

Shills and Snipes

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

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

May 27, 2017

Forthcoming: Games and Economic Behavior

Abstract

Online auctions with a fixed end-time often experience a sharp increase in bidding towards the end (“sniping”) despite using a proxy-bidding format. We provide a novel explanation of this phenomenon under private values. We show that it is closely related to shill bidding by the seller. Late-bidding by buyers arises not to snipe each other, but to snipe the shill bids. We allow the number of bidders in the auction to be random and model a continuous bid arrival process. We show the existence of late-bidding equilibrium. Next, we characterize all equilibria under a natural monotonicity condition and show that they all involve sniping with positive probability. We characterize the time at which such late bidding occurs and discuss welfare implications.

JEL CLASSIFICATION: D44

KEYWORDS: Online auctions, correlated 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 fixed end (“hard end”) time, and in such auctions there is often a noticeable spike in bidding activity right at the end, a phenomenon called “sniping” or “last minute bidding”.¹ In an English auction in which bidding is meant to be done incrementally, such behavior makes sense: by bidding just before the auction closes, a bidder might be able to foreclose further bids and win at a low price. However, to prevent such behavior, eBay allows bidders to use a proxy bidding system in which a bidder submits a maximum price, and the system then bids incrementally on behalf of the bidder up to the maximum price. The advantage of this system is that the proxy-bot cannot be sniped: so long as the highest bid of others is lower than the maximum price that a bidder has submitted to the proxy bid system, the latter wins.

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 hide the information content of their bids from other bidders.² However, a large fraction of auctions on online platforms such as eBay fit the private values paradigm well, and yet experience significant amount of sniping.³ What explains such bidder behavior in a private values setting? This is the question we address in this paper, and suggest a novel solution.

Our analysis starts by considering another phenomenon that occurs in online auctions.

¹It is the fixed ending that makes sniping possible. One way to submit a late bid is to use a sniping service. Several online sites offer this service, and have active user bases. See sites such as auctionsniper.com, gixen.com, ezsniper.com, bidsnapper.com. From site-provided lists of recent auctions won using its service, comments on the discussion forum, or user testimonials it is clear that there is an active market for sniping services.

²See Bajari and Hortacısu (2003), Ockenfels and Roth (2006).

³See, 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. 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. These items are fairly standardized products and would seem to fit the private values framework better. While the quality of, say, a laptop may indeed vary affecting the payoff of anyone who buys it in a similar fashion the crucial point is that it is unlikely that some bidders are better informed about the quality than others. With items such as coins, on the other hand, some bidders may have greater expertise than others in recognizing the true worth of the items. In such auctions, bidding behavior of experts may give away valuable information to the non-experts, prompting late bidding by the experts.

Sellers often put in bids assuming different identities (and/or by getting others to bid on their behalf). While the practice – known as “shilling” or “shill bidding” – is illegal, and frowned upon by the online auction community, prevention requires verification which is obviously problematic. Legal or not, shill bidding is reported to be widespread in online auctions.⁴

The principal characteristic of a shill bid—the one that presumably generates all the passion surrounding the issue—is that the seller submits bids above own value in order to raise the final price. In this sense, any non-trivial reserve price (i.e. 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 typically higher than the seller’s own value for the object. However, in a standard private-value auction with a known distribution of values, the optimal reserve price is also the optimal shill bid; there is no other higher bid that the seller can submit (openly or surreptitiously) that would improve revenue. Put differently, in a standard private values model there does not seem to be any rationale for shill bidding.⁵

In our model, a seller uses an online auction site (like eBay) to try to sell an item where the auction format used is proxy bidding. The important point of departure is that the seller faces some uncertainty about the distribution from which bidders’ values are drawn. In this setup bids convey information regarding the true distribution, creating an incentive for the seller to raise the reserve price. Since it is not possible to openly adjust the reserve price mid-auction, there is now scope for profitable shill bidding. And late bidding by bidders is directly related to shill bidding by the seller: the bidders bid late not because they want to snipe the bids of other bidders but because they want to *snipe the shill bids*.

The specific model we consider incorporates many of the features of real life online auctions. The set of (participating) bidders is random and their arrival at the auction is allowed to be random as well. A consequence is that neither the seller nor bidders observe the actual number of bidders, a feature that fits well with actual online auction environ-

⁴See, for example, the The Sunday Times (2007) report on shill bidding on eBay. See also the BBC Newsbeat report Whitworth (2010). In Walton (2006) the author describes how he and his colleagues placed a large number of shill bids on their eBay auctions.

⁵There might be scenarios – for example if cancelling bids is not costly – where the seller would have an incentive to shill bid even when the distribution is known. While this is not the focus here, it is worth pointing out that the bid-time choice problem of bidders in such scenarios is likely to be similar to that in our model.

ments. The auction proceeds in continuous time. Importantly, the bid arrival process is continuous and random. The auction has a fixed end time, and as bids get pushed later and later they start losing (smoothly) some chance of arrival. Note that in this game the actions of the players (submission of bids) are not directly observable - what is observable is a public signal (movement of the auction price) with a stochastic lag.

We show that there exists an equilibrium that exhibits sniping. In this equilibrium all bidder types delay their bids till the very last moment such that any further delay would result in their own bids arriving with a probability less than one. The seller submits shill bids whenever it is optimal to do so but the crucial point is that the shill bids fail to arrive with strictly positive probability.

Our second main result shows that in this environment sniping - in particular, the strategy bidders follow in the equilibrium mentioned above - is a general phenomenon. While in some equilibria there might be types who do not have any need to delay bids, it is always the case that there are types who gain from delaying the seller from submitting shill bids. However, it might be possible to sustain an equilibrium where some types bid early simply because the bidders themselves follow (somewhat strange) strategies that “punish” late bidding. We show, however, that under a natural “monotonicity” assumption such strategies can be ruled out in which case every equilibrium exhibits sniping with strictly positive probability.

Relating to the Literature

In our paper, bidders want to delay bids 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 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. all bids arrive, but no 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. Rasmusen (2006) models a private values setting in which a high value bidder hides information from a bidder who does not know own value by bidding at a discontinuous last minute.

Ockenfels and Roth (2006) consider a private values model and show that there is an equilibrium with last minute bidding. They assume a “last point” in time (let us call it t_L) such

that a bid made at t_L reaches with probability $0 < p < 1$, and importantly, 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 ; by doing so each takes a chance that his own bid will reach while the other bidder’s bid will not - allowing the former to win and pay a low price. Deviations are not profitable so long as the collusive price is low enough. Note, however, that if we drop the discontinuity in bid arrival and make the arrival probability of bids a continuous function of time (bid made at $t < t_L$ reaches with a higher probability than bids made at $t = t_L$ but the difference goes to zero as $t \rightarrow t_L$), then starting from the situation where bidders are supposed to be bidding at time t_L , each bidder would have an incentive to bid “a little early,” which then unravels the sniping equilibrium.

Ockenfels and Roth (2006) study a second model of last minute bidding with the same bid arrival timing structure but set in a common values environment with two bidders: an expert and a non-expert. Only the expert knows whether an item is genuine. They show an equilibrium in which the expert bids only if the item is genuine and bids only at the “last point of time” t_L to deny the non-expert any chance to react to this information.

In contrast to the above literature, we have a standard private values setting and bidders have no incentive to hide any information from other bidders; the reason for late bidding is to try to snipe the seller’s shill bids. A further difference is that we consider continuous bid times to study the optimal bid times.

Regarding shill bidding, Graham, Marshall and Richard (1990) investigate the question of phantom bids and model a fixed number of distributionally heterogeneous IPV bidders. In this case the auctioneer waits until bidding is over, observes the second highest value and updates the reserve price using a phantom bid. In our setting the incentive to shill bid arises from the fact that the value distribution is unknown to the seller. The seller, however, is not the auctioneer and the shill bids have to be placed in the same manner as the bids of the other (genuine) bidders. Also, we allow for a random number of bidders and a time dimension, so the specific updating mechanism is different.

Engelberg and Williams (2009) analyze an incremental shill-bidding strategy to discover the high value when bidders – presumably due to behavioral biases – bid in predictable units. Here too late bidding would be beneficial in reducing the scope for successful

shill bidding; however, such calculations need not apply when behavioral biases or naive decision-making dictate bid-time selection. In such contexts, our work can be seen as a benchmark model with rational bidders.

Chakraborty and Kosmopoulou (2004), Lamy (2009) examine shill bidding in environments with common 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. Kosmopoulou and De Silva (2007) provide experimental evidence of this phenomenon.

2 The Model

A seller is interested in selling a single unit of an indivisible object and uses an online auction site to try to sell the item. The seller's own value for the object is zero. The auction format is proxy bidding with a hard (i.e. fixed) end time. The seller can post a reserve price at the beginning and also submit shill bids during the auction.

Bidders⁶ are drawn randomly from some set of potential bidders and arrive randomly at the auction according to some stochastic process. The seller as well as each bidder therefore faces a random set of bidders that (possibly) changes over time. For $i = 0, 1, \dots, N$, let λ_i be the prior that the number of participating bidders is i , where $\sum_i^N \lambda_i = 1$. We assume $\lambda_i > 0$ for all i . The exact nature of the random arrival process is inessential to the subsequent analysis. We assume that the set of participating bidders as well as their arrival process are independent of the distribution of values as well as the actual values.

Values of bidders Let \mathcal{F} be a set of distributions F_1, \dots, F_H on the support $[\underline{v}, \bar{v}]$. We assume the following monotone likelihood ratio property of distributions.

Assumption 1 (Monotone likelihood ratio property) *The distributions in \mathcal{F} are ordered in terms of likelihood ratio property: a higher value of v is more likely to have been generated from a distribution $F_{k'}$ than from the distribution F_k for $k' > k$.*

⁶By bidders we mean genuine buyers. The seller assumes the identity of a buyer in order to submit shill bids but in what follows our use of the term bidder does not include the seller.

The assumption implies that the optimal reserve price is higher for distribution $F_{k'}$ than for F_k for $k' > k$. This provides motivation for the seller to shill bid: in so far as higher bids reflect higher values (the extent of which depends on the specific equilibrium), increase in the auction's current price (current second highest bid) results in updated posterior beliefs, which might, in turn, induce the seller to want to *raise* the reserve price through a shill bid.

Let $\mu_k > 0$, where $\sum_{k=1}^H \mu_k = 1$, be the probability with which nature chooses distribution F_k from the set \mathcal{F} . The bidders' values are then determined according to independent draws from the distribution F_k . Each bidder privately observes own value. Neither the bidders nor the seller observe F_k but have the same prior belief over \mathcal{F} given above.⁷




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

Timing of bids and arrivals The auction starts at $-T < 0$ and ends at time 1. A crucial element of our model is the continuous and stochastic bid arrival process. The arrival time of any bid submitted at time t is uniformly distributed on $[t, t + 1]$, so long as $t + 1 \leq 1$, i.e. $t \leq 0$. If $t > 0$, the bid gets lost⁸ with probability t , and with probability $(1 - t)$, the arrival time is now distributed uniformly over $[t, 1]$.

Note that bids submitted at time $t \in [-T, 0)$ arrive with certainty; such bids are “early bids.” Bids submitted at $t \in [0, 1]$ are “last minute” bids (we use the expressions “last

⁷We assume that bidders do not know the distribution only because we think it might be more realistic; none of our results are affected if we assume instead that bidders do know the distribution. What is crucial is that the seller does not know the distribution.

⁸Being “lost” simply means that the bid fails to arrive by the time the auction ends.

minute bidding”, “late-bidding” and “sniping” interchangeably).⁹ 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 $t > 0$ (inside the last minute) is lost with probability t .^{10,11}

Since we want to examine the optimal choice of time of bidding we assume that bidders arrive randomly over $[-T, 0]$. Therefore any bid placed at time $t > 0$ is due to strategic reasons (i.e., the bidder chose to delay submitting a bid) and not because it would not have been possible for the bidder to have bid earlier.

For any $t \in [-T, 1]$, let p_t denote the current auction price and h_t denote the public history of auction prices up to (but not including) time t . The public history h_t is thus a step function over the interval $[-T, t)$. When the first bid above the reserve price arrives, the reserve price becomes active and is shown as public history.¹² We define t as an **active period** if the auction price changes at t ; every instance that is not an active period is an inactive period.

Below, we describe strategies of the bidders and the seller somewhat informally to convey the essential ideas of our model without requiring the reader to wade through too much notation. Since we have a continuous time game there are the usual issues such as

⁹Note that given the continual improvement of technology and connection speeds, the “last minute” represented here by the unit interval should be thought of as representing a short period of time over which the bidder can choose to make a bid which might fail to arrive.

¹⁰Evidence of stochastic bid arrival abounds online. A Google search of the phrase “my eBay bid didn’t go through” brings up a large number of results including eBay community forum posts, where bidders complain about non-arrival of bids and replies suggesting they had bid too close to the end. Technology sites such as TechRadar advises bidders that sniping services might cut it too fine. It is also worth mentioning that some sites explicitly mention sniping the seller as motivation, which is the idea in this paper. The eBay buying guide site features an article titled “Sniping, The Intelligent Way to Bid!” which advises bidders to not show their hand early “or others, including the seller, will take advantage.”

¹¹There is another way to interpret this model of bid arrival. Suppose the eBay countdown clock is not necessarily synchronized with any computer’s clock so that a bidder who cuts it fine going by own computer clock might suddenly find the auction has ended before they could bid. Numerous online forums suggest such a possibility. See, for example, the eBay community site <http://community.ebay.com/t5/Archive-Bidding-Buying/Countdown-clock-is-wrong/qaq-p/6552937>. Suppose the clock is known not be off by more than 10 seconds. Within the last 10 seconds, the chance of a mismatch being present increases for any bidder as time passes. This would lead to exactly the same bid-arrival-timing structure as presented in the model.

¹²In some auctions, the first activity that is registered is when the second bid above the reserve price arrives. We assume the other variation as it is the more general one, but nothing in our analysis depends on whether the first activity occurs when the first or second bid above the reserve price arrives.

existence of well-defined strategies and dealing with “sequential” actions taking place at the “same instance.” These are dealt with in Appendix B.

2.1 Strategy of bidders

At every instant t , the feasible set of actions for a bidder (henceforth, bidders are assigned the pronoun “he”) is to either remain inactive or be active and submit a bid, which is a number in $[0, v]$, with two additional restrictions: later bids must exceed earlier bids and a bid at any time t must also be higher than the current auction price p_t .¹³

Bidder i arriving at time $\underline{t}_i \in [-T, 0]$ can submit one or more bids over time $t \in [\underline{t}_i, 0]$ and can also submit a bid at some point inside the “last minute,” i.e. at some time $q \in (0, 1]$.

Formally, we model a bidder’s actions as choosing his *bid level*. Let $b_{i,t}$ denote the bid level of bidder i at time t . We normalize the initial bid level to be 0. After that, for all instances where the bidder remains inactive, the bid level does not change whereas submitted bids are reflected by *upward jumps* in the bid level.

Bidder i can observe h_t for all $t \geq \underline{t}_i$. At every $t \geq \underline{t}_i$ bidder i also observes own value, arrival time, as well as the history of own bid levels for all time periods $\tau \in [\underline{t}_i, t)$. These, along with the public history h_t , form a bidder’s private history $h_{i,t}$ at t . **A strategy of bidder i** is a function that maps the set of all possible private histories of the bidder to the bid-level choice set, with the restriction that there can be no more than 1 upward jump in the bid level over $(0, 1]$. See Appendix B.1 for a more formal statement.

¹³We restrict the upper limit of bidding at the value v for the following reason. Note that any bid above true value is dominated by a bid of true value (the usual Vickrey-auction reasoning applies) - thus the restriction does not limit equilibrium behavior. However, since we are considering a dynamic game, without the restriction there would be (off-equilibrium) histories where a bidder has himself bid more than his value. The restriction avoids the problem of deciding optimal action after such “irrational” histories.

2.2 Shill bidding environment

The seller (“she”) starts with a reserve price of R_0 .¹⁴ The actual value of R_0 is inessential to our analysis as long as it is not so high that the seller does not have *any desire* for shill bidding in the future. Specifically, let r^1, \dots, r^H be the corresponding optimal reserve prices if the seller believed with certainty that the distribution was F_k for $k = 1, 2, \dots, H$. We assume that the seller’s prior over the set \mathcal{F} is such that if the seller were to choose a reserve price *under the assumption that it would be impossible for her to shill bid*, the chosen reserve price would be strictly less than r^H . It follows that in the actual model in which the seller does have the option to increase the reserve price later (with positive probability), the seller’s choice of R_0 cannot be equal to r^H under the same prior over \mathcal{F} . Importantly, this also means that the seller would have an incentive to shill bid.

We assume the following about bid timing: if the auction price moves at time t in a way so that the seller updates her reserve price, she can submit a shill bid at time t . This simultaneity of price movement and shill bid at t causes no interpretation problems (see Appendix B.1.4).¹⁵

Note that, similar to any bidder, the seller can submit any number of shill bids before time 0. Over the last minute any bidder can submit at most one bid. We do not place such a restriction on the seller - so that even here, the seller can submit one or more shill bids. Our results remain unchanged even if the seller could only bid exactly once over the last minute - but we do not *require* this in order to show that our results are not driven by any such restriction.¹⁶

The starting reserve price R_0 is part of the stated mechanism. The more interesting aspect of the seller’s strategy is the submission of shill bids. Since shill bidding is illegal, we assume that the seller shill bids through multiple accounts to avoid detection.¹⁷

¹⁴Choice of R_0 follows from standard dynamic programming principles. For any given R_0 , the seller calculates the expected revenue using her priors regarding the buyer-value distributions F_k , buyer arrival process, and her knowledge of the strategies - including her own shill bidding strategies - and then chooses the R_0 that maximizes this expected revenue. We assume such a maximum exists. In particular, since the strategies may depend on R_0 , optimal choice of R_0 may involve solving for a fixed point.

¹⁵Note that this is a statement about timing of *bid submission* by the shill bidder. Once the seller submits a bid, arrival of the shill bid is of course stochastic and is according to the same process specified previously.

¹⁶Also, in reality a seller might indeed ask other agents to bid on her behalf - and so restricting the ability to put multiple bids over the last minute may be unrealistic.

¹⁷In practice, she could also ask others—agents carrying out her instructions—to bid on her behalf. Since

2.3 Strategy of the seller

The seller chooses the initial reserve price R_0 equal to the optimal reserve price in a second price auction given the seller's prior over the set of distributions \mathcal{F} . As the auction progresses, the seller acts as an “updater”: whenever the auction price moves resulting in the seller receiving some information, the seller uses it to update the reserve price and decides whether to submit a shill bid. Thus the seller becomes active only at active periods. The seller remains inactive at all inactive periods.

It follows that at every instant $t \in [-T, 1]$, the feasible set of actions for the seller is to either remain inactive or be active and submit a shill bid, which is a number in $[R_0, v_H]$, with the additional restrictions that the seller is active at t only if auction price moves at t , later bids must exceed earlier bids and a bid at any time t must also be higher than the current auction price p_t .

Let $b_{s,t}$ denote the shill-bid level of the seller at time t . We normalize the initial bid level of the seller to be R_0 . After that, for all instances where the seller remains inactive, the shill-bid level does not change, whereas submitted shill bids are reflected by *upward jumps* in the shill-bid level.

The seller can observe public history h_t for all $t \in [-T, 1]$. At every t , the seller also observes the history of own bid levels for all time periods $\tau \in [-T, t)$. These, along with the public history h_t form the seller's private history, $h_{s,t}$ at t . **A strategy of the seller** is a function that maps the set of all possible private histories of the seller to the shill-bid level choice set. See Appendix B.1 for a more formal statement.

Further, the seller's bid level choice in active periods is defined recursively. Let τ_1 be the first instance of time $t > -T$ that is an active period. If the auction price at τ_1 is p , the seller updates her prior (over the set of distributions \mathcal{F}) given the information content of the auction price and calculates an updated optimal reserve price using the updated prior.¹⁸ Let the updated reserve be denoted by r_{τ_1} . If $r_{\tau_1} > R_0$, then the seller submits

modelling these agents and the communication between them and the seller is unnecessary for our purposes, we simply assume that the seller creates multiple accounts herself and shill bids using them. While we assume that the seller manages to avoid detection by the auction platforms, none of our results would be qualitatively affected if we assumed instead that the seller's illegal activities are detected with positive probability provided the punishment for shill bidding, upon being detected, is not too large.

¹⁸See Appendix B.2 for details of updating for the information arriving given the bidders' strategies specified in Theorem 1.

a shill bid equal to r_{τ_1} ; otherwise she remains inactive. Define $S_1 = \max[R_0, r_{\tau_1}]$. The seller's strategy is defined recursively. Consider any active period $t = \tau_k$. The seller again updates her (posterior) beliefs over \mathcal{F} (given public and her private histories) and calculates the updated reserve price r_{τ_k} . If $r_{\tau_k} > S_{k-1}$, seller submits a shill bid equal to r_{τ_k} ; otherwise she remains inactive. The "counter" S is updated to $S_k = \max[S_{k-1}, r_{\tau_k}]$.

3 Late-bidding Equilibrium

In this section we show existence by constructing an equilibrium with late-bidding and sniping. In this equilibrium bidders delay submitting their bids so as to reduce the chance of a successful shill bid but do not delay so much that they incur the risk that their own bids may not arrive. In the next section, we show that this is not an isolated special case. Under a reasonable restriction on strategies, sniping is a pervasive phenomenon: *every* equilibrium involves late-bidding with strictly positive probability.

Define bidding **truthfully** as a bidder submitting a bid equal to own value. Recall that R_0 is the seller's initially chosen reserve price (the official reserve price of the auction).

Theorem 1 *There exists an equilibrium in which every arriving bidder of any type $v \in (R_0, \bar{v}]$ bids once, and truthfully, at time 0.*

The rest of the section constructs the proof of this result. First, we set up strategies for the bidders and the seller. Next, we show that these are mutual best responses.

Proof: Consider the following strategies for bidders and the seller. We show that these form an equilibrium.

Bidders' strategies: All bidders of all types remain inactive – i.e., do not submit any bids – for all histories for $t \in [-T, 0)$. At $t = 0$, and for any history, bidder with value v submits a bid equal to v if the auction price at time 0 is less than v ; otherwise the bidder remains inactive. For any $t \in (0, 1]$, a bidder remains inactive if the history of the bidder is such that the bidder has submitted bid equal to own value v at time 0. For any history such that the bidder has not submitted bid equal to v at time 0, the bidder immediately submits a bid equal to v if the current auction price is strictly less than v , and remains inactive if the auction price is (weakly) greater than v .

Seller's strategy: The seller's strategy is as stated in section 2.3. In addition, we need to specify the seller's beliefs at any active period for our proposed equilibrium. Following the notation introduced in section 2.3, consider any active period $t = \tau_k$. If the auction price at τ_k is p , the seller updates her current prior (over the set of distributions \mathcal{F}) using the belief that a bidder with type $v \geq p$ has arrived. As stated previously, the seller then uses her updated beliefs over \mathcal{F} and calculates the updated reserve price r_{τ_k} . See Appendix B.2 for details of updating for the information arriving over the course of the auction given the strategies of bidders specified above.

It is clear that if the above profile of strategies is an equilibrium, the resulting outcome would be as stated in the result: all arriving bidders would bid for the first time at time $t = 0$ and bid truthfully and not make any further bids. Hence, the remaining task is to check that the above is indeed an equilibrium.

Consider first the bidders.

The bidders' problem

To show that bidding at time 0 is an equilibrium, we need to rule out possible deviations. There are three types of possible deviations: bid (at or below true value) before time 0, bid lower than true value at time 0 then raise the bid at some point after 0 (incremental bidding), bid only after time 0. We rule out these in the following three steps.

- **Step 1** (deviation to bidding before time 0): It is obvious that deviating and submitting a bid lower than true value at time $t = 0$ or $t < 0$ is worse than submitting a bid equal to true value at $t = 0$. Since bids submitted at $t < 0$ and $t = 0$ both reach with certainty, the usual weak dominance argument applies. Further, submitting a bid equal to true value at some time $t < 0$ is not a profitable deviation. Given bidder strategies, the early bid does not change the behaviour of any other bidder. Bids submitted at $t < 0$ and $t = 0$ both reach with certainty but the earlier bid triggers a shill bid from the seller earlier with (at least weakly) higher probability, reducing expected payoff.

Steps 2 and 3 are completed using two results, the proofs of which are relegated to the Appendix.

- **Step 2** (deviation to incremental bidding): We now need to rule out the possibility that

a bidder deviates by bidding some number less than his true value (i.e., bids untruthfully) at $t = 0$ and then bidding some higher number less than or equal to true value again at some $t > 0$. Proposition 1 below shows that such incremental bidding is unprofitable: the bidder should bid truthfully either at $t = 0$ or at some $t > 0$.

Proposition 1 *Given the strategy profile specified above, it is optimal for any bidder to submit a single bid of true value v either at time 0 or at some point of time $q \in (0, 1)$. In other words, incremental bidding is suboptimal.*

The formal proof is in the Appendix. The intuition is quite simple. Consider an incremental bidding strategy in which a bidder with value v bids $v_1 < v$ at $t = 0$ and bids again $v_2 \leq v$ at some point in time $q \in (0, 1)$. If v_1 is a winning bid, adding a bid later can only reduce expected payoff. This is because the bid of v_1 arrives with certainty - so the second bid adds nothing to arrival probability. However, with strictly positive probability the second bid arrives before the first bid, and when it does, with strictly positive probability it triggers a shill bid. But this shill bid is triggered earlier than necessary (i.e. earlier than the time at which v_1 arrives), thus raising the probability that the shill bid actually arrives, which in turn reduces expected payoff. The second point to note is that - and this follows from the standard property of second price auctions - raising the winning bid does not change the auction price and so payoff from v_1 given that it is a winning bid is the same as the payoff from v . Thus if it is optimal not to sacrifice any probability of bid reaching, it is best to bid v at 0 and nothing further. If, on the other hand, it is optimal to sacrifice some probability of winning, it is best to bid v at some $q > 0$. In this case adding a bid of v_1 at 0 reduces payoff, as, with strictly positive probability, it arrives earlier than the arrival time of the bid at q and triggers a shill bid.

- **Step 3** (deviation to bidding only at $t > 0$): Proposition 2 below shows that remaining inactive at $t = 0$ and bidding at $t > 0$ is not a profitable deviation.

Proposition 2 *Given the strategy profile specified above, deviating to bidding at some $t > 0$ is not profitable for any bidder.*

The result shows that delaying bidding beyond 0 sacrifices some chance of arrival but gains nothing. The intuition follows from two crucial observations. First, bid submission is unobservable and the resulting payoff depends only on the time at which the bid ar-

rives, not on when the bid was submitted.¹⁹ Second, for any arrival-time s , the payoff from a bid arriving at s cannot be negative (since the bidder bids at most true value) and is in fact strictly positive since there is a strictly positive chance of winning. Thus instead of bidding at time 0 if a bidder deviates and delays bidding till $t > 0$, the payoffs are identical for any arrival time $s > t$, but bidding at 0 allows additional opportunities for arrival $s < t$ and consequent positive payoffs $\pi(s)$ that are lost if the bid is delayed to t , making the deviation unprofitable.

Step 1 and Proposition 1 in step 2 narrow the optimal bidding strategy of a bidder of type v to two options: bid v once at $t = 0$ or bid v once at some time $t > 0$. This still leave open the possibility that starting from the stated strategies, deviating to $t > 0$ could be profitable. Proposition 2 in step 3 then rules this out by showing that remaining inactive at $t = 0$ and bidding at $t > 0$ is not a profitable deviation. Thus bidders cannot profitably deviate from the stated strategy profile.

The seller's problem

To complete the proof we need to show that the seller's strategy is optimal. Note that any active period $t < 0$ is clearly off-equilibrium-path and Perfect Bayesian Equilibrium puts no restriction on beliefs and resulting action by the seller at those periods. Therefore the beliefs specified are compatible with equilibrium.²⁰ For any active period $t > 0$ the seller's posterior beliefs are consistent with bidders' strategies and the assumption that the set of bidders who arrive at the auction is independent of the distribution of values. So, the last thing to check is whether the seller with updated reserve r_{τ_k} benefits from refraining from shill bidding r_{τ_k} even if $r_{\tau_k} > S_{k-1}$. Note however that such action is beneficial only if submitting the shill bid at $t = \tau_k$ will prevent the seller from taking some profitable action in the future. However, that is not possible. Ability to successfully shill bid in the future depends on time remaining in the auction, not on past shill bids. Finally, any increase in auction price (weakly) increases the updated reserve and hence there cannot be a future event at which the seller regrets a past shill bid and would like

¹⁹The overall expected payoff depends of course on the timing of the bid since that affects the probabilities of the bid reaching at various points in time. The point however is that the payoff $\pi(s)$ resulting from the bid reaching at time s does not depend on when the bid was made.

²⁰Given our assumption that the set of participating bidders, as well as the arrival process of bidders are independent of (distribution of) values, seller's posterior beliefs as postulated seem natural.

to lower it.||

4 Sniping across all equilibria

Our second main result argues that in the environment we study, sniping is a “general” phenomenon in the sense that all equilibria involve sniping with positive probability.

Part of the intuition behind this result is obvious. While it is possible that there are bidder types (values) who do not have any strict incentive to delay submitting their bids (since it is possible that the seller does not submit shill bids when the auction price moves within some certain ranges), it is not the case that all types will have such indifference. Therefore there should be types who, like in the equilibrium above, would like to submit bids at time $t = 0$ but not earlier. The question then is what can happen to make also these types bid early in equilibrium.

This can happen only if the bidders themselves follow (somewhat strange) strategies that “punish” delayed bidding. We discuss this possibility in section 4.1 below and show that a natural restriction on strategies (“monotonicity”) rules out the possibility of any equilibrium involving such self-punishing strategies. Armed with this restriction, we show that types who could trigger shill bids never bid before time 0 (Proposition 3). We then use slight variations of Propositions 1 and 2 (namely, Propositions 1A and 2A) to show that no type would want to delay bidding to some $t > 0$. Our next main result, Theorem 2 then follows.

4.1 Monotonicity

As noted above, the reason we might not have last minute bidding is if bidders themselves use strategies in which they “punish” delayed bidding (which of course acts against their own collective self-interest). To see how this might be possible, fix a time $t_1 < 0$ and a price level $p_1 > R_0$. Consider the following strategy of a type v : If by time t_1 the auction price has not reached p_1 , bid v at t_1 . If the auction price has reached p_1 , do nothing at t_1 (i.e. wait further to bid). If higher types follow such strategies, it might be optimal for some types to bid early (to avoid being “punished” by other types who would bid early and would also, in turn, trigger higher shill bids early). The strategies of higher types

might be optimal if they, in turn, faced such punishments from even higher types.²¹ Such strategies might not be part of any equilibrium, but in our general setting, we cannot rule them out without imposing some restrictions on strategies.

The strategies mentioned above are non-monotonic in the sense that at a certain time, lower prices would trigger bids but higher prices would not. As we show below, if we impose a monotonicity requirement on strategies, all equilibria involve late bidding.

Definition 1 (Monotonicity) *A strategy of a bidder of type v is monotonic if the following property holds: if the bidder submits a bid of $v' \leq v$ at time t if the auction price at t is $p < v'$, then the bidder also submits a bid of at least v' at t for any higher auction price $p' \in (p, v']$.*

Implication of monotonicity Monotonicity has the following implication that is used repeatedly to prove the results that follow. Suppose in some proposed equilibrium, a bidder (say bidder 1) is supposed to submit a bid at time t , but deviates and submits the bid at $t' > t$. The deviation weakly reduces the chance of the auction price crossing any given threshold, which in turn weakly delays the next bid being triggered, which again weakly reduces the chance of the auction price crossing any threshold and so on. Therefore, imposing monotonicity implies that the deviation, involving delaying the bid, cannot strictly increase the chance of a bid (by some other bidder) being triggered at any future point of time.

4.2 Last minute bidding across all equilibria

To see whether last minute bids must occur, we first consider the incentive to bid before time 0. As noted at the start of this section, it is possible that there are bidder types who do not have any strict incentive to delay submitting their bids in equilibrium, since it is possible that the seller does not submit shill bids when the auction price moves within certain ranges. However, our model ensures that there are types with values high enough so that when price movements lead the seller to update upwards the probability that such types have arrived, shill bids would be triggered. We call these types “shill-positive,” defined below.

²¹In other words, any early bidding equilibrium necessarily involves threats to each bidder from others saying in effect “bid early or face a higher chance that we will bid earlier than otherwise and facilitate shill bids.”

Definition 2 (Shill-positive types) *Given any equilibrium, a type v is said to be shill-positive in that equilibrium if the following is true. If at any time $t \in (-T, 1)$ during the auction the seller believes that a type at least equal to v has arrived, she revises her beliefs over \mathcal{F} such that she submits a shill bid at t .*

It follows directly from the monotone likelihood ratio property (assumption 1) that if v is shill-positive, so is any type $v' > v$. Therefore, the set of shill-positive types in an equilibrium is a non-empty interval of the form $[v_*, v_H]$. The result below shows that such types never bid early.

Proposition 3 *Suppose bidders other than 1 use monotonic strategies. In any equilibrium, for bidder 1 of any “shill-positive” type $v > R_0$, bidding at $t < 0$ is suboptimal.*

The intuition is straightforward: bids at time 0 as well as times $t < 0$ reach with certainty. Staring from any proposed equilibrium in which a shill-positive type bids before time 0, consider a deviation to time 0. From the implication of monotonicity noted above, this does not trigger any punishment from other bidders. Further, this does not sacrifice of arrival probability, but later arrival triggers shill bids later, reducing the probability that the shill bid arrives.

Since any arriving bidder draws a shill-positive type with strictly positive probability, the following corollary is immediate.

Corollary 1 *Last minute bids occur with strictly positive probability in all equilibria.*

4.3 Equilibrium strategies of bidders: further characterization

Let us now show that if others use monotonic strategies, the best response of a bidder with a shill-positive type cannot involve incremental bidding. Proposition 3 rules out bidding before time 0 by such types. However, this still leaves open the possibility that such a type submits a bid at time 0 and another inside the “last minute.”

Proposition 1 rules out this possibility for the strategies we constructed to show existence of late-bidding (Theorem 1). The reason is that the bidder strategies constructed for Theorem 1 are monotonic. As the next result shows, the same result in fact applies to all

monotonic strategies. The proof is also essentially the same - there is exactly one step where the argument needs to be modified slightly to take monotonicity into account - we do this in the Appendix.

Proposition 1A *Suppose bidders other than 1 use monotonic strategies. In any equilibrium, for bidder 1 of skill-positive type $v > R_0$, it is optimal to submit a single bid of v either at time 0 or at some point of time $q \in (0, 1)$. In other words, incremental bidding is suboptimal, and in any equilibrium a bidder bids exactly once, and submits a truthful bid, at some point of time in $[0, 1)$.*

The intuition is the same as that for Proposition 1. Next, a result similar to Proposition 2 rules out waiting beyond time 0.

Proposition 2A *There is no equilibrium in which any type of any bidder bids after 0.*

The intuition is the same as that for Proposition 2. For the sake of completeness, we include a proof in the Appendix. We now prove the main result of this section.

Theorem 2 *If we restrict attention to monotonic strategies, in all equilibria, all bidders of all skill-positive types above R_0 bid their true value exactly at time 0.*

Proof: Proposition 3 rules out bid times before time 0. From Proposition 1A, we know that bidders submit truthful bids once either at 0 or at some point of time in $(0, 1)$. Finally, Proposition 2A rules out the latter. Therefore in all equilibria, all bidders of all skill-positive types above R_0 bid their true value exactly at time 0. ||

Note that the above result characterizes the bid time for types that are skill-positive; for those that are not, they can bid at any time in $[-T, 0]$. Thus there could be equilibria involving both early and late bids as well as ones with only late bids (e.g. if all types above R_0 are skill-positive). The fact that skill-positive types bid exactly at the last possible instant of time at which their bids would arrive with certainty is also reminiscent of the widespread use of sniping services. A good quality sniping service behaves exactly in this way: cuts it as fine as possible but ensures bid arrives.

5 Conclusion

Late-bidding is a widely observed phenomenon in online auctions, many of which fit the private values model well. We provide an explanation for such bidding behavior connecting it to another commonly observed phenomenon in online auctions - shill bidding. The main result is that the bidders bid late not to snipe each other but to snipe the shill bids. Our model incorporates many of the features of real life online auctions and the framework we develop, we believe, should be useful for analyzing online auctions under common values or other richer valuation environments.

We conclude by noting possible welfare implications of our results. As mentioned in the introduction, shill bidding is illegal and universally frowned on. The literature on shill bidding in common value auctions justify such status by showing how shill bids might impede information aggregation (and lead to a reduction in the seller's equilibrium payoff). However, with private values the conclusions are less clear cut. Obviously a successful shill bid is an increase in reserve price and can therefore result in additional loss of efficiency for the usual reasons. However, an option of shill bidding might lead the seller to post a lower initial reserve price as compared to the case where such an option is not available.²² Given that in equilibrium bidders delay submitting their bids resulting in some shill bids failing to reach, it is possible that a lower initial reserve price combined with shill bids sometimes failing to arrive could result in overall higher welfare than when the seller starts with a higher reserve price knowing that she will not have a chance to update it later. However, a conclusive answer to whether shill bidding is overall harmful depends on specifics of any actual situation (e.g. the precise nature of value distributions, bidder arrival process and other aspects of the model).

Finally, in addition to considering overall welfare, it might also be interesting to examine this at a disaggregated level and see how the welfare of buyers, sellers and online auction sites are affected. An interesting question for future research would be to study shill bidding and bid-time choice in a framework of multiple online auction sites such that buyers and sellers can choose which auction sites to use, with sites choosing different rules that affect the bidding possibilities for buyers and sellers as instruments for competing for customers.

²²The maintained assumption in this paper is that it is not possible to stop shill bidding. The discussion here is under the counterfactual of what would happen if stopping shill bidding were possible.

A Proofs

A.1 Proof of Proposition 1

Consider the problem of bidder 1 of type v . Consider an incremental bidding strategy v_1 at time 0 and v_2 at time $q \in (0, 1)$, where $v_1 < v_2 \leq v$.

In what follows, the term “positive probability” means probability strictly greater than zero.

Step 1: Let $P_0(v_1, v_2)$ be the expected payoff given that v_1 is a winning bid and given that v_1 reaches.

Since v_1 reaches with certainty, the bid of v_2 serves in this case only to trigger a shill bid earlier than necessary with positive probability. To see this, suppose v_1 arrives at $t > q$ (an event that occurs with positive probability). In the absence of the bid of v_2 at q , a shill bid would be triggered by bidder 1's bid only at t . However, if the bid of v_2 arrives before t (which happens with positive probability), a shill bid is triggered earlier than necessary (note that an earlier shill bid has a greater chance of reaching, thereby reducing the expected payoff of bidder 1).

Further, dropping the bid at q must also weakly delay the shill bids triggered by arrival of bids by other bidders. It follows that the payoff given v_1 is a winning bid can be improved by dropping the bid at q :

$$P_0(v_1, v_2) < P_0(v_1). \quad (\text{A.1})$$

Here $P_0(v_1)$ is the expected payoff of bidder 1 given v_1 , submitted at time 0, wins (and there are no other bids by bidder 1).

Next, note that if v_1 is a winning bid, so is any bid greater than v_1 . In particular, v is a winning bid. Further, if we raise v_1 , the payoff given v_1 wins does not change. This is a standard property of second price auctions - raising the winning bid does not change auction price (in other words, any higher bid is observationally equivalent: it has the same impact on auction price and triggers shill bids in exactly the same way). So the payoff given v_1 wins ($P_0(v_1)$) is the same as the payoff given v wins, i.e.

$$P_0(v_1) = P_0(v). \quad (\text{A.2})$$

Step 2: Next, let $EP_q(v_1, v_2)$ denote the expected payoff when v_1 is a losing bid and v_2 reaches. Note that v_2 gets lost with positive probability - so the bidder's expected payoff is a product of the actual expected payoff given v_2 wins and v_1 loses, and the probability of arrival of the bid of v_2 .

Now, if $v_2 < v$, $EP_q(v_1, v_2)$ can be raised by setting $v_2 = v$ since when v_2 wins, whether $v_2 < v$ or $v_2 = v$ makes no difference to payoff, but in cases where v_2 loses and v wins, the payoff gets raised. It follows that

$$EP_q(v_1, v_2) \leq EP_q(v_1, v). \quad (\text{A.3})$$

Step 3: The expected payoff from the incremental bidding strategy is given by:

$$\pi_{\text{inc}} = \Pr(v_1 \text{ wins})P_0(v_1, v_2) + \Pr(v_1 \text{ loses}, v_2 \text{ wins})EP_q(v_1, v_2).$$

Using the inequalities (A.1) to (A.3),

$$\pi_{\text{inc}} < \bar{\pi} = \Pr(v_1 \text{ wins})P_0(v) + \Pr(v_1 \text{ loses}, v \text{ wins})EP_q(v_1, v).$$

Next, let

$$\alpha(v_1) = \frac{\Pr(v_1 \text{ wins})}{\Pr(v \text{ wins})}.$$

Note that as $v_1 \rightarrow v$, $\alpha \rightarrow 1$. Also, if v_1 is dropped (or, equivalently, v_1 is set to a value lower than the initial reserve price R_0), $\alpha = 0$. For the purpose of the proof, let us represent dropping v_1 as setting $v_1 = 0$. Then $v_1 \in \{0\} \cup [R_0, v]$. Further, if v_1 is raised to v , $\alpha = 1$. Then we can write

$$\bar{\pi} = \Pr(v \text{ wins}) \left[\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \right] \quad (\text{A.4})$$

If we can show that the convex combination inside the square brackets is maximized either if $v_1 = 0$ or $v_1 = v$, that would prove that incremental bidding is suboptimal and the optimal strategy is either to bid only at 0 or to bid only at some point $q > 0$. Further, since bidding exactly once is optimal, the optimal bid is indeed v . This would complete the proof.

We now show that the convex combination inside the square brackets is indeed maximized at a corner.

Consider the term inside the square brackets. As v_1 decreases, α decreases. Further, the payoff given v_1 is a losing bid ($EP_q(v_1, v)$) increases. This is because the bid of v_1 , which reaches with certainty, reaches before the winning bid with positive probability, and therefore triggers shill bids earlier than necessary (i.e. earlier than shill bids triggered by arrival of the winning bid) with positive probability. A lower v_1 still reaches before the winning bid with the same probability, but upon reaching triggers a shill bid with lower probability, thereby raising payoff.

It follows that the payoff given v_1 is a losing bid and v arrives is maximized when v_1 is dropped altogether:

$$EP_q(v_1, v) < EP_q(v) \quad (\text{A.5})$$

Consider the convex combination $\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v)$. As specified above, $v_1 \in \{0\} \cup [R_0, v]$, where $v_1 = 0$ represents dropping the bid of v_1 . As v_1 varies, $P_0(v)$ does not change while $EP_q(v_1, v)$ is strictly decreasing in v_1 .

The two corner values of the convex combination are $P_0(v)$ and $EP_q(v)$, corresponding to $v_1 = v$ and $v_1 = 0$ respectively. Now suppose there is an interior value of v_1 (i.e. some $v_1 \in [R_0, v)$) that maximizes the convex combination. For that interior value, we have $\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \geq \max[P_0(v), EP_q(v)]$, or writing it out, we have,

$$\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \geq P_0(v), \quad (\text{A.6})$$

$$\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \geq EP_q(v). \quad (\text{A.7})$$

Since $1 - \alpha(v_1) \neq 0$, we have from (A.6) that

$$P_0(v) \leq EP_q(v_1, v). \quad (\text{A.8})$$

From (A.7), we have

$$\alpha(v_1)(P_0(v) - EP_q(v_1, v)) \geq EP_q(v) - EP_q(v_1, v) > 0,$$

where the last inequality follows from (A.5). Since v_1 is an interior point, we have $\alpha(v_1) > 0$. Therefore, the above implies $P_0(v) > EP_q(v_1, v)$, which contradicts (A.8).

It follows that the expression $\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v)$ cannot be maximized at an interior value of v_1 - it is maximized either by bidding only at 0 or by bidding only at some point $q > 0$. Therefore incremental bidding is suboptimal, and the only relevant question is whether the bidder should bid at 0 or wait until some point $q > 0$ to bid.

Since bidding exactly once is optimal, it also follows from standard second-price auction logic that truthful bidding is optimal (bidding anything other than true value is weakly dominated). This completes the proof. ||

A.2 Proof of Proposition 2

Suppose some type v of bidder 1 (say) deviates to bid at time $t > 0$ under some history. Note that the expected payoff of the bidder conditional on the bid arriving at time $s > t$ depends only on s . (The time of bid submission is not part of public history, so the continuation histories conditional on the bid arriving at s are exactly the same whether the bid is submitted at t or some other t' .) Let $\pi(s)$ be the expected payoff conditional on the bid reaching at s . Also, recall from section 2 that a bid at $t > 0$ arrives with probability $1 - t$. The expected payoff after deviation can then be written as

$$P(t) = (1 - t) \int_t^1 \pi(s) \frac{1}{1 - t} ds = \int_t^1 \pi(s) ds.$$

The payoff gain over bidding at $t = 0$ is

$$P(t) - P(0) = - \int_0^t \pi(s) ds.$$

Consider the sign of the expected payoff $\pi(s)$ for any arrival time $s \in [0, t)$. Since bidder 1 bids at most own value v , $\pi(s)$ cannot be negative for any value of s . Further, with strictly positive probability all other arriving bidders draw values below v , and even in the worst case in which the seller submits a shill bid strictly greater than v , it fails to reach with strictly positive probability. Thus $\pi(s)$ is always nonnegative, and is strictly positive with strictly positive probability. It follows that for any arrival time $s \in [0, t)$, $\pi(s) > 0$. Therefore, $P(t) - P(0) < 0$, and the deviation lowers payoff. This completes the proof. ||

A.3 Proof of Proposition 3

The seller's strategy involves submitting a shill bid at active times (i.e. times when the auction reserve price is met or auction price jumps). It follows that if the bid-arrival-time distribution shifts to the right, the distribution of shill bidding times also shifts to the right, which strictly reduces the probability of successful arrival of shill bids and strictly improves the expected surplus of any serious bidder.

Now suppose there is an equilibrium in which some types of shill-positive bidders submit bids before time 0. Then it must be that in equilibrium shill bids are triggered before time 0 with strictly positive probability.

Suitably rename bidders so that bidder 1 of type $v > R_0$ is shill-positive, and submits a serious bid of $v' \leq v$ (a bid that exceeds the current auction price if the reserve price has already been met, or exceeds the reserve price if the auction is not yet active) at time $t < 0$. Consider a deviation by bidder 1 in which the bid of v' is submitted at time 0. In both cases the bid arrives with certainty before the end of the auction. So the deviation does not lose any probability of arrival. Further, as noted in section 4.1, the assumption of monotonicity implies that, starting from an equilibrium strategy, if a bidder deviates and bids later, this cannot increase the chance of a bid by some other bidder type being triggered at any future point of time. It follows that shifting bidder 1's bidding time to 0 shifts the bid-arrival-time distribution to the right, which raises bidder 1's expected surplus. Therefore the deviation is profitable, which is a contradiction. ||

A.4 Proof of Proposition 1A

Consider the problem of bidder 1 of type v . Consider an incremental bidding strategy v_1 at time 0 and v_2 at time $q \in (0, 1)$, where $v_1 < v_2 \leq v$.

In what follows, the term “positive probability” means probability strictly greater than zero.

Step 1: Let $P_0(v_1, v_2)$ be the expected payoff given that v_1 is a winning bid and given that v_1 reaches.

As shown in the proof of Proposition 1, since v_1 reaches with certainty, the bid of v_2 serves in this case only to trigger a shill bid earlier than necessary with positive probability.

Further, monotonicity of strategies of bidders other than 1 implies that if a bid from 1 arrives later, this cannot increase the chance of a bid (by some other bidder) being triggered at any future point of time. Therefore dropping the bid at q must also weakly delay the shill bids triggered by arrival of bids by other bidders. It follows that the payoff given v_1 is a winning bid can be improved by dropping the bid at q :

$$P_0(v_1, v_2) < P_0(v_1). \quad (\text{A.9})$$

Here $P_0(v_1)$ is the expected payoff of bidder 1 given v_1 , submitted at time 0, wins (and there are no other bids by bidder 1).

The rest of the proof is identical to that of Proposition 1. ||

A.5 Proof of Proposition 2A

Suppose there is an equilibrium in which some type v of some bidder is supposed to bid at time $t > 0$ under some history. Note that the expected payoff of the bidder conditional on the bid arriving at time $s > t$ depends only on s . The time of bid submission is not part of public history, so the continuation histories conditional on the bid arriving at s are exactly the same whether the bid is submitted at t or some other t' .²³ Let $\pi(s)$ be the expected payoff conditional on the bid reaching at s .²⁴ Also, recall from section 2 that a bid at $t > 0$ arrives with probability $G_t(1) < 1$. The expected payoff in the purported equilibrium can then be written as

$$P(t) = (1 - t) \int_t^1 \pi(s) \frac{1}{1 - t} ds = \int_t^1 \pi(s) ds.$$

Now consider a deviation to bidding at an earlier time $t - \Delta > 0$. As noted above, given any arrival time $s > t$, the payoff is the same as before. Therefore

$$P(t - \Delta) - P(t) = \int_{t-\Delta}^t \pi(s) ds$$

Now consider the expected payoff $\pi(s)$ for any arrival time $s \in [t - \Delta, t)$. The worst case for bidder 1 is when such a deviation is detectable with certainty.²⁵ In that case, the worst possible punishment is that other bidders all bid their values at s .²⁶ Since bidder 1 bids at most own value v , the expected payoff $\pi(s)$ cannot be negative for any value of

²³Bidding at t or t' would result in different probabilities that the bid reaches at any given $s > \max\{t, t'\}$ but conditional on the bid reaching at s , the expected payoff of the bidder would be exactly the same.

²⁴The realized payoff depends on the history at time s and the subsequent future histories following the history at time s . The expected payoff $\pi(s)$ is an expectation of the realized payoffs taken with respect to all these histories.

²⁵For example, suppose t is the earliest equilibrium bid time for any bidder, or the equilibrium bid times are such that no bid is supposed to arrive at any point of time in $[t - \Delta, t)$. In such cases, a deviation is detected with certainty.

²⁶Such a punishment might or might not be credible, but we are simply showing that even under the worst possible punishment the payoff exceeds zero. Therefore, for any other strategy by other bidders the payoff exceeds zero as well.

s. Further, with strictly positive probability no other bidder draws a value above R_0 , and even when some others do draw values above R_0 , with strictly positive probability no such bid of \bar{v} arrives. Similarly, the worst shill bid - a shill bid strictly greater than v - fails to reach with strictly positive probability. Thus the payoff is always nonnegative, and is strictly positive with strictly positive probability. Since this is true in the worst case, this is true for all cases. It follows that for any arrival time $s \in [t - \Delta, t)$, $\pi(s) > 0$. Therefore, $P(t - \Delta) - P(t) > 0$, and the deviation is beneficial, which gives us a contradiction. This completes the proof. ||

B Technical Appendix

In this appendix we provide a formal description of strategies and equilibrium used in the paper. In particular, the following issues are resolved.

1. Since our auction is a continuous time game some care needs to be taken to ensure that strategies are defined in a *coherent* manner. We first define strategies formally and then note the usual problem: since the real line is not well ordered, for any time t there is no natural notion of “last time before t ” nor the “next instance after t .” So, for example, continuous time makes it difficult to interpret a strategy that says “bid at the next instant after t .” The problem has well-known “fixes” in the literature, which, as we show, apply here.
2. Second, we must consider modelling events that take place “sequentially” and yet “at the same instance.” If a bidder arrives at t , we allow the bidder to bid at any time in $[t, 1]$, so two events: bidder arrival and bidding could happen at the same time, although the two are naturally ordered. Similarly, if a bid arrives at t resulting in a price jump at t , we allow the seller to shill bid at t . Two events take place at the same instant though they are naturally ordered. It is not difficult to resolve this formally - we do this by taking limits of small inertias.²⁷

²⁷Indeed, the issue of coincident but sequential events is no more difficult or challenging than that involved in other continuous time auctions such as the standard Japanese type English auction (see, for example, Milgrom and Weber (1982)) where price rises continuously with each bidder keeping a button pressed to signal their participation, and the instant of time at which a bidder’s willingness-to-pay is reached, the bidder releases the button. So two events - price reaching a specific point and a bidder exiting by releasing

3. Finally, our existence result (Theorem 1) involves the seller updating beliefs. We give an informal description in the main text; in Section B.2 we provide formal details.

B.1 Strategies and continuous time issues

B.1.1 Bidders' strategies

We first define bidders' strategies formally. In doing so, we follow the strategy construction of Bergin and Macleod (1993).²⁸ While the bidder's and the seller's strategies are not the same, the technical issues regarding continuous time are the same for both. We therefore define bidders' strategies in some detail and briefly describe the seller's strategies.

Recall (from Section 2.1) that the action choice of bidder i arriving at \underline{t}_i for all $t \geq \underline{t}_i$ is to either remain inactive or be active and submit a bid. For now we ignore the issue that a bidder may bid *immediately* upon arrival, which implies two events - arrival and bidding happen at the same instance. We resolve this in section B.1.4 below.

Denote the arrival time of i by $\underline{t}_i \in [-T, 0]$. Recall that $b_{i,t}$ denotes the bid level of bidder i at time t . Let $B_{i,t}$ denote the set of feasible bids for bidder i at time $t \in [\underline{t}_i, 1]$. This evolves with time as follows.

For bidder i arriving at time \underline{t}_i , let bid level $b_{i,\underline{t}_i} = 0$. Let $t' \geq \underline{t}_i$ be the first instance such that bidder i submits a bid at time t' , where $B_{i,t'} = [0, v_i]$. Let the bid be $b' \in B_{i,t'}$. In that case *actions* of bidder i are $b_{i,t} = 0$ for $t \in [\underline{t}_i, t')$, and $b_{i,t} = b'$ for $t = t'$. More generally, if bidder i submits a bid \tilde{b} at time \tilde{t} , submits, at time $\hat{t} > \tilde{t}$, the next bid $\hat{b} \in (\max\{\tilde{b}, p_{\hat{t}}\}, v_i]$ and remains inactive in the intervening period, then his actions are $b_{i,t} = \tilde{b}$ for $t \in [\tilde{t}, \hat{t})$ and $b_{i,t} = \hat{b}$ for $t = \hat{t}$. Further, for $t \in (\hat{t}, 1]$, $B_{i,t} = (\max\{\tilde{b}, p_t\}, v_i]$.

Finally, the following restriction applies to actions: if bidder i submits a bid at any time $t > 0$, then he must be inactive over the intervals $(0, t)$ and $(t, 1]$.

The above implies the following restrictions on the path of actions: a bidder either remains inactive (bid level remains unchanged), or submits a bid (bid level jumps *upwards*). Further, any bid has to be greater than the current auction price. Finally, only a single bid

a button, takes place at the same instant of time - but it is implicit that there is a natural sequence and this causes no interpretation problem.

²⁸See also Simon and Stinchcombe (1989) for a related approach.

is allowed for $t > 0$.

Private action path of player i is the function $b_{i,t}$ for all $t \in [\underline{t}_i, 1]$. Private history, $h_{i,t}$ at time t includes, value of the object v_i , arrival time \underline{t}_i , private action path till time t , i.e., $b_{i,\tau}$ for $\tau \in [\underline{t}_i, t)$ as well as public history h_t . Let $H_{i,t}$ be the set of all possible histories at time t .

Let B_i denote the union of $B_{i,t}$ for $t \in [\underline{t}_i, 1]$. Let H_i denote the union of $H_{i,t}$ for all $t \in [\underline{t}_i, 1]$. Following the construction of Bergin and Macleod (1993), the strategy of bidder i is

$$\sigma_i : H_i \times [\underline{t}_i, 1] \rightarrow B_i$$

where for any h'_i and h''_i in H_i , if they agree (with respect to the metric introduced by Bergin and Macleod) on the interval $[\underline{t}_i, t)$, the action $b_{i,t}$ at time t is the same for h' and h'' . The last restriction simply states that strategies cannot use information from the future.

B.1.2 Continuous time issues

As pointed out by well-known work by Simon and Stinchcombe (1989), Bergin and Macleod (1993), the strategy defined above can be problematic for a game in continuous time: it may not define a unique outcome. To deal with the issue that strategies in extensive form games should define outcome paths uniquely, we make use of the concept of **inertial strategies** introduced by Bergin and Macleod (1993). The intuitive idea is that submission of bids require a reaction time (however small): any bidder (but this applies to the seller also) who has submitted a bid at time t may not submit the next bid till some time has elapsed. More formally, consider a time t such that $b_{i,t} > \sup_{\tau \in [\underline{t}_i, t)} b_{i,\tau}$. An ε -**inertial strategy** is a strategy such that for all t as defined above $b_{i,\tau} = b_{i,t}$ for $\tau \in [t, t + \varepsilon)$. Let Σ_i^ε denote the collection of all ε -inertial strategies for player i and Σ^ε the cartesian product. It follows from Theorem 1 in Bergin and Macleod (1993) that inertial strategies define unique outcomes. Strategies for the actual continuous time auction game then are the closure of the set of inertial strategies.²⁹

²⁹See Bergin and Macleod (1993) for details regarding the construction of the metric with regard to which convergence of strategies are defined. Action sets in our context are closed, bounded intervals; constructions in Bergin and Macleod allow for actions sets to be compact subsets of arbitrary metric spaces.

B.1.3 Seller's strategies

Recall that $b_{s,t}$ denotes the shill bid level of the seller at time t . This is a step-function as for any bidder's bid level with the following extra restriction. Let \mathcal{T}_p denote the set of active times in the auction (the set of times at which the price moves up). The set of active times for the seller is a subset of \mathcal{T}_p . Further, $b_{s,t}$ is determined by calculating the optimal reserve price given the information content of public history h_t at t . The updating process given a specific equilibrium is described in section B.2. Let $B_{s,t}$ denote the set of feasible shill bids for the seller at time $t \in [-T, 1]$.

Private action path of the seller is the function $b_{s,t}$ for all $t \in [-T, 1]$. Given the updating described above, this automatically has the property that the bid level can stay unchanged or jump upwards. Further, as for bidders, a bid at any t is required exceed the current auction price p_t . Private history $h_{s,t}$ at time t includes private action path till time t , i.e., $b_{s,\tau}$ for $\tau \in [-T, t)$ as well as public history h_t . Let $H_{s,t}$ be the set of all possible histories at time t .

Let B_s denote the union of $B_{s,t}$ for $t \in [-T, 1]$ and let H_s denote the union of $H_{s,t}$ for all $t \in [-T, 1]$. The strategy of the seller is given by a function

$$\sigma_s : H_s \times [-T, 1] \rightarrow B_s$$

where for any h'_s and h''_s in H_s , if they agree on the interval $[-T, t)$, the action $b_{s,t}$ at time t is the same for h' and h'' . The last restriction simply states that strategies cannot use information from the future.

The same problem as for bidders' strategies arises here under continuous time, and can be resolved in the same way by considering strategies for the actual continuous time auction game as the closure of the set of inertial strategies.

B.1.4 Coincident events

The one remaining issue is that we allow for naturally ordered events that, nevertheless, take place at the "same instance" of time in the model.

First, consider the issue of bid arrival. The arrival time of a bid (by either a bidder or the seller) submitted at time t is uniformly distributed over $[t, t + 1]$ (or over $[t, 1]$ if $t > 0$), thus arriving after t with probability 1. Since arrival of a bid coincides with bid time with

0 probability, this causes no formal issue. However, we could also assume that there is a minimum technology inertia so that a bid at t arrives only over $[t + \delta, t + 1]$ (or over $[t + \delta, 1]$ if $t > 0$) where δ is a small positive number, but since setting δ to zero is without loss of generality here, and avoids unnecessary notation, we do so.

Second, a bidder arriving at time t_i can also submit a bid at time t_i . Similarly, following a change in auction price at period t , the seller submits a shill bid “instantaneously” at time t . To interpret, introduce small inertias as follows. A bidder arriving at t requires time $\varepsilon_1 > 0$ to “understand the state of the auction and settle down” and can start observing history and take any action not before $t + \varepsilon_1$. The seller is around from the very start of the auction and not subject to this inertia. Second, after a price movement at t , every agent (bidders and the seller) is subject to a “cognitive inertia” (i.e. time to process the price movement and decide on any action) of $\varepsilon_2 > 0$ so that the agent cannot submit a bid before $t + \varepsilon_2$.

Given these small inertias, all events are ordered in real time: if a bidder arrives at t , he cannot bid before $t + \varepsilon_1$ and if price changes at $t + \varepsilon_1$, cognitive inertia pushes earliest bidding time to $t + \varepsilon_1 + \varepsilon_2$. Further, given a price movement at any t , the earliest time at which the seller can shill bid is $t + \varepsilon_2$. Thus all events - bidder arrival and bidding, price movement and bidding - are ordered in real time.

We could carry these inertias in the analysis without changing any result, but this encumbers notation unnecessarily. To simplify, we consider the limit as ε_1 and ε_2 go to 0.

B.2 Updating the seller’s beliefs

Recall from the model specification that $\mu_i, i = 1, 2, \dots, H$ is the seller’s prior belief that buyer values are drawn from distribution F_i . For any t , let $\mu_{i,t}$ be the (updated) posterior beliefs. Recall that the seller’s strategy for Theorem 1 involves updating beliefs given the specified bidders’ strategies. In this section, we illustrate the belief-updating by the seller by explicitly describing a few updates.

Before we do that, it is worth noting that the specific details of the updating procedure do not matter for our results. The results simply require that shill bids are submitted with positive probability. For the bidders to delay submitting their bids all that is required is the specification in section 2.2 that there is scope for the seller to do *some* updating so

that she has the incentive to shill bid. Also recall that, while the bidders delay submitting their bids, they do not delay beyond the time $t = 0$ (Proposition 2, Proposition 2A). Again, these results do not make any use of the exact updating rule used by the seller.

For ease of exposition, it is useful to consider separately the beliefs on and off the equilibrium path. In fact we break up the analysis for time $t \in [-T, 0)$ and for $t \in [0, 1)$.

Updating for time periods $t < 0$ Since the strategies of all bidders is to remain inactive—irrespective of history—till time 0 and submit bids at time $t = 0$, over the time interval $[-T, 0)$ all inactive periods are on equilibrium path and any active period is off-path. Further, given bidders’ strategies, all (on-path) inactive periods $t \in [-T, 0)$ are uninformative events. All active periods are off-path and Perfect Bayesian Equilibrium (PBE) imposes no restriction on seller’s beliefs; the updating we describe is consistent with a theory that posits that all types have the same probability of a “tremble” of deviating from their equilibrium strategy and placing a bid equal to their value.

Specifically, let τ_1 be the first active period in $(-T, 0)$. Then a straightforward application of Bayes’ Rule shows that μ_{i,τ_1} is given by (recall that R_0 is the initial reserve price)

$$\mu_{i,\tau_1} = \frac{[\mu_i][1 - F_i(R_0)]}{\sum_{k=1}^H [\mu_k][1 - F_k(R_0)]}. \quad (\text{B.1})$$

There are now two possibilities: the seller submits a shill bid or she doesn’t. Suppose first the case when the seller does not submit a shill bid at time τ_1 and let the next active period be $t = \tau_2 > \tau_1$. As mentioned earlier, for all $k = 1, 2, \dots, H$, the seller’s updated beliefs remain μ_{k,τ_1} for all inactive periods between τ_1 and τ_2 . Now, since τ_2 is an active period, the auction price changes and let the price at τ_2 be denoted as p_{τ_2} . Then again applying Bayes’ Rule the seller’s updated beliefs are given by

$$\mu_{i,\tau_2} = \frac{[\mu_{i,\tau_1}][1 - F_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][1 - F_k(p_{\tau_2})]}. \quad (\text{B.2})$$

At this point, again, the seller may or may not submit a shill bid. If she doesn’t then her posterior beliefs continue to remain μ_{k,τ_2} till another possible active period $\tau_3 \in (\tau_2, 0)$ and the updating follows the same procedure as shown above given the auction price p_{τ_3} at time τ_3 , and so on.

Next, consider the situation when the seller submits a shill bid at time τ_1 , the first active period. Let s_{τ_1} denote the value of seller’s shill bid. As before, we let the next active

period be denoted by τ_2 ; note that if the seller shill bids at τ_1 then the latest point in time for τ_2 is $\tau_1 + 1$.

Now, given that the seller shill bids an amount s_{τ_1} at time τ_1 , the price p_{τ_2} can be greater than, equal to, or less than s_{τ_1} .

If $p_{\tau_2} < s_{\tau_1}$, then the seller's updated posterior beliefs are given by

$$\mu_{i,\tau_2} = \frac{[\mu_{i,\tau_1}][f_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][f_k(p_{\tau_2})]} \quad (\text{B.3})$$

and if $p_{\tau_2} = s_{\tau_1}$, then the seller's updated posterior beliefs are given by

$$\mu_{i,\tau_2} = \frac{[\mu_{i,\tau_1}][1 - F_i(s_{\tau_1})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][1 - F_k(s_{\tau_1})]}. \quad (\text{B.4})$$

If however, $p_{\tau_2} > s_{\tau_1}$, then the seller updates her beliefs to

$$\mu_{i,\tau_2} = \frac{[\mu_{i,\tau_1}][1 - F_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][1 - F_k(p_{\tau_2})]}. \quad (\text{B.5})$$

And the updating process continues similarly for future active periods $t < 0$.

A remark before we continue. Notice that for $\tau_2 < \tau_1 + 1$, it is possible that the seller's shill bid has not reached yet and the activity is caused by yet another bidder bidding before time 0. Off-equilibrium beliefs are unrestricted under PBE. We choose these by having the posterior beliefs satisfy a *consistency* condition (in the same spirit as Sequential Equilibrium) as follows. Notice that the seller's shill bid reaches by τ_2 with strictly positive probability whereas "trembles" of the bidders converge to zero. This explains the right hand side of equation B.3 and B.4; for these two cases, the seller's updating reflects the fact that the activity at time τ_2 is "infinitely more likely" to be caused by arrival of the seller's shill bid than by yet another tremble (i.e., a tremble other than the one that created the first active period τ_1). It is only when $p_{\tau_2} > s_{\tau_1}$ that the seller has to believe that there is some other bidder who has also trembled. This explains the right hand side of equation B.5.

Next we consider the time period beyond $t = 0$.

Updating at time periods $t \geq 0$ Now, provided there has been no active period prior to $t = 0$, all active and inactive periods for $t > 0$ are on the equilibrium path.³⁰ Let τ_1 be the

³⁰If there have been active periods prior to $t = 0$ then formally, everything that happens afterwards is off-equilibrium path. However, since the requirements of PBE is to use strategies (along with Bayes' Rule)

first active period. Along the equilibrium path, $\mu_{i,t} = \mu_i$ for $t < \tau_1$ and hence the seller updates her belief exactly as in equation B.1.

Once again, the seller may or may not optimally decide to shill bid. If she doesn't then she updates her beliefs as in equation B.2. If the seller does shill bid at τ_1 then her updating at τ_2 is similar in many ways to what we discussed earlier but with some important differences. So, suppose that the seller does shill bid at τ_1 , and let the bid be equal to s_{τ_1} . Again, we consider three cases based on whether the price p_{τ_2} is greater than, equal to, or less than s_{τ_1} . The first two cases are relatively simple, the seller's updating in these cases are as shown in equation B.5 and equation B.4.

The case that is different from what we had earlier—and is a bit more involved—is when $p_{\tau_2} < s_{\tau_1}$. It is possible that the price movement at τ_2 is caused by the seller's shill bid (which she had submitted at time τ_1) reaching or it could be due to arrival of some other bidder's bid. And unlike the case for $t < 0$, now the latter events are non-negligible (in terms of probability of occurring).

Note that *if* the seller knew that the price movement at τ_2 is caused by the seller's shill bid, then her updated posterior belief would be equal to

$$\frac{[\mu_{i,\tau_1}][f_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][f_k(p_{\tau_2})]}.$$

On the other hand, if she knew that it is caused by the arrival of the bid by some other bidder, then her updated posterior belief would be equal to

$$\frac{[\mu_{i,\tau_1}][1 - F_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][1 - F_k(p_{\tau_2})]}.$$

To obtain the posterior, we require a few more steps. In what follows, recall that the density of the uniform distribution - the bid arrival process - is unity.

First, suppose the number of participating bidder is $n \geq 1$. The prior probability that the activity at time τ_2 is caused by the seller's bid reaching before the bids of the $n - 1$ other bidders reach is³¹ equal to $(1 - \tau_2)^{n-1}$. On the other the prior probability that one of the

to derive beliefs "whenever possible", it is easy to see—from the discussion that follows—what the seller's updated posteriors are. Since bidders who have not deviated are supposed to follow their prescribed equilibrium strategies, the seller's updating is similar to what she would do if there had been no deviation but with a major difference that her belief at $t = 0$ in that case are not her initial priors μ_i .

³¹Recall that bid of one of the n has already reached at time $t = \tau_1$.

$n - 1$ bids has reached and neither the seller's shill bid nor the other $n - 2$ has reached is $(n - 1)(1 - \tau_2)^{n-2}(1 - (\tau_2 - \tau_1))$. Therefore, conditional on the number of participating bidders being n , the seller's posterior belief, let us call it H_n , is

$$H_n = \frac{(1 - \tau_2)^{n-1}}{[(1 - \tau_2)^{n-1}] + [(n - 1)(1 - \tau_2)^{n-2}(1 - (\tau_2 - \tau_1))]} \frac{[\mu_{i,\tau_1}][f_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][f_k(p_{\tau_2})]} \\ + \frac{(n - 1)(1 - \tau_2)^{n-2}(1 - (\tau_2 - \tau_1))}{[(1 - \tau_2)^{n-1}] + [(n - 1)(1 - \tau_2)^{n-2}(1 - (\tau_2 - \tau_1))]} \frac{[\mu_{i,\tau_1}][1 - F_i(p_{\tau_2})]}{\sum_{k=1}^H [\mu_{k,\tau_1}][1 - F_k(p_{\tau_2})]}.$$

Of course the seller does not know the number of bidders who have arrived, but her updated belief that the number is n can be derived as follows.

Recall from section 2 that $\lambda_j > 0$ denotes the prior that j bidders have arrived, $j \in \{0, \dots, N\}$. For $t > 0$, as bids arrive, the seller updates this based on the arrival-time sequence (τ_1, τ_2, \dots) . Let y_0 denote the (null) arrival-time sequence when no bids have arrived, and y_ℓ denote the history of arrival-times up to time τ_ℓ , the time of the ℓ -th price movement. Then $\lambda_n \equiv \lambda_n(y_0)$. Further,

$$\lambda_n(y_1) \equiv \Pr(n \text{ bidders} | \text{price movement at } \tau_1) = \frac{\lambda_n(y_0) \Pr(\tau_1 | n)}{\sum_{j=1}^N \lambda_j(y_0) \Pr(\tau_1 | j)} \\ \lambda_n(y_2) \equiv \Pr(n \text{ bidders} | \text{price movements at } \tau_1, \tau_2) = \frac{\lambda_n(y_1) \Pr(\tau_2 | n)}{\sum_{j=1}^N \lambda_j(y_1) \Pr(\tau_2 | j)}$$

where $\Pr(\tau_1 | n) = n(1 - \tau_1)^{n-1}$ and $\Pr(\tau_2 | n) = (n - 1) \frac{1}{1 - \tau_1} \left(\frac{1 - \tau_2}{1 - \tau_1} \right)^{n-2}$. It follows that, in the case under consideration (the seller shill bids at τ_1 and we have $p_{\tau_2} < s_{\tau_1}$), the seller's updated posterior belief that the buyers' values are drawn from F_i is

$$\mu_{i,\tau_2} = \sum_{n \geq 1} \left(\frac{\lambda_n(y_2) H_n}{\sum_{j=1}^N \lambda_j(y_2)} \right).$$

Note that at this point the seller knows that at least 1 bidder has definitely arrived. This explains why the calculation excludes the case $n = 0$.³²

A similar updating procedure applies for future active periods.

³²In general, if the seller knew at any time t that at least $m \geq 1$ bidders had definitely arrived (this could happen for example if, after the reserve price becoming active, the auction price moved a further $m - 1$ times by time t and the seller had not submitted any shill bids before t) the calculations for updating would only consider cases in which number of bidders were at least m .

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, 3
- Bergin, James and W. Bentley Macleod**, “Continuous Time Repeated Games,” *International Economic Review*, 1993, 34 (1), 21–37. 27, 28
- Chakraborty, Indranil and Georgia Kosmopoulou**, “Auctions with shill bidding,” *Economic Theory*, 2004, 24, 271–287. 5
- Engelberg, Joseph and Jared Williams**, “eBays proxy bidding: A license to shill,” *Journal of Economic Behavior and Organisation*, 2009, 72, 509–526. 4
- Graham, Daniel A., Robert C. Marshall, and Jean-Francois Richard**, “Phantom bidding against heterogeneous bidders,” *Economics Letters*, 1990, 32 (1), 13–17. 4
- Kosmopoulou, Georgia and Dakshina G. De Silva**, “The effect of shill bidding upon prices: Experimental evidence,” *International Journal of Industrial Organization*, 2007, 25, 291–313. 5
- Lamy, Laurent**, “The Shill Bidding Effect versus the Linkage Principle,” *Journal of Economic Theory*, 2009, 144, 390–413. 5
- Milgrom, Paul R. and Robert J. Weber**, “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 1982, 50 (5), 1089–1122. 26
- Myerson, Roger**, “Optimal Auctions,” *Mathematics of Operations Research*, 1981, 6, 58–63. 2
- 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, 3, 4
- 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). 3

- 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
- Simon, Leo K. and Maxwell B. Stinchcombe**, "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 1989, 57 (5), 1171–1214. 27, 28
- The Sunday Times**, "Revealed: how eBay sellers fix auctions," 2007. January 28. 2
- Walton, Kenneth**, *Fake: Forgery, Lies, & eBay*, Gallery Books, 2006. 2
- Whitworth, Dan**, "Man fined over fake eBay auctions (BBC Newsbeat)," 2010. July 5. 2
- Wintr, Ladislav**, "Some Evidence on Late Bidding in eBay auctions," *Economic Inquiry*, 2008, 46 (3), 369–379. 1