auctions and 22% of monitor auctions receive their last bid in the last 10 seconds.

atic. Legal or not, shill bidding is reported to be widespread in online auctions.[2]

The principal characteristic of a shill bid–the one that presumably generates all the passion surrounding the issue–is that the seller is submitting a bid above own value in order to raise the final price. In this sense, of course, any non-trivial reserve price (i.e. any reserve price that is strictly higher than the seller's own value) in a standard auction is an openly-submitted shill bid. We know from Myerson (1981) that the optimal reserve price is positive even when the seller has no value for the object for sale. However, in a standard private-value auction with a known distribution of values, the optimal reserve price is also the optimal shill bid. In other words, there is no other higher bid that the seller can submit (openly or surreptitiously) that would improve revenue.

In our model, a seller uses an online auction site (like eBay) to try to sell an item. The auction format used is proxy bidding[3]. The important point of departure is that the seller faces some uncertainty about the distribution from which bidder valuations are drawn. In such a setup, bids convey useful information to the seller and, since it is not typically possible to adjust the reserve price mid-auction, submitting shill bids can be profitable for the seller.

We show that this practice of shill bidding by the seller is directly related to the submission of bids at the last minute by bidders in private value auctions. The bidders bid at the last minute (defined below) not because they want to snipe the bids of other buyers–indeed, as noted above, this is not possible under proxy bidding–but because they want to *snipe the shill bids*. We show that the *unique equilibrium outcome* is that the seller submits shill bids once, but only *after* certain information events occur, and the bidders targeted by the shill bids optimally bid at the last possible time such that their bids reach with

[2]See, for example, the The Sunday Times (2007) report on shill bidding on eBay.

[3]This allows actual bidding to be as in an English auction while allowing asynchronous bidding since the bidder simply submits a maximum price and does not have to monitor the progress of the auction. As noted above, proxy bidding rules out the possibility that a bidder is sniped. Note however that the proxy bidding format is neither a truly English auction (it reveals less information than a truly open ascending auction), nor is it a true Vickrey auction (which is a sealed bid auction and reveals no information to bidders before they submit their bids). However, while it is not a *standard* second price auction, proxy bidding format *is* a second price auction.

probability 1.[4,5] The bids from the bidders then arrive randomly "inside" the last minute, and therefore the relevant information event also occurs "inside" the last minute. This, in turn, implies that any response from the seller (the shill bid) arrives with probability strictly less than 1. Thus the shill bid is sniped in the sense that there is a positive probability that any shill bid does not arrive. The fact that targeted bidders delay bids in equilibrium, but not so much that they sacrifice any probability of bid arrival, together with the fact that the shill bid is optimally made at a time such that it might not arrive, have the interesting implication that the auction with shills features lower revenue but higher efficiency compared to the optimal auction (when the seller faces no uncertainty about the distribution of values). The higher efficiency is because the optimal reserve price under the usual Myerson conditions is inefficiently high and results in no trade for some (buyer) values that are higher than the seller's true value. The fact that the optimal reserve price might not arrive therefore raises efficiency, and at the same time reduces revenue.

A technical contribution of our work is that, as far as we know, we are the first to solve for a continuous choice of the time of bidding with a continuous bid arrival process, both natural features of online auctions. As discussed later, the literature simply assumes some form of discontinuity to define a last point of bidding. Therefore the literature can only address the question of incentive for delayed bidding, but cannot determine the equilibrium time of bidding satisfactorily. We believe that our technique can be used in other models of auctions, or more general games with a continuous choice of timing.

We assume that the auction takes place over a time interval stretching from $-T$ to 1, where $-T < -1$. Bidders enter randomly over $[-T, 0]$, and can submit proxy bids at any time after they enter. The arrival time of a bid made at time t is uniformly distributed on $[t, t+1]$. If $t + 1 < 1$, i.e. if $t < 0$, any bid made at t arrives eventually with certainty. The last point of time at which this property holds is time 0. This is the cusp of the "last minute." For a bid made "inside" the last minute, i.e. at a time $\hat{t} > 0$, there is a prob-

[4]This is akin to using a sniping software or online sniping service. These ensure that the bid is delayed as much as possible, but not so much that it might get lost.

[5]The matter is of course more delicate than the above description suggests. There must be some pooling in timing strategies across types drawn from different distributions. For example, it can't be an equilibrium for bidders to bid early only if they come from a distribution for which the starting reserve price is already optimal. In that case the lack of bidding over some interval of time would reveal information to the seller that may undermine the usefulness of sniping.

ability of $\hat{t}$ that it does not arrive. Bidders could bid at time 0, or they could sacrifice a small amount of arrival probability by pushing their bid times to slightly later than 0, or they could sacrifice their probability of arrival a lot and bid close to time 1. These strategies have different implications for revenue and efficiency. We are therefore interested in determining the precise time of bidding.[6] We show that the unique outcome, in *any* equilibrium, is that the targeted bidders optimally choose to bid exactly at time 0. When, and only when, at least two such bids arrive, the seller optimally submits a shill bid. Since this necessarily happens inside the last minute, there is a positive probability that the shill bid does not arrive. This is the sense in which the shill bid gets sniped.

To clarify our setting of continuous choice of bid times and arrival process, and how this rules out collusive equilibria which rely on discontinuous arrival, it helps to compare our setting to that of Ockenfels and Roth (2006) who also provide a rationale for last minute bidding under private values. They assume that there exists a "last point" in time (let us call it t_L) with the following property: a bid made at the point t_L reaches with probability $0 < p < 1$; further, no one can react to such a bid if it reaches. On the other hand, a bid made at time $t_L - \varepsilon$ for *any* $\varepsilon > 0$, reaches with probability 1 *and* the other bidder has time to react and submit a counter bid which also reaches with probability 1. Given this setup, they show that there is a "collusive" equilibrium in which the bidders bid at time t_L because by doing so each takes a chance that his own bid will go through while the other bidder's bid will not - allowing the former to win and pay a low price. If anyone deviates and bids before t_L, the other retaliates and bids before t_L also, and a standard outcome follows. So long as the collusive price is low enough deviations are not profitable. Note however, that if we drop the discontinuity in bid arrival and make the arrival probability of bids a continuous function of time (a bid made at $t < t_L$ reaches with a probability that goes to zero as $t \to t_L$), then starting from the situation where bidders are supposed to be bidding at time t_L, each bidder will have an incentive to bid "a little early," which then unravels the sniping equilibrium.

Turning to a different feature of the auctions we consider, note that it is crucial that these have a fixed end time. Given such a "hard" ending, the bidders targeted by the shill bid-

[6]Of course, if we allow bidders to arrive after time 0, they would necessarily bid after time 0. To determine the optimal bid time relative to the "last minute" starting at 0, we allow entry up to 0, but not later. It is also worth noting that our results do not depend in any way on this restriction. If we did allow entry after time 0, that would not change the equilibrium bid time of bidders who arrive before time 0.

der can snipe the shill bid by delaying their bids. Alternatively, an auction could have a "soft" ending so that if a bid arrives in the last 10 minutes (say), the end time is automatically extended until there is no bidding activity for 10 continuous minutes - a format used by uBid.com. In auctions with a soft ending, even in the absence of proxy bidding, sniping becomes impossible. As Roth and Ockenfels (2002) report, the contrast between the two formats gives rise to interesting differences in observed bidding behavior across auctions. Auctions on eBay (which uses a hard end time), have substantially greater last minute bidding compared to Amazon auctions (these auctions, now defunct, used a soft ending). Since the purpose of late bidding in our model is to snipe the shill bid, and since this is not possible under a soft ending, our results are consistent with this finding.[7]

Finally, note that our central characterization result on the optimal bid time is independent of the actual number of bidders so long as there are at least two. Therefore we can allow for random entry implying that final number of entrants is random – and its realization unknown to bidders when they submit bids. The result fits the environment of online auctions, where a bidder typically does not have precise knowledge about how many others are bidding. There is thus a "robustness" feature of our result similar to that of the weakly dominant bid function in a second-price auction where a bidder's strategy is invariant to the number of opponents the bidder faces in the auction. Later, in section 5, we make use of these robustness properties to construct a pure late bidding equilibrium for which the knowledge of actual number of bidders is immaterial.

Relating to the Literature

In our paper, the bidders targeted by the shill bids want to delay their bids optimally to hide information from the seller. Other papers have considered reasons for bidders to delay bids to hide information from other bidders. Bajari and Hortaçsu (2003) consider a common values setting and assume a (discontinuous) timing structure that implies that an eBay auction is a two stage auction: up to time $t_L - \varepsilon$ it is an open ascending auction, and for the rest of the time it is a sealed bid auction (i.e. in this stage all bids arrive, but no

[7] In our formal model, all buyers arrive before time $t = 0$; see footnote 6. If buyers can arrive after time 0, then there would be some late bidding even in soft end proxy auctions. Holding constant buyer arrival rate, one would nevertheless expect more last minute bidding in fixed end than in soft end auctions which is what the data seem to suggest.

one can respond to any one else's bid). Under this structure, they show that all bidders bidding only at the second stage is an equilibrium. Ockenfels and Roth (2006) study a second model of last minute bidding set in a common values environment with two bidders: an expert and a non-expert. Only the expert knows whether an item is genuine or fake. However the non-expert has a higher value for a genuine item compared to the expert. The expert does not bid if the item is fake. If the item is genuine, it is then clear that the expert might not want to bid early as such bids might reveal to the non-expert that the item is genuine. Assuming the same time structure as in their "collusion" theory discussed above, they show that if the prior probability that the object is fake is high enough, there is an equilibrium in which only the expert bids, and bids only at a "last point of time" time t_L, thus not giving the non-expert the chance to react to this information.

Rasmusen (2006) presents a similar idea in a private values context, with two bidders: one knows own value while the other does not, but could find out at a cost. There is a last point of time t_L and all bids made up to the last point t_L arrive. Under certain values of cost of value discovery, if the auction price rises to a certain point before t_L, the uninformed bidder pays to discover own value. This raises the auction price for a high-valued informed bidder, who therefore wants to bid at the last point t_L to prevent valued discovery by the uninformed bidder.

As noted above, our approach differs from these ideas in that, first, we have a standard private values setting in which (ex ante) symmetric bidders know their own valuations and have no incentive to hide information from other bidders. The reason for their late bidding is their desire to reduce the probability that the seller gets a chance to successfully submit a shill bid. Another difference arises from our modeling of timing. We adopt a timing structure in which bid times and arrival times are chosen continuously, and solve for the optimal bid timing. The continuous timing structure also gives rise to a further comparison with the literature. We could ask how the results in the literature would be modified in a continuous timing framework.

As we discussed above, the "collusive" equilibrium of Ockenfels and Roth (2006) depends crucially on the discontinuous time structure, and cannot arise in a framework of continuous bid and arrival times. However, the other types of equilibria discussed above, in which bidders bid late in a common values framework to discourage aggressive bidding, or where an expert hides information from a non-expert in a common values setting, or a private values setting in which a high value bidder hides information from a bidder who does not know own value, could also arise using our continuous-bid-and-arrival-time framework. In that case, an interesting question would be to determine the optimal bid time for the bidders whose bids have information content, given that they stand to lose if this information is revealed to others *and* now allowing for (at least some of) the latters' subsequent bids to arrive before the auction ends.

Finally, consider the question of shill bidding. We show that shill bidding naturally provides an incentive for targeted bidders to delay their bids. Given that shill bidding is reported to be widespread, last minute bidding arises as a natural counterpart in online auctions. In this work, we are interested in shill bidding only in so far as it provides a natural reason for last minute bidding, and our principal focus is then on deriving optimal bid timing and its consequences for revenue and efficiency. Therefore we chose the simplest setting in which the seller has an incentive to shill bid. In our setting, values are conditionally (conditional on the realization of the distribution) independent. Papers such as Chakraborty and Kosmopoulou (2004), Lamy (2009) examine shill bidding in environments with common values or interdependent values, and show that the presence of shill bidding can reduce the information content of the observed auction prices, and reduce the seller's revenue. Lamy shows how a "shill bidding effect" arises with interdependent values that goes against the usual linkage principle, and therefore favors first price auctions (immune to shill bids) against second price auctions in revenue ranking. However, unless the seller can credibly commit to not submitting shill bids, such bids arise in equilibrium.

The rest of the paper is organized as follows. The next section presents the model. Section 3 presents two preliminary results characterizing the equilibrium bid timing. Section 4 then presents the main result on bid timing by the bidders targeted by shill bidding. Section 5 explicitly constructs an equilibrium to establish existence. Finally, section 6 concludes.

2 The Model

A seller is interested in selling a single unit of an indivisible object. The seller's own value for the object is zero. There are $n \geq 2$ potential bidders. We allow for random bidder arrival and hence n is a random variable (on the set of integers greater than 1). While most of standard auction theory assumes the number of bidders to be common knowledge, it is difficult to justify this assumption for online auctions. We assume here that bidders do not observe n. While our model involves standard Bayesian rational agents who necessarily have priors over n, the equilibrium outcome we are interested in is present in all equilibria *irrespective* of n. Thus the exact nature of the prior belief is unimportant.[8]

Let v_i denote the value of bidder i, $i \in \{1, \ldots, n\}$. The bidders' valuations of the object are private but correlated as follows. With probability $1 - \mu \in (0, 1)$ all bidders' valuations are determined according to independent draws from a distribution F_L on the interval $[0, a]$, and with probability μ, the values are determined as independent draws from a distribution F_H on $[a, b]$. Here $b > a > 0$.[9] We assume that F_L and F_H are continuous distributions with continuous density functions (denoted by f_L and f_H respectively) that are strictly positive on $[0, a]$ and $[a, b]$ respectively. Note that if μ were equal to either 0 or 1, we would be in the independent private value (IPV) setting.

Let Ψ_L and Ψ_H denote the virtual values associated with F_L and F_H respectively, so that

$$\Psi_k(v) = v - \frac{1 - F_k(v)}{f_k(v)}$$

where $k \in \{L, H\}$. We assume Ψ_L and Ψ_H are increasing, i.e. F_L and F_H are regular distributions in the sense of Myerson (1981).

When (it is known that) the types are drawn from distribution F_k, for $k \in \{L, H\}$, the optimal reserve price for the seller is given by R_k where $R_L = \Psi_L^{-1}(0)$ and $R_H = \max\{a, \Psi_H^{-1}(0)\}$. We assume that $R_H > a$, i.e. the optimal reserve price for the distribution F_H is non-trivial.

Let $Y_L(R_L)$ be the expected revenue from setting optimal reserve price when the distribu-

[8]An implication of this is that our result has a certain robustness property not often found in (perfect) Bayesian Nash equilibria of many other auction models.

[9]We could allow for support of F_L to be $[0, a]$ and of F_H to be $[\alpha, b]$ with $\alpha > a$. This generality would, however, add another piece of notation without providing any additional insight.

tion is F_L and $Y_H(R_H)$ the expected revenue from setting the optimal reserve price when the distribution is F_H. We let $Y_H(a)$ denote the expected revenue when valuations are in $[a, b]$ but the auction is without a reserve price (i.e. reserve price trivially set to a).

The seller uses an online auction site to try to sell the item. The auction format is proxy bidding with a hard (i.e. fixed) end time. The seller can post a reserve price at the beginning.

Buyer arrival in the auction is allowed to be random. However, as will be clear from the discussions and analysis below, we do not need to formally model this process. We do assume however, that the buyer arrival is invariant with respect to whether the true valuation distribution is F_H or F_L.

To ensure a non-trivial problem, we need the seller's prior μ to be such that the seller chooses R_L as the initial reserve price. However, this is not enough. To see why, suppose that when the first bid above the reserve price R_L arrives, the seller observes this[10] and thus updates her prior to posterior $\tilde{\mu}$ where $\tilde{\mu} > \mu$.[11] If the seller now wants to switch to the reserve price R_H immediately, then in the actual auction, the seller shill bids (by placing a bid R_H) as soon as the first bid is made in the auction. In that case bidders would have no strict incentive to participate when their values are drawn from F_L , and under even the smallest cost of participation would strictly prefer non-participation. If the bidders do not participate when their valuations are drawn from F_L then the seller would simply choose the optimal reserve price to be R_H and the model becomes completely trivial. In order to rule out this uninteresting case we therefore need to assume that the seller does not immediately want to jump to reserve price R_H when the first bid above the reserve price R_L arrives.

We now give a sufficient condition for the seller not to have such an incentive to switch to R_H immediately. Let $\widehat{\mu} \equiv \dfrac{\mu}{\mu + (1-\mu)\,(1 - F_L(R_L))}$.

This is the updated probability that the bid comes from a bidder drawn from the high

[10]If the seller does not observe this, our problem would be easier. However, the continuous price increase case is only a limiting case, and under any price increment the arrival of the first bid above the reserve price would be observable.

[11] $\widetilde{\mu}$ depends on, among other things, the time when the bid arrives, the stochastic process of buyer arrival, the bidding strategies (with respect to timing of placing of bids) of the bidders, as also the stochastic process according to which bids arrive after being made. We will shortly give a condition, sufficient for our purpose, that does not depend on these factors.

distribution F_H. The (sufficient) condition is given by

$$\widehat{\mu} Y_H(a) + (1 - \widehat{\mu}) Y_L(R_L) \geqslant \widehat{\mu} Y_H(R_H) \tag{2.1}$$

Note that this condition is sufficient but not necessary. To see why note that since the seller can change the reserve price by shill bidding at any time before the auction ends, not changing the reserve price is therefore akin to keeping an option alive (unexercised). If the seller does not shill bid immediately after the first bid, but waits till the second bid arrives from the high distribution so that the price in the auction moves above a, the seller would then know for sure that the distribution is F_H. Switching to R_H at this point results in expected payoff $Y_H(R_H)$ with a positive probability (the probability that the shill bid does not get lost). Given the updated belief $\widehat{\mu}$, while the right hand side of (2.1) is the expected revenue from switching to R_H immediately, the expected revenue for a seller who waits till the second bid above a arrives is higher than the left hand side.

The timing of the auction is as follows. The auction starts at $-T \leqslant -1$ and ends at 1. Bidders arrive randomly over $[-T, 0]$. A bidder arriving at time $s \in [-T, 0]$ can bid (i.e. submit a proxy bid) at any time $t \in [s, 1]$. A crucial part of the model is the continuous bid arrival process. We assume that the arrival time of a bid submitted at time $t \in [-T, 1]$ is uniformly distributed on $[t, t+1]$.

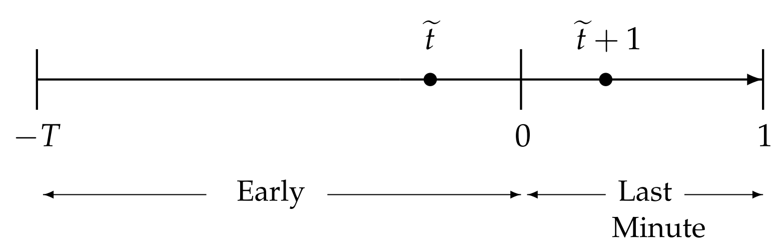

Figure 1: Bid timing and arrival. The auction starts at $-T \leqslant -1$ and ends at 1. Bidders arrive randomly over $[-T, 0]$. The arrival time of a bid made at time $\widetilde{t} \in [-T, 1]$ is uniformly distributed on the time interval $[\widetilde{t}, \widetilde{t}+1]$. Early bids arrive with certainty, while a bid at any time q inside the "last minute" period gets lost with probability q and with probability $(1-q)$ the arrival time is distributed uniformly on $[q, 1]$.

A bid submitted at time $t \in [-T, 0)$ is called an "early bid" and early bids therefore arrive with certainty. Bids submitted at $t \in [0, 1]$ are "last minute" bids. A last minute bid submitted at $t = 0$ (at the cusp of the last minute period) still arrives with probability 1, but any bid at $t = q > 0$ (inside the last minute) is lost[12] with probability q, and with probability $(1 - q)$ the arrival time is distributed uniformly on $[q, 1]$.

Let p_t denote the auction price at time t, and h_t the history of auction prices up to (but not including) time t. Note that h_t is thus a step function over the interval $[-1, t)$.

The seller can submit a bid (the shill bid). Let $\sigma_0(t)$ be the seller's strategy of submitting the shill bid. Formally, $\sigma_0(t) : \{h_t, p_t\} \rightarrow \{0, 1\}$ such that if there is a τ such that $\sigma_0(\tau) = 1$ then $\sigma_0(t) = 1$ for all $t \geq \tau$. In words, τ is the instance when the seller decides to submit the shill bid[13]. Note that the range of the seller's strategy is $\{0, 1\}$, since if the seller decides to shill bid, the amount of the bid should be R_H. Hence the crucial aspect of the shill bid is not the amount, but the *timing*, of the shill bid.[14] As we show later, in equilibrium, seller waits till the auction price p_t goes above a (in other words two bids above a reach) to submit the shill bid. Of course it is possible that the first two bids above a are also above R_H, in which case the seller will not need to shill bid.

We now describe bidder strategies. As is standard, h_{-T} is the null history and let $p_{-T} = R_L$. For any t, let $\mathcal{P}_t$ be the set of numbers (strictly) greater than p_t and define the set $B_t = \{\emptyset \cup \mathcal{P}_t\}$. Strategy of bidder i, having valuation v_i and having arrived at time s is then denoted by $\sigma_i^s(t\,|v_i)$ where $\sigma_i^s(t\,|v_i) : \{h_t, p_t\} \rightarrow B_t$. Time periods when the bidder chooses $\emptyset$ are periods when the bidder is *inactive*, in other words, these are the periods when the bidder places no bids. Choosing any $b \in B_t$ where $b \in \mathcal{P}_t$ refers to submitting a proxy bid equal to the amount b. As in standard auctions, a bidder does not

[12] Being "lost" simply means that the bid fails to arrive by the time the auction ends.

[13] Using the terminology from the mechanism design literature, the (officially) chosen reserve price is part of the mechanism. The seller's strategy refers to her action only with regard to the shill bidding.

[14] We are making an implicit assumption that the seller submits shill bids using one other identity. In general, the seller could use multiple alternative identities to submit several bids of R_H when the relevant information event occurs. However, since shill bidding is illegal, using multiple alternative identities also involves costs that we leave out in our formal model. There are at least two reasons for such costs to arise. First, submitting a bid requires having an account with the auction site backed by a bank account or credit card, and hence multiple identities bidding would require having multiple accounts requiring different credit cards or bank accounts that cannot easily be traced to a single identity. Second, multiple bids equal (or even approximately equal) to R_H submitted around the same time would raise the possibility of triggering an investigation by the auction site. Finally, allowing multiple shill bids would only reinforce our main result (Theorem 1) since any increase in the chances of the shill bid being successful only reduces the incentives of the bidders to delay submitting their bids beyond the time instant $t = 0$.

have any incentive (in equilibrium) to submit a bid strictly greater than his valuation, but may submit bids that are strictly less than his valuation. Anticipating this equilibrium property, we call a bid **serious** if the bid is equal to the bidder's valuation, and bids less than that are termed **non-serious** bids.

3 Equilibrium Characterization

We start analyzing the auction specified above by establishing several properties of the equilibrium bid times for both the buyers and the seller. We then establish the main result on the timing of bids by buyers in the next section. Finally, section 5 demonstrates existence by explicitly constructing an equilibrium.

In this section, we record two preliminary results. First, Proposition 1 shows the difference in the incentives of bidders with respect to the timing of their serious bids when they are from the two different distributions. When the distribution is F_L, the only (real) concern for the bidders is that their bids reach with probability 1. Hence the latest they should bid is at $t = 0$. When the distribution is F_H, bidders now would want to lower the chances of a successful shill bid by the seller and hence any serious bid at time $t < 0$ is dominated by a bid placed at $t = 0$ since each of these will reach with probability 1 but the earlier bid, by reaching earlier, also increases the chances that the seller is able to shill bid successfully. Proposition 1 leaves open the possibility that when the distribution is F_H, some serious bids may be submitted at some time $t > 0$. However, we show in our main result, Theorem 1, below that this does not happen in equilibrium: in all equilibria bidders with types drawn from the high distribution F_H submit serious bids at and only at time $t = 0$.

Proposition 1. *In any equilibrium, types drawn from F_L bid at some $t \leq 0$ and types drawn from F_H submit serious bids at $t \geq 0$.*

The next result, in addition to its importance in understanding how shills and snipes are connected, also helps to reduce dramatically the complexity of the task of finding the seller's optimal strategy. It shows that, in equilibrium, the seller submits the shill bid only when she is certain that the true distribution is F_H.

Proposition 2. *In any equilibrium, the seller submits the shill bid of R_H only when the auction price moves above a, i.e. when two bids above a arrive. That is* $\inf\{\tau : \sigma_0(\tau) = 1\}$ *is such that* $p_\tau \geqslant a$.

Proof: Let Z_t denote the set consisting of histories h_t and auction prices p_t such that $p_t < a$. Suppose on the contrary that there is time $\widetilde{t} > 0$ and a subset $\widetilde{Z}_{\widetilde{t}}$ of $Z_{\widetilde{t}}$ such that if the history and prices belong to the set $\widetilde{Z}_{\widetilde{t}}$ then the seller submits the shill bid. Now, for any bidder from the distribution F_H, and for history h_s, where $s \geqslant \tilde{t} - 1$, if some bid b[15] creates a positive probability of reaching some node $z_{\widetilde{t}} \in \widetilde{Z}_{\widetilde{t}}$, then the best response of this bidder is not to place the bid b. But then the node $z_{\widetilde{t}}$ being reached must be due to the bid b by types from distribution F_L. This should, if anything, reduce the seller's posterior belief that types are drawn from the high distribution. If it was not the seller's best response, with higher beliefs, to shill bid earlier, it can't be a best response to shill bid now. The seller should thus deviate from the proposed strategy and not submit the shill bid at time $\widetilde{t}$, which contradicts the supposition that this was part of the seller's equilibrium strategy.$\|$

Note that if in equilibrium the bidders from distribution F_L submit serious bids before $t = 0$, then bidders from distribution F_H may have to submit non-serious bids before $t = 0$. Otherwise, depending on the arrival process of the bidders, lack of early bidding may lead the seller to believe with high enough probability that the true distribution is F_H and hence submit a shill bid even before the auction price increases beyond a. Hence, some equilibrium may require bidders from F_H to submit non-serious bids to mimic the actions of the bidders from F_L over the period $[-T, 0)$ (the early period).

4 Bid Timing: the Main Result

Proposition 1 tells us all we need about the bid timing of bidders drawn from the distribution F_L. Recall that we call a bid serious if it is equal to the valuation of the bidder. We now consider the equilibrium bid-time strategies for submitting serious bids by types drawn from F_H.

[15]The bid b might be a serious bid or a non-serious bid including a zero bid (which is the same as not placing any bid at all).

Note that the only reason a bidder would submit his serious bid at time $t > 0$ is to delay the time at which the auction price exceeds a (i.e. the time of arrival of the second bid) which is the time when the seller submits the shill bid. So if bidder 1 (say) has a plan to submit his serious bid at $t > 0$, but the auction price exceeds a for the first time at $t' < t$ (two bids from other bidders arrive before t, with the second bid arriving at t'), bidder 1 should optimally submit his bid at t', as further delay reduces payoff (by reducing the probability of reaching) but has no further impact on the probability of arrival of the shill bid.[16] Also, of course, if the auction price exceeds that value of bidder 1 before time t, bidder 1 plays no further role in the auction process.

Let $Q_i(v_i)$ denote the planned time to submit the serious bid for bidder i with value v_i, so long the auction price has not already exceeded a. That is

$$Q_i(v_i) = \inf\{t : \sigma_i(v_i) = v_i\} \quad \text{for } \{h_t,\, p_t\} \text{ such that } p_t \leq a.$$

From Proposition 1, we know that $Q_i(v_i) \geqslant 0$ for all i.

The following theorem is the central result of the paper.

Theorem 1. *Suppose bidders draw types from the distribution F_H. For any $n \geq 2$, for all $i \in \{1, \ldots, n\}$, the only bid-time consistent with equilibrium is for all bidders to submit a serious bid exactly at time 0, i.e. for all $i \in \{1, \ldots, n\}$, $Q_i(v_i) = 0$ for all $v_i \in [a, b]$.*

Note, as mentioned earlier, that in the auction, the number of bidders, n, is not known. However, since the above result is true for any n, it follows that $Q_i(v_i) = 0$ continues to be the best response in the actual auction.

Terminology In the rest of this section, we discuss serious bids by types drawn from the distribution F_H. Since there is no scope for confusion, for economy of expression we refer to these simply as bids.

[16]Note that with $n = 2$, any planned bid time is also the actual bid time as there is only one other bidder. In this case calculations are simpler than the more general $n > 2$ case.

Outline of the proof The proof involves several steps, which we present below. Before continuing with the proof, let us first sketch informally the main steps of the argument.

• For any two instances of time $t' > t'' \geq 0$, there are two effects with respect to bidding at time t' versus at time t''. The earlier bid, i.e. the one placed at time t'' reaches with higher probability. However, it also increases the probability of a successful shill bid. Now, it seems intuitive that bidders with higher valuations have more to lose from their bids not arriving. Further, it is clear that bidders with valuations below R_H have more to lose if the shill bid arrives. Hence, it is useful to categorize types of bidders into two groups: those whose valuations are greater than R_H and those whose valuations are less than R_H. We consider these two cases separately.

• The derivative of expected payoff with respect to bid time $q \geqslant 0$ is the marginal benefit (which could be negative) from bid delay. First, we show that for types $v \geqslant R_H$, the marginal benefit from bid delay decreases in bidder valuation. Second, we show that for types $v \in [a, R_H]$, delaying the bid beyond time $t = 0$ is strictly worse than bidding at $t = 0$,. In other words, the marginal benefit from bid delay is strictly negative for types up to R_H. This, coupled with the fact that the marginal benefit for values above R_H is decreasing in value, implies that the marginal benefit of bid delay is strictly negative for all values. Therefore it must be the case that no bidder would want to wait beyond $t = 0$ to submit his bid.

• To show the first fact above is straightforward. Once we write down the expected payoff, and prove some properties of a few functions, we can establish that the marginal benefit from bid delay decreases in bidder valuation by straightforward differentiation.

• The proof of the second fact, that the marginal benefit from bid delay is strictly negative for types $v < R_H$, requires a slightly longer string of arguments and we prove it in two steps. Proposition 3 shows that if all bidders other than (say) bidder 1 bid at time $t = 0$, bidder 1's strict best response is to bid at $t = 0$. The argument is now completed by showing the following. First, *if* there is an equilibrium where some type of some bidder j wait beyond $t = 0$ to submit his bid, then there is a smallest type $v_* > a$, such that for all bidders, types $v < v_*$ strictly prefer to bid at $t = 0$ and hence only types v_* and above for some bidders may have an incentive to wait if at all. Suppose $v_* < R_H$. Given that all types of all bidders with valuations less than v_* place their serious bid at $t = 0$, it then follows (from using arguments in the proof of Proposition 3) that for any bidder, type

v_* has a strict incentive not to delay the bid beyond $t = 0$. By continuity, types slightly above v_* have a strict incentive not to delay their bids beyond $t = 0$ as well. Thus it is not possible to have $v_* < R_H$. But if v_* is at least R_H, this shows that the marginal benefit from bid delay is strictly negative for all types $v \leqslant R_H$, which establishes the second fact and completes the proof.

4.1 Preliminaries

We now proceed with the analysis. Let N be the set of bidders and let N_{-1} denote the set of bidders other than bidder 1. Consider a subset of this set containing $k \leqslant n-1$ bidders. There are $\binom{n-1}{k}$ such subsets. Let Γ_k denote the set of these subsets, and let γ_k denote a typical element of Γ_k, i.e. γ_k is a particular subset of k bidders.

Let $\vec{Q}_{-1} \equiv (Q_2(\cdot), \ldots, Q_n(\cdot))$.

Let $\beta(q, \vec{Q}_{-1}, \gamma_k)$ denote the probability that the bids of the k bidders $j \in \gamma_k$ arrive and the bids of the $n-k-1$ bidders $j \in N_{-1} \backslash \gamma_k$ get lost. Let $\phi(q, \vec{Q}_{-1}, \gamma_k)$ denote the probability that bidder 1's bid arrives, and let $\psi(q, \vec{Q}_{-1}, \gamma_k)$ denote the probability that bidder 1's bid arrives and the shill bid does not arrive.

For economy of notation, let

$$\beta_k \equiv \beta(q, \vec{Q}_{-1}, \gamma_k)$$
$$\phi_k \equiv \phi(q, \vec{Q}_{-1}, \gamma_k)$$
$$\psi_k \equiv \psi(q, \vec{Q}_{-1}, \gamma_k)$$

The following Lemma shows some useful properties of the functions β_k and ϕ_k.

Lemma 1. $\dfrac{\partial \beta_k}{\partial q} \leqslant 0$ *and* $\dfrac{\partial \phi_k}{\partial q} \leqslant 0$.

Proof: First, it is obvious that ϕ_k is non-increasing in q. Now, any bidder $j \neq 1$ bids either at $Q_j(v_j)$ or at the point of arrival of the second bid, if this is earlier. Therefore if the arrival distribution of the second bid is shifted to the right, the bidding time of bidder j increases weakly, and therefore the probability that the bid of j arrives decreases weakly. If $q \geqslant \max_{j \in \gamma_k} Q_j(v_j)$, clearly increasing q further has no impact on β_k. If however $q < \max_{k \in \gamma_k} Q_k(v_k)$, a rise in q shifts the distribution of arrival time of the second bid to the right, reducing the probability that the bid of j arrives. Therefore β_k is non-increasing in q. $\|$

Let

$$y_k \equiv \max_{j \in \gamma_k} \{v_j\}$$

i.e. y_k is the highest value amongst the k bids that reach out of the $n-1$ bids of bidders other than bidder 1.

4.2 Properties of payoff for types above R_H

Let us first write down the expected payoff of bidder 1 for the case $v_1 \geqslant R_H$. Define the function S^U as follows.

$$\begin{aligned} S^U\left(q, \vec{Q}_{-1}, \gamma_k\right) &= \Pr\Big(R_H \leqslant y_k \leqslant v_1\Big) E\left(\sum_{\gamma_k \in \Gamma_k} \Big[\beta_k \phi_k (v_1 - y_k)\Big] \mid R_H \leqslant y_k \leqslant v_1\right) \\ &+ \Pr\Big(a \leqslant y_k \leqslant R_H\Big) E\left(\sum_{\gamma_k \in \Gamma_k} \Big[\beta_k \Big((\phi_k - \psi_k)(v_1 - R_H) + \psi_k (v_1 - y_k)\Big)\Big] \mid a \leqslant y_k \leqslant R_H\right) \end{aligned} \tag{4.1}$$

$S^U(q, \vec{Q}, \gamma_k)$ is the expected payoff of bidder 1 for $v_1 \geqslant R_H$, when $k > 0$ other bids reach. The first term captures the case $y_k \geqslant R_H$ (in which case whether the shill bid arrives is irrelevant), and the second term captures the case $a \leqslant y_k \leqslant R_H$, and the shill bid either arrives (in which case bidder 1 pays R_H) or it does not arrive (in which case the price is y_k).

The expected payoff of bidder 1 for any $v_1 \geqslant R_H$ is given by

$$\begin{aligned} \pi_1^U \equiv \pi_1 | v_1 \geqslant R_H &= \sum_{k=1}^{n-1} S^U\left(q, \vec{Q}_{-1}, \gamma_k\right) \\ &+ \Pr\left(a \leqslant v_j \leqslant b \; \forall j \neq 1\right) E\left(\prod_{j=2}^{n} Q_j(v_j)(1-q)(v_1 - R_L) \,\Big|\, a \leqslant v_j \leqslant b \; \forall j \neq 1\right) \end{aligned} \tag{4.2}$$

where the second term captures the case in which no other bids reach. Obviously, given that we are considering the types drawn from F_H, $\Pr\left(a \leqslant v_j \leqslant b \; \forall j \neq 1\right) = 1$, and therefore the second term is simply $E\left(\prod_{j=2}^{n} Q_j(v_j)(1-q)(v_1 - R_L)\right)$.

The following result now shows that the marginal benefit from bid delay given by $\frac{\partial \pi_1^U}{\partial q}$ (which could be negative) is decreasing in v_1.

Lemma 2. $\dfrac{\partial \pi_1^U}{\partial q}$ *decreases in* v_1*, i.e.*

$$\begin{cases} \dfrac{\partial}{\partial v_1}\left(\dfrac{\partial \pi_1^U}{\partial q}\right) \leqslant 0 & \text{if } \dfrac{\partial \pi_1^U}{\partial q} \geqslant 0 \\ \\ \dfrac{\partial}{\partial v_1}\left|\dfrac{\partial \pi_1^U}{\partial q}\right| \geqslant 0 & \text{if } \dfrac{\partial \pi_1^U}{\partial q} < 0 \end{cases}$$

The proof, which is relegated to the appendix, requires straightforward (if somewhat lengthy) differentiation, and then making use of Lemma 1.

4.3 Properties of payoff for types below R_H

Next, define the function S^L as follows.

$$S^L\left(q, \vec{Q}_{-1}, \gamma_k\right) = \Pr\Big(a \leqslant y_k \leqslant v_1\Big) E\left(\sum_{\gamma_k \in \Gamma_k} \Big[\beta_k \psi_k (v_1 - y_k)\Big] \mid a \leqslant y_k \leqslant R_H\right) \quad (4.3)$$

$S^L(q, \vec{Q}, \gamma_k)$ is the expected payoff of bidder 1 with value $v_1 < R_H$ when k other bids reach. The expected payoff of bidder 1 for any such $v_1 < R_H$ is given by

$$\pi_1^L \equiv \pi_1 | v_1 < R_H = \sum_{k=1}^{n-1} S^L\left(q, \vec{Q}_{-1}, \gamma_k\right) + E\left(\prod_{j=2}^{n} Q_j(v_j)(1-q)(v_1 - R_L)\right) \quad (4.4)$$

where the second term captures the case in which no other bids reach. Therefore

$$\frac{\partial \pi_1^L}{\partial q} = \sum_{k=1}^{n-1} \Pr\Big(a \leqslant y_k \leqslant v_1\Big) E\left(\sum_{\gamma_k \in \Gamma_k} \Big[\Big(\beta_k \frac{\partial \psi_k}{\partial q} + \psi_k \frac{\partial \beta_k}{\partial q}\Big)(v_1 - y_k)\Big] \mid a \leqslant y_k \leqslant v_1\right)$$
$$- E\left(\prod_{j=2}^{n} Q_j(v_j)(v_1 - R_L)\right) \quad (4.5)$$

Since $\frac{\partial}{\partial v_1}\left(\frac{\partial \pi_1^U}{\partial q}\right) \leqslant 0$, it follows that if the derivative with respect to q is always negative for $v \leqslant R_H$, then it is always negative for all values of v_1 and therefore bidding at $t = 0$ is the only solution.

To show that this is the case, we first show that if others bid at $t = 0$, bidding at $t = 0$ is the unique best response for any bidder.

4.3.1 Step 1: Bidding at $t = 0$ is a strict equilibrium

The following Lemma records the expressions for two important probability terms, g_1 and g_2.

Lemma 3. *Suppose the $n-1$ bidders other than bidder 1 bid at $t = 0$ and bidder 1 bids at time $q \geqslant 0$. The probability that the second bid arrives at time t is given by $g_1(t, n)$ for $t < q$ and $g_2(t, q, n)$ for $t \geqslant q$ where*

$$g_1(t, n) = (n-1)(n-2)t(1-t)^{n-3}$$
$$g_2(t, q, n) = \frac{(n-1)(1-t)^{n-2}}{1-q}(nt - q)$$

The proof is in the appendix. Here, let us simply show the simpler case of $n = 2$. In this case, the only relevant case is that the second bid arrives at some time $t \geqslant q$, so that we need to derive only $g_2(t, q, 2)$.

Let $Z_1(t, q)$ denote the distribution of arrival time of the bid by bidder 1. Since bidder 1 submits the bid at $q \geqslant 0$, Z_1 has support $[q, 1]$. Similarly, let $Z_2(t)$ denote the distribution of arrival time of the bid by bidder 2, which has support $[0, 1]$. Let $\xi_1(t, q)$ and $\xi_2(t)$ denote the corresponding density functions, respectively. Note that $\xi_1(t, q) = 1/(1-q)$ and $\xi_2(t) = 1$. Further, $Z_1(t, q) = \frac{t-q}{1-q}$ and $Z_2(t) = t$. The probability that the second bid arrives at t (i.e. one bid arrives before t and the second at t) is $g_2(t, q, 2) = Z_2(t)\xi_1(t, q) + Z_1(t, q)\xi_2(t) = \frac{t}{1-q} + \frac{t-q}{1-q} = \frac{2t-q}{1-q}$. One can similarly derive g_1 and g_2 for $n > 2$ bidders, as shown in the appendix.

To continue with the general ($n \geq 2$) case, since all bidders $j \neq 1$ bid at $t = 0$, all such bids reach. We now calculate the probability ψ_{n-1} that bidder 1's bid arrives and the shill bid does not arrive. Note that the shill bid is submitted at the time t, i.e. the instance when the second bid arrives, and therefore it gets lost with probability t. Bidder 1's bid is submitted at t if $t < q$ and at q otherwise.

$$\begin{aligned}\psi_{n-1} &= \int_0^q (1-t)tg_1(t,n)dt + \int_q^1 (1-q)tg_2(t,q,n)dt \\ &= (n-1)(n-2)\int_0^q t^2(1-t)^{n-2}dt + (n-1)\int_q^1 (1-t)^{n-2}\, t\,(nt-q)dt \\ &= \frac{1}{n(n+1)}\left(2(n-2) + (1-q)^{n-1}\left(4(1-q) + q(3n-1) + (n-1)^2q^2\right)\right) \qquad (4.6)\end{aligned}$$

Proposition 3. *If bidders other than 1 bid at $t = 0$ then bidding at $t = 0$ is the unique best response for bidder 1.*

Proof: If other bidders bid at $t = 0$, the payoff of bidder 1 for $v_1 \leqslant R_H$ is given by

$$\pi_1^L = \int_a^{v_1} (v_1 - y_{n-1})\psi_{n-1}dG_{n-1}$$

where G_{n-1} is the distribution of y_{n-1} (which is the highest value among $(n-1)$ other bidders). Therefore

$$\frac{\partial \pi_1^L}{\partial q} = \int_a^{v_1} (v_1 - y_{n-1})\frac{\partial \psi_{n-1}}{\partial q}dG_{n-1}$$

Now, from the expression for ψ_{n-1} given by equation (4.6),

$$\frac{\partial \psi_{n-1}}{\partial q} = -(1-q)^{n-2}\left((n-1)q^2 + \frac{(1-q)(1+(n-1)q)}{n}\right) < 0$$

Therefore $\frac{\partial \pi_1^L}{\partial q} < 0$ for $v_1 \in (a, R_H]$. At $v_1 = R_H$, $\pi_1^L = \pi_1^U$. It follows that $\frac{\partial \pi_1^U}{\partial q} < 0$ at $v_1 = R_H$. Further, for $v_1 \geqslant R_H$, we already know that $\frac{\partial}{\partial v_1}\left(\frac{\partial \pi_1^U}{\partial q}\right) \leqslant 0$. It follows that the derivative of payoff with respect to q is strictly negative for all values of v_1 and therefore bidding at $t = 0$ is the only solution.$\|$

4.3.2 Step 2: Every equilibrium involves serious bidding no later than at $t = 0$

The next Proposition shows that *if* there are some types of bidder j who bid later than at $t = 0$, then these types are above a lower bound $v_j^* > a$.

Proposition 4. *For all $j \in \{1, \ldots, n\}$, there exists $v_j^* > a$ such that it is a dominant strategy for bidder j to bid at time 0 for $v \leqslant v_j^*$.*

Proof: From Proposition 3 above, we know that if $Q_j(v_j) = 0$ for all $j \neq 1$ and all $v_j \in [a, b]$, then the best response is $q = 0$ for all $v_1 \in [a, b]$. Next, suppose for some j, $Q_j(v_j) > 0$ on any sub-interval of $[a, b]$. Consider the derivative of π_1^L with respect to q, given by equation (4.5). Note that as v_1 goes to a, the first term goes to zero, and therefore

$$\lim_{v_1 \to a} \frac{\partial \pi_1^L}{\partial q} = -E\left(\prod_{j=2}^{n} Q_j(v_j)(a - R_L)\right) < 0.$$

Therefore, for values of v_1 close to a, bidder 1 bids optimally at $t = 0$. In other words, for v_1 close to a, the best response is $q = 0$.

It follows that for all $j \in \{1, \ldots, n\}$, there exists $v_j^* > a$ such that irrespective of the other bidders' bidding time, bidder j optimally bids at time 0 for $v \leqslant v_j^*$.‖

Let

$$v_* = \min_j v_j^*$$

We know that for all $j \in \{1, \ldots, n\}$, bidder j bids at $t = 0$ for values $v_j \leqslant v_*$. Let M be the subset of the set of bidders such that for any bidder $\ell \in M$, $v_\ell^* = v_*$. Clearly, M is non-empty.

Suppose $v_* < R_H$.

Consider bidder $\ell \in M$ of type v_*. The second term in equation (4.5) is non-positive, and is negative whenever some $Q_j(v_j) > 0$. Further, since, for all other bidders, types weakly lower than v_* bid at time 0, it follows from Proposition 3 (by renaming bidder 1 as bidder ℓ) that the first term in negative as well. Hence, $\frac{\partial \pi_\ell^L}{\partial q} < 0$ at $v_\ell = v_*$. And, from equation (4.5), $\frac{\partial \pi_\ell^L}{\partial q}$ is continuous is v_ℓ. Therefore, by continuity, for values of v_ℓ just above v_* it is still the case that $\frac{\partial \pi_\ell^L}{\partial q} < 0$, and therefore the best response of these types of bidder ℓ is to bid at $t = 0$. Since this is true of any bidder $\ell \in M$, this implies that $\min_j v_j^* > v_*$. Contradiction.

Therefore we must have $v_* \nless R_H$, i.e. v_* is at least equal to R_H. Since v_* is at least R_H, it follows that all bidders, and therefore all bidders other than 1, bid at time $t = 0$ for all types in $(a, R_H]$. It follows, using Proposition 3, that $\frac{\partial \pi_1^L}{\partial q} < 0$ at $v_1 \in (a, R_H]$.

Since, for $v_1 \geqslant R_H$, $\frac{\partial \pi_1^U}{\partial q}$ is decreasing in v_1, and since we have now shown that the derivative is always negative for $v \leqslant R_H$, it follows that the derivative is always negative for all values of v_1 and therefore all types of all bidders bidding at $t = 0$ is the only equilibrium. This completes the proof of Theorem 1.

4.4 Revenue and Efficiency

An *optimal auction* maximizes the seller's expected revenue. Further, an auction is *fully efficient* if the object is sold to the highest value buyer whenever this highest value exceeds the value of the seller. Full efficiency obtains in an auction in which the highest value bidder wins and the reserve price is the same as the seller's value. Note that an optimal auction is typically not fully efficient, as the optimal reserve price is typically higher than the seller's value.

Our main result shows that in all equilibria, bidders drawn from the distribution F_H submit serious bids at and only at $t = 0$. We also know that in all equilibria, bidders drawn from F_L submit serious bids at or before $t = 0$. It follows that, even though the seller's shill bid is sniped in equilibrium in the sense that it is necessarily made after time 0 and therefore gets lost with positive probability, all bids from actual bidders reach with probability 1. Hence in all equilibria, when the object is sold to a bidder, it goes to the bidder with the highest value. Further, the shill bid gets lost with positive probability implying that full efficiency obtains for value distribution F_H with positive probability. Thus the auctions being considered in the paper are more efficient, but generate less expected revenue, than the optimal auction when the seller knows the distribution from which values are drawn. Now, an auction with a flexible end time (the auction duration is automatically extended by, say, 5 minutes, if a bid arrives in the last 5 minutes of the auction) allows the seller to implement the optimal auction. This is because with a flexible end time it is impossible to prevent the seller from successfully shill bidding. It follows that auctions with a hard end-time – the format considered here – generate less expected revenue than auctions with a soft ending. Thus, for a *given distribution of buyers*, auctions with a flexible end time are less efficient, but generate greater revenue, than auctions with a fixed end time.[17]

5 Constructing an Equilibrium

The discussion so far has focused on properties of equilibria. In this section, we establish existence by explicitly constructing an equilibrium. We know that types drawn from the high distribution F_H necessarily place serious bids at $t = 0$. We construct a simple equilibrium below which features pure late bidding in the sense that all bidders from either distribution bid at $t = 0$.

Consider the following strategies.

- Irrespective of the distribution from which values are drawn, each bidder submits a serious bid equal to true value at time $t = 0$, and no bidder (again from either distribution) submits any bid prior to $t = 0$.
- The seller posts a reserve price R_L initially (at time $-T$). Subsequently, if no bids reach before time $t = 0$,[18]
  - the seller submits a shill bid of R_H at the instance the auction price goes (weakly) above a (i.e. two bids above a reach), and
  - the seller does not submit a shill bid if the auction price does not reach a.
- If one or more bids above R_L arrive at some time $t < 0$ (i.e. the auction price of R_L becomes "active" or moves above R_L), the seller updates her posterior belief that the distribution is F_H to 1 and submits the shill bid R_H immediately.

[17] However, entry incentives might differ. See the conclusion for more on this issue.

[18] When the first bid at or above R_L reaches, the auction price of R_L becomes "active" (this is just a way of saying that the seller observes the first bid – see footnote 10.). With subsequent bids the auction price goes above R_L. Therefore no bids reaching before time 0 is the event that the auction price of R_L does not become active before time 0.

Let us show that these strategies form an equilibrium.

First, consider the case in which types are drawn from the low distribution F_L. In this case the bidders face a second price auction with reserve price R_L, and submitting a bid equal to true value is the weakly dominant strategy. The specified bid time of $t = 0$ is such that all bids reach with certainty. Any later bid time for a serious bid is strictly worse as bids get lost with positive probability. Given the seller's strategy, any earlier bid time (for either serious or non-serious bids above R_L) is also strictly worse, as this gives rise to a positive probability that a shill bid of R_H is placed and reaches (in which case there is no sale, and each bidder gets a zero payoff). Therefore, given the seller's strategy, the specified strategy for bidders with types drawn from F_L is a best response.

Next, consider the case in which types are drawn from the high distribution F_H. As shown earlier, in this case the unique equilibrium bid time for serious bids is $t = 0$. Further, submitting one or more non-serious bids above R_L at any time before $t = 0$ is strictly worse as this simply raises the probability that the shill bid is placed earlier, reducing expected payoff. Finally, these bidders face a second price auction with either a reserve price R_H if the shill bid reaches, or no reserve price if the shill bid is lost (i.e. a trivial reserve price of a). In either case, bidding true value is as usual the weakly dominant strategy.

Finally, consider the seller's strategy. Condition (2.1) ensures that the seller does not want to submit a shill bid after only one bid is observed (this is when the auction price R_L becomes active) at some time $t > 0$. When the auction price rises above a (i.e. when two bids above a arrive), the seller knows for sure that the actual distribution is F_H, and therefore submitting a shill bid at this instance is a best response. Further, the auction price never goes above a only if the distribution is F_L, in which case it is optimal not to place a shill bid. The seller's strategy of shill bidding (and hence updating reserve price) if a bid arrives before $t = 0$ is optimal given her posterior belief in the off-the-equilibrium event when a bid has been observed before $t = 0$.

It follows from these that the strategies above form an equilibrium.

6 Conclusion

Last minute bidding is a widely observed phenomenon in online auctions, many of which fit the private values model well. We provide an explanation for such biding behavior that does not rely on any discontinuity of the bid arrival process. Our work also clarifies the role of shill bidding in a private values environment: while it is irrelevant in the independent private values setup, shill bidding can be useful to the seller when values are correlated since it essentially allows the seller to adjust the reserve price mid-auction. As we show, bidders who are targeted by the shill bid then have an incentive to delay events that lead the seller to learn about the value distribution, reducing the chances of a successful shill bid. In other words, bidders bid late because they want to snipe the shill bid. We allow a continuous choice of bid times and a continuous arrival process for submitted bids. Our main result shows that there is a unique equilibrium bid-time for targeted bidders to submit serious bids: all such bidders bid at the last point of time (which, in our model, is time 0) at which their bid still reaches with probability 1. The seller's shill bid is then triggered at a time so that it necessarily gets lost with positive probability. In this sense last minute bidding serves to snipe the shill bid.

The result is true for any number of bidders above 1, so that knowledge of actual number of bidders is immaterial. We can therefore allow for random entry, so that no bidder knows the precise number of other bidders - a setting natural for online auctions.

As noted above, we show that while bidders do not place their bids early in the auction, they still submit the bids at the latest instance such that their bids arrive with probability 1. An interesting implication of this is that in all equilibria, when the object is sold to a bidder, it goes to the bidder with the highest value. Further, since the shill bid does not arrive with strictly positive probability, these auctions are more efficient but generate less revenue than the optimal auction when the distribution from which bidder valuations are drawn is commonly known. In contrast to auctions with a fixed end time that we consider here, those with a flexible end time always allows the seller to shill bid successfully. It follows that the fixed end time format is more efficient and generate less expected revenue compared to the flexible ending format.

In other words, if we fix the distribution of buyers, auctions with a fixed end time are less attractive to sellers, and therefore also generate less revenue for the auction site. An interesting question then emerges as to why these auctions are used by behemoths of online auctions like eBay. While this is beyond the scope of our formal analysis, we believe that the key to this lies in entry incentives. Since, ceteris paribus, buyers obtain more surplus from auctions with a fixed end-time, it is clear that such an auction would be more attractive for them. With endogenous buyer entry, it is then not clear that the flexible end-time auction sites necessarily generate more expected revenue for sellers. In general, we would expect factors like diversity in the nature of the objects being sold, as well as differences in prior beliefs regarding distributions from which buyer valuations are drawn to affect how different groups of buyers and sellers self-select themselves to participate in a variety of online auction sites using different auction formats. Studying competition amongst online auction sites using different auction mechanisms in environments characterized by rich heterogeneity of buyers and sellers should be an interesting and important topic for future research.

7 Appendix

7.1 Proof of Lemma 2

To save on notational clutter, let

$$S_k^U \equiv S^U\left(q, \vec{Q}_{-1}, \gamma_k\right)$$

First, consider the derivative of S_k^U with respect to q:

$$\begin{aligned}\frac{\partial}{\partial q}S_k^U = \Pr\Big(R_H \leqslant y_k \leqslant v_1\Big)E\left(\sum_{\gamma_k\in\Gamma_k}\left[\left(\beta_k\frac{\partial\phi_k}{\partial q}+\phi_k\frac{\partial\beta_k}{\partial q}\right)(v_1-y_k)\right] \mid R_H \leqslant y_k \leqslant v_1\right)\\ +\Pr\Big(a \leqslant y_k \leqslant R_H\Big)\left[E\left(\sum_{\gamma_k\in\Gamma_k}\left[\frac{\partial\beta_k}{\partial q}\Big(\phi_k(v_1-R_H)+\psi_k(R_H-y_k)\Big)\right] \mid a \leqslant y_k \leqslant R_H\right)\right.\\ \left.+E\left(\sum_{\gamma_k\in\Gamma_k}\left[\beta_k\left(\frac{\partial\phi_k}{\partial q}(v_1-R_H)+\frac{\partial\psi_k}{\partial q}(R_H-y_k)\right)\right] \mid a \leqslant y_k \leqslant R_H\right)\right]\end{aligned} \tag{7.1}$$

Suppose that $\frac{\partial}{\partial q}S_k^U \geqslant 0$. Then

$$\begin{aligned}\frac{\partial}{\partial v_1}\left(\frac{\partial S_k^U}{\partial q}\right) = \Pr\Big(R_H \leqslant y_k \leqslant v_1\Big)E\left(\sum_{\gamma_k\in\Gamma_k}\left[\left(\beta_k\frac{\partial\phi_k}{\partial q}+\phi_k\frac{\partial\beta_k}{\partial q}\right)\right] \mid R_H \leqslant y_k \leqslant v_1\right)\\ +\Pr\Big(a \leqslant y_k \leqslant R_H\Big)\left[E\left(\sum_{\gamma_k\in\Gamma_k}\left[\frac{\partial\beta_k}{\partial q}\phi_k\right] \mid a \leqslant y_k \leqslant R_H\right)\right.\\ \left.+E\left(\sum_{\gamma_k\in\Gamma_k}\left[\beta_k\frac{\partial\phi_k}{\partial q}\right] \mid a \leqslant y_k \leqslant R_H\right)\right]\end{aligned} \tag{7.2}$$

Using Lemma 1 above, it follows that

$$\frac{\partial}{\partial v_1}\left(\frac{\partial S_k^U}{\partial q}\right) \leqslant 0. \tag{7.3}$$

Next, suppose $\frac{\partial S_k^U}{\partial q} < 0$. Then $\left|\frac{\partial S_k^U}{\partial q}\right| = -\frac{\partial S_k^U}{\partial q}$, and therefore, $\frac{\partial}{\partial v_1}\left|\frac{\partial S_k^U}{\partial q}\right| \geqslant 0$.

Now,

$$\frac{\partial \pi_1^U}{\partial q} = \sum_{k=1}^{n-1} \left(\frac{\partial S_k^U}{\partial q} \right) - E \left(\prod_{j=2}^{n} Q_j(v_j)(v_1 - R_L) \right)$$

Case 1. Suppose, first, that $\frac{\partial \pi_1^U}{\partial q} \geqslant 0$. Since the second term is clearly negative, it follows that in this case $\frac{\partial S_k^U}{\partial q} \geqslant 0$ (for all values of k, since the sign of this term cannot flip across values of k). In this case,

$$\frac{\partial}{\partial v_1} \left(\frac{\partial \pi_1^U}{\partial q} \right) = \sum_{k=1}^{n-1} \frac{\partial}{\partial v_1} \left(\frac{\partial S_k^U}{\partial q} \right) - E \left(\prod_{j=2}^{n} Q_j(v_j) \right) \leqslant 0 \quad (7.4)$$

where the last inequality follows using the inequality (7.3).

Case 2. Next, suppose $\frac{\partial \pi_1^U}{\partial q} < 0$. In this case, $\frac{\partial S_k^U}{\partial q}$ could be positive or negative. Now,

$$\frac{\partial}{\partial v_1} \left| \frac{\partial \pi_1^U}{\partial q} \right| = -\sum_{k=1}^{n-1} \frac{\partial}{\partial v_1} \left(\frac{\partial S_k^U}{\partial q} \right) + E \left(\prod_{j=2}^{n} Q_j(v_j) \right)$$

If $\frac{\partial S_k^U}{\partial q} \geqslant 0$, the right hand side is positive. Further, if $\frac{\partial S_k^U}{\partial q} < 0$, then $-\sum_{k=1}^{n-1} \frac{\partial}{\partial v_1} \left(\frac{\partial S_k^U}{\partial q} \right) = \sum_{k=1}^{n-1} \frac{\partial}{\partial v_1} \left| \frac{\partial S_k^U}{\partial q} \right| \geqslant 0$, and therefore the right hand side is again positive. Therefore, in this case $\frac{\partial}{\partial v_1} \left| \frac{\partial \pi_1^U}{\partial q} \right| \geqslant 0$, which completes the proof.‖

7.2 Proof of Lemma 3

Suppose all bidders other than 1 bid at $t = 0$, i.e. $Q_k(v_k) = 0$ for all $v_k \in [a, b]$ and all $k \in \{2, \ldots, n\}$. Suppose bidder 1 bids at $q > 0$. Let $g_1(t, n)$ denote the probability that the second bid arrives at time $t < q$ and let $g_2(t, q, n)$ denote the probability that the second bid arrives at time $t \geqslant q$.

Let $Z_1(t, q)$ denote the distribution of arrival time of the bid by bidder 1. Since bidder 1 submits the bid at $q \geqslant 0$, Z_1 has support $[q, 1]$. Let $\xi_1(t, q)$ denote the corresponding density function. Similarly, let $Z_k(t)$ denote the distribution of arrival time of the bid by bidders $k = \{2, \ldots, n\}$ which has support $[0, 1]$. Let $\xi_k(t)$, $k = \{2, \ldots, n\}$ denote the corresponding density function.

Note that $\xi_1(t,q) = \dfrac{1}{1-q}$ and $\xi_k(t) = 1$ for $k \neq 1$. Further, $Z_1(t,q) = \dfrac{t-q}{1-q}$ and $Z_k(t) = t$ for $k \neq 1$.

For any $t < q$, the probability that the second bid arrives at time t is given by

$$g_1(t,n) = \sum_{k=2}^{n} \left(\xi_k(t) \sum_{\substack{\ell=2 \\ \ell \neq k}}^{n} \left(Z_\ell(t) \prod_{\substack{j=2 \\ j\neq k \\ j \neq \ell}}^{n} (1 - Z_j(t)) \right) \right)$$

and for $t \geqslant q$, the probability that the second bid arrives at time t is given by

$$\begin{aligned} g_2(t,q,n) = Z_1(t,q) \sum_{k=2}^{n} \left(\xi_k(t) \prod_{\substack{j=2 \\ j\neq k}}^{n} (1 - Z_j(t)) \right) + \xi_1(t,q) \sum_{k=2}^{n} \left(Z_k(t) \prod_{\substack{j=2 \\ j\neq k}}^{n} (1 - Z_j(t)) \right) \\ + (1 - Z_1(t,q)) \sum_{k=2}^{n} \left(\xi_k(t) \sum_{\substack{\ell=2 \\ \ell \neq k}}^{n} \left(Z_\ell(t) \prod_{\substack{j=2 \\ j\neq k \\ j \neq \ell}}^{n} (1 - Z_j(t)) \right) \right) \end{aligned}$$

Using the values of Z_k and ξ_k for $k \in \{1, \ldots, n\}$, the expressions for g_1 and g_2 reduce to

$$\begin{aligned} g_1(t,n) &= (n-1)(n-2)t(1-t)^{n-3} \\ g_2(t,q,n) &= \frac{t-q}{1-q}(n-1)(1-t)^{n-2} + \frac{1}{1-q}(n-1)t(1-t)^{n-2} + \frac{1-t}{1-q}(n-1)(n-2)t(1-t)^{n-3} \end{aligned}$$

where g_2 simplifies further to

$$g_2(t,q,n) = \frac{(n-1)(1-t)^{n-2}}{1-q}(nt - q)$$

This completes the proof.‖

References

Bajari, Patrick and Ali Hortaçsu, "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *Rand Journal of Economics*, 2003, *34* (2), 329–355. 1, 5

Chakraborty, Indranil and Georgia Kosmopoulou, "Auctions with shill bidding," *Economic Theory*, 2004, *24*, 271–287. 7

Lamy, Laurent, "The Shill Bidding Effect versus the Linkage Principle," *Journal of Economic Theory*, 2009, *144*, 390–413. 7

Myerson, Roger, "Optimal Auctions," *Mathematics of Operations Research*, 1981, *6*, 58–63. 2, 8

Ockenfels, Axel 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*, 2006, *55*, 297–320. 1, 4, 6

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

Roth, Alvin E. and Axel Ockenfels, "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 2002, *92* (4), 1093–1103. 1, 5

The Sunday Times, "Revealed: how eBay sellers fix auctions," 2007. January 28. 2

Wintr, Ladislav, "Some Evidence on Late Bidding in eBay auctions," *Economic Inquiry*, 2008, *46* (3), 369–379. 1