Birkbeck Working Papers in Economics & Finance

ISSN 1745-8587

Birkbeck UNIVERSITY OF LONDON

Department of Economics, Mathematics and Statistics

BWPEF 1510

# Shills and Snipes

***Subir Bose***
*University of Leicester*

***Arup Daripa***
*Birkbeck, University of London*

February 2015

▪ Birkbeck, University of London ▪ Malet Street ▪ London ▪ WC1E 7HX ▪

# Shills and Snipes

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

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

February 2015

**Abstract**

Online auctions with a fixed end-time often 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 study a correlated private values environment in which the seller bids in her own auction (shill bidding). Bidders selected randomly from some large set arrive randomly in an auction, then decide when to bid (possibly multiple times) over a continuous time interval. A submitted bid arrives over a continuous time interval according to some stochastic distribution. The auction is a continuous-time game where the set of players is not commonly known, a natural setting for online auctions. We show that there is a late-bidding equilibrium in which bids are delayed to the latest instance involving no sacrifice of probability of bid arrival, but shill bids fail to arrive with positive probability, and in this sense optimal late bidding serves to snipe the shill bids. We show conditions under which the equilibrium outcome is unique. Our results suggest that under private values, the case against shill-bidding might be weak.

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 time (“hard 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.[1]

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.[2] However, a large fraction of auctions on online platforms such as eBay fit the private values paradigm well, and experience significant amount of sniping.[3,4] 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

[1] Under 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, the latter wins.

[2] See Bajari and Hortaçsu (2003), Ockenfels and Roth (2006).

[3] 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. 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 much better than a common values one. While the quality of, say, a laptop may indeed vary affecting in a similar fashion the payoff of anyone who buys it, 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.

[4] It 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.

using a framework in which bid times are chosen in a continuous manner and bid arrival times are continuous but random. Further, we allow for random bidder arrival and, as in the case with real life online auctions, neither the identities nor the actual number of buyers is assumed to be common knowledge.

We show that there is an equilibrium with a "late-bidding" outcome that is both natural and intuitive. Moreover, such an outcome is unique under a monotonicity condition (discussed below) on bidder strategies. Further, in the absence of monotonicity, if there is any equilibrium with a different outcome, it is Pareto dominated for bidders compared to an equilibrium with late bidding. Importantly, the results are unaffected by either the set of bidders who participate or by the exact nature of the random arrival process. Nor are they affected by the exact nature of the priors over these. Therefore the results have a "robustness" with respect to aspects of the environment about which at least the modeler, and also perhaps the seller and the bidders, are unlikely to have precise information.

While we consider private value online auctions, the framework that we develop can be adapted usefully to other online auction settings, or more generally to settings like bargaining with deadlines and randomly arriving outside options. Let us clarify how the framework differs from the literature. The literature on sniping in online auctions (considering private and common values and other richer settings) typically assume a discontinuous timing setting where a certain "last point of time" plays a special role. All bids arrive with probability 1 before the last point of time, irrespective of how close to this last point of time they have been made. Bids made at the last point of time fail to arrive with some exogenously given probability. Further, other bidders can respond to a bid made before the last point in time with certainty but are completely unable to respond if the bid is placed at the last point. In contrast, in our model, there is a fixed end time for the auction which is no different from other, regular, times, and the bidders choose how close to the fixed end time they would like to place their bid. In other words, the arrival probability of bids - as also how much time one leaves for one's rivals to react to one's action - is a matter of endogenous choice. The questions addressed in the literature could be analyzed in the more natural continuous setting of our framework and while some of the results would continue to hold, there are others that depend crucially on a discontinuity between regular time and the last point of time and would unravel if this discontinuous arrival process of bids is removed. We discuss this further later.

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 problematic. Legal or not, shill bidding is reported to be widespread in online auctions. Direct evidence arises from prosecutions of or admissions by shill bidders. [5] However, the ease of creating alternate identities implies that it is hard to verify shill bidding to a standard of evidence required for prosecutions. Nevertheless, shill bidding is a much discussed topic in online discussion forums for bidders, and such forums provide a rich seam of evidence that shill bidding (real or perceived) is a matter of significant concern to bidders.[6] It is therefore natural to expect that they would bear this in mind and try not to cede any advantage to a potential shill bidder when formulating bidding strategies.

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, 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.[7]

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. 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 openly adjust the reserve price mid-auction–allows scope for profitable shill bidding.

[5]See, 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.

[6]See, for example, results for "shilling" on the site community.ebay.com, or `http://forums.moneysavingexpert.com/showthread.php?t=4583383`.

[7]There 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. Footnote 26 comments further on this issue.

We show that 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 continuous choice gives rise to interesting results on bid-timing and efficiency. To clarify, consider the model we analyze. We assume that the auction takes place over a time interval stretching from $-T$ to 1. The interval $[-T, 0)$ is the "early" period and $[0, 1]$ is the "last minute" (with evolving technology this is in practice a short interval of time, perhaps a few seconds rather than an actual minute). A random selection (from some large set) of bidders enter randomly over $[-T, 0]$, and can submit proxy bids at any time after they enter. Bidders can submit one or more bids leading up to time 0, and also choose to make a bid inside the last minute (at some point $\hat{t} \in (0, 1)$). The arrival time of a bid made at time $t$ is uniformly distributed on $[t, t+1]$ so long as $t \leqslant 0$ (so that $t+1$ is within end time 1). In other words, early bids made at some $t \leqslant 0$ arrive eventually with certainty. The last point of time at which this property holds is time 0. This is the cusp of the last minute. A bid made after time 0 is made at some time $\hat{t}$ "inside" the last minute (i.e. at a time $\hat{t} \in (0, 1)$). Any such bid fails to arrive with probability $\hat{t}$. Bidders could bid at time 0, or sacrifice some arrival probability by pushing their bid times to later than 0. These strategies have different implications for revenue and efficiency. We are therefore interested in determining the precise time of bidding.[8] The seller chooses an initial reserve price; this is part of the description of the auction. The seller's strategy consists of submitting shill bids at a finite number of time points.

We show that in any equilibrium, so as to reduce the chances of the seller submitting shill bids successfully, the bidders optimally bid at the very last possible time such that their bids reach with probability 1. As noted above, in our model this is time 0. Indeed, bidding exactly at time 0 is an equilibrium outcome, one that is unique under a monotonicity condition. Therefore, whenever the auction results in a sale, there is no loss of efficiency. Further, we show that as a result of the bid-time selection by bidders, all shill bids necessarily arrive with probability strictly less than 1. Thus the shill bidder is sniped in the sense that there is a positive probability that some or all shill bids do not arrive.

As mentioned earlier, shill bidding is generally considered harmful to bidders' interests. However, it follows from our results that a clear conclusion cannot be drawn. As already noted, all genuine bids eventually arrive resulting in the sale being efficient when

[8] Of course, if bidders do not arrive by time 0, they would necessarily bid after time 0. Since our focus is on studying the *choice* to bid early or late rather than forced late bidding from late arrival, we assume that all bidders arrive before time 0.

it happens (efficiency "at the top"). Second, the fact that the seller can increase (probabilistically) the reserve price while the auction is ongoing means that the seller might choose a *lower* initial reserve price compared to the case where shill bids are ruled out. This, combined with the fact that some or all shill bids may fail to arrive implies that in some cases the auction with shills realizes trading gains that would be lost in the standard case without shill bids. Thus the auction with shills may be more efficient "at the bottom."

Let us now discuss the specific results of our paper. We first show that the seller's equilibrium strategy involves submitting shill bids only at times when the auction registers some activity (the reserve price is met or auction price jumps up). To see why this matters, consider the following example of a "threat" strategy by the seller to induce early bidding by the bidders: the seller submits a "high" shill bid if there has been no activity till some time $t < 0$. If the bidders believe this threat then that might induce them to bid early, allowing the seller to shill bid more successfully. Our result shows that such threats are not credible. In addition to being of some independent interest, this result plays an important role in the analysis that follows.

Next, we show that it is suboptimal for the bidders to only bid at a time $t > 0$ (i.e. only bid "inside" the last minute). This establishes the crucial result that it is not worthwhile to sacrifice probability-of-bid-arrival in order to snipe shill bids.

Given these results characterizing equilibria, the natural next step is to establish that an equilibrium does exist. We show that there is indeed an equilibrium in simple strategies: submit a truthful bid at time $t = 0$ irrespective of history up to time 0, and remain inactive thereafter. We also define a monotonicity condition under which bidding at time 0 by all bidders (of types above the initial reserve price) is the *unique* equilibrium outcome. If there is any equilibrium with early bidding (bidding at some time $t < 0$), it must therefore involve non-monotonic strategies. The final result shows that from the viewpoint of the bidders, any such equilibrium is Pareto dominated by the one involving bidding at time 0.

Since the property of monotonicity of strategies plays a crucial role in some of our results, we now describe it and explain the intuition behind the role it plays.

Suppose strategies have the following simple structure: at any time $t$, there is a trigger price $p_t^i$ for bidder $i$, who bids at $t$ if and only if the auction price at $t$ equals or exceeds $p_t^i$. We call such strategies "monotonic." Consider a simple setting with two bidders. Let "bidding early" denote bidding at any time before 0. If bidder 2 follows a monotonic strategy, bidder 1 cannot be induced to bid early. To see why, note that bidding early implies 1's bid would arrive early and trigger a shill bid early with positive probability. Also, bidding early implies 2 might bid early (if the auction price reaches $p_t^2$ at an early $t$), which in turn might again trigger a shill bid early. Since delaying bidding till time 0 involves no loss of bid-arrival probability, but delays triggering shill bids, early bidding by bidder 1 is suboptimal.

Can bidders be induced to bid early in an equilibrium? This can happen only if the expected payoff from deviating (and not bidding early) is sufficiently bad. But how can that happen? We have argued that it is not possible for the seller to induce early bidding in a credible way. Hence the only way for a bidder (1 say) to bid early in equilibrium is for bidder 2 to have a strategy in which 2 bids at some time $t < 0$ if the auction price is *lower* than some threshold level but not bid if it is *higher*. The first part acts as the punishment against deviation. A sufficiently harmful punishment would induce bidder 1 to bid early to raise the probability that the auction price crosses the threshold at time $t$. Note however, that this (somewhat strange) strategy of player 2 violates monotonicity.[9]

Since it is suboptimal for bidders to delay bidding beyond time 0 and since under monotonicity they do not bid before time 0, it follows that the unique outcome is to bid at time 0. Further, the strategies we propose to show existence are monotonic by construction. Therefore any deviation by a bidder is contemplated in a setting where others adopt monotonic strategies. It follows that the logic that guarantees uniqueness also proves existence.

To see how our results differ from the literature, 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

[9] In 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."

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 a "hard" end time is crucial here as it allows bidders to snipe the shill bidder by delaying bids. Alternatively, an auction could have a "soft" ending so that if a bid arrives in the last 5 minutes (say), the end time is automatically extended until there is no bidding activity for 5 continuous minutes - a format used by uBid.com. A soft ending precludes sniping. As Roth and Ockenfels (2002) report, the contrast between the two formats gives rise to interesting differences in bidding data across auctions. Auctions on eBay, which uses a hard end time, have substantially greater late 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.[10]

## Relating to the Literature

In our paper, bidders want to delay 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 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

[10]In our formal model, all buyers arrive before time $t = 0$ to avoid trivial last minute bidding. If buyers can arrive after time 0, there would be some late bidding even in soft end proxy auctions. Holding constant buyer arrival rate, one would nevertheless expect more late bidding in hard end rather than soft end auctions which is what the data seems to suggest.

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 timing 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.

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 late bidding is the desire to reduce the probability that shill bids arrive successfully. Another difference arises because we allow bid times to be chosen continuously. This implies a further connection to the literature. We could ask how the results in the literature would be modified in our continuous timing framework.

As noted above, the "collusive" equilibrium of Ockenfels and Roth (2006) cannot arise in a framework of continuous bid and arrival times. But other types of equilibria discussed above–late bidding in a common values framework to discourage aggressive bidding, or an expert bidding late to hide information–could also arise using our continuous-bid-and-arrival-time framework. An interesting question then is to determine the optimal bid time in these cases.[11]

Finally, consider the question of shill bidding. Graham, Marshall and Richard (1990) and Bag, Wang and Dinlersoz (2000) consider the question of reserve price updating under private value settings.[12] 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, but obviously such calculations need not apply when behavioral biases or naive decision-making dictate bid-time selection. In such contexts,

[11] 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. This, too, could be analyzed using our continuous approach.

[12] See also Deltas (1999) who considers bid raising by a dishonest seller.

our work can be seen as a benchmark model with rational bidders, and explaining any observed departures from our conclusions would then require incorporating more complex environments or behavioral biases.

Chakraborty and Kosmopoulou (2004) and 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 abstain from shill bidding, such bids arise in equilibrium.[13]

The rest of the paper is organized as follows. The next section presents the model. Section 3 presents characterization results under monotonicity. Section 4 then proves our main results on equilibrium bid-timing under shill bidding. Finally, section 5 concludes. Proofs not in the body of the paper are collected in the appendix.

## 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. This is a second-price auction. The seller can post a reserve price at the beginning and also submit shill bids during the auction.

Bidders are drawn randomly from some large set $\mathcal{N}$ of potential bidders. Bidder arrival in the auction is allowed to be random. The seller as well as each bidder therefore faces a random subset of other bidders. The arrival process is independent of the distribution of valuations or actual bidder valuations.

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. In our model the arrival of bidders to the auction is stochastic and private information and consequently neither the seller nor the bidders observe the actual number of bidders. While our model

[13] See also Kosmopoulou (2011) who analyzes "off the wall" bidding by the seller under common values.

involves standard Bayesian rational agents who therefore have priors over different subsets of arriving bidders, the equilibrium outcome we are interested in holds *irrespective* of which subset arrives. Thus the prior belief plays no role in our proofs implying that the exact nature of the prior belief is unimportant. Therefore our results have a certain robustness property not often found in (perfect) Bayesian Nash equilibria of many other auction models.

**Bidder valuation** We analyze shill bidding under a correlated private values setting. Let $\mathcal{F}$ denote a set of distributions $F_1, \ldots, F_H$, where each distribution has support $[\underline{v}, \overline{v}]$, where $0 \leqslant \underline{v} < \overline{v} < \infty$.[14] Nature chooses a distribution $F_k$ from the set $\mathcal{F}$ and the bidders' values are determined according to independent draws from the distribution $F_k$. Each bidder privately observes its own value. Neither the bidders nor the seller observes $F_k$ but has some prior belief over $\mathcal{F}$.[15] We assume that the distributions are ordered in terms of likelihood ratio property. In other words, a higher value of $v$ is more likely to have been generated from a distribution $F_{k''}$ than the distribution $F_{k'}$ for $k'' > k'$. Since dominance in terms of likelihood ratio implies dominance in terms of hazard rates, this also implies that the optimal reserve price will be higher for distribution $F_{k''}$ than for $F_{k'}$. This is important in providing a reason for shill bidding. If–as in second price auctions–bids reflect true values, increase in current price (current second highest bid) in the auction results in updated posterior beliefs inducing the seller to either maintain status quo or, with positive probability, to want to *raise* the reserve price.

[14] We use the word "support" a bit loosely. We allow for intervals $[\underline{v}_k, \overline{v}_k]$ for $k \in \{1, 2, \cdots, H\}$ to be different, and then define $\underline{v} = \min \underline{v}_k$ and $\overline{v} = \max \overline{v}_k$, the *union* of the supports of the distribution. However we continue to call $[\underline{v}, \overline{v}]$ as the support and this should cause no confusion.

[15] We need to emphasize that we assume that bidders do not know the distribution only because we think it is more realistic; however, none of our results would be affected if we had assumed that bidders do know the distribution. Bidder valuations in our model is private and correlated; the alternative assumption would have made them, in addition, conditionally independent. What is crucial is that the seller does not know the distribution.

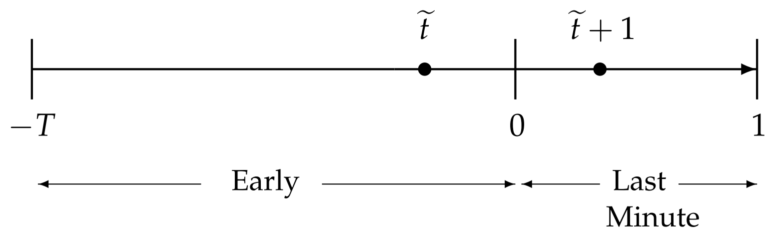


Figure 1: Bid timing and arrival. The auction starts at $-T < 0$ and ends at 1. Bidders arrive randomly over $[-T, 0]$. The arrival time of a bid made at time $\widetilde{t} \in [-T, 1]$ is uniformly distributed on the time interval $[\widetilde{t}, \widetilde{t}+1]$. Early bids arrive with certainty,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

**Timing of bids and arrivals** The auction starts at $-T < 0$ and ends at time 1. A crucial part of the model is the continuous bid arrival process. The arrival time of any bid submitted at time $t$ is uniformly distributed on $[t, t+1]$, so long as $t+1 \leqslant 1$, i.e. $t \leqslant 0$. If $t > 0$, the bid gets lost[16] with probability $t$, and with probability $(1-t)$, the arrival time is now distributed uniformly over $[t, 1]$.

Note that a bid 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. 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 $t$.

Since we want to examine the optimal choice of time of bidding in equilibrium (whether the bidders bid early to ensure arrival of their bids or sacrifice some probability of arrival and bid late, or combine these in some way by incremental bidding), we assume that bidders arrive randomly over $[-T, 0]$. Therefore any bid placed at time $t > 0$ is due to strategic reasons (and not because it would not have been possible for the bidder to have bid earlier).

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

### 2.1 Strategy of bidders

A bidder arriving at time $s \in [-T, 0]$ can bid one or more (finite number of) bids over time $t \in [s, 0]$. Further, a bidder can also submit a bid at some point inside the "last minute," i.e. at some point $q \in (0, 1)$ which then reaches with probability $(1-q)$. Note 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.

For any $t \in [-T, 1]$, let $h_t$ denote the public history of auction prices up to (but not including) time $t$. When the first bid above the reserve price arrives, the reserve price becomes active.[17] Note that the public history $h_t$ is thus a step function over the interval $[-T, t)$.

Upon arriving at the auction at $s$, a bidder can observe $h_t$ for all $t \geqslant s$. At any $t \geqslant s$ bidder also observes own valuation, arrival time, as well as the history of own bids up to $t$. These along with the public history $h_t$ form a bidder's private history at $t$.

At every instant $t$, a bidder's feasible set of actions is to either remain inactive or be active and submit a bid higher than the current auction price. A strategy of a bidder is a sequence of maps, for each time $t$, from the bidder's private history to the set of feasible actions.

Note that types below $R_0$ never win. Further, since type $R_0$ necessarily gets a 0 payoff, it is indifferent across all bid-times. Since we cannot possibly impose any equilibrium restrictions on the bid-time choice of type $R_0$, we simply break indifference of type $R_0$ in favor of non-participation so that only types $v > R_0$ (types with a strictly positive expected surplus) participate, and derive results about the bidding behavior of participating types $(R_0, \overline{v}]$.

[17] In some auctions, this may not be the case. The first activity that is registered in that case is when the second bid above the reserve price arrives. We assume the other variation as 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.

## 2.2 Strategy of the seller

Next, consider the strategy of the seller. The seller starts with a reserve price of $R_0$. This is the seller's optimal reserve price given the prior information on the distribution from which bidders draw values.

The starting reserve price is part of the stated mechanism. The strategy of the seller refers to submission of shill bids. We assume upfront that the seller submits bids through a finite number of agents and the number of shill bids that can be submitted over the time interval $[0, 1]$ is finite.[18]

In general, the strategy of the seller is similar to that of any buyer. At any time $t$ the seller observes the public history $h_t$ defined above. The seller also observes the history of own bids up to $t$. These form the seller's private history at $t$.

At every instant $t$, a seller's feasible set of actions is to either remain inactive or be active and submit a shill bid higher than the current auction price. A strategy of the seller is a sequence of maps, for each time $t$, from the seller's private history to the set of feasible actions.

Note that history at time $t$ consists of a set of discrete instants of time at which either the auction became active (reserve price was met) or the auction price jumped. At all other times the auction was inactive. Proposition 1 below shows that the seller optimally submits shill bids only when some activity occurs (either reserve met or auction price jumps). This result is useful in understanding how bidders behave in equilibrium.

**Proposition 1.** *Let any time t be called "active" if either the reserve price is met at t or the price in the auction changes at t. All other times are called "inactive." In any equilibrium, the seller submits shill bids only at active times.*

**Proof:** Suppose there is an equilibrium in which the seller's strategy involves submitting a shill bid $S$ at some time $t$ even if no bids arrive until $t$ (so that the reserve price is not met at $t$, i.e. there is no discernible activity in the auction). Let $R_1$ be the first shill bid the seller makes when the reserve price is met. We consider two separate cases, $S = R_1$ and $S > R_1$.

[18] The seller presumably operates through a number of agents - submitting several bids through multiple accounts from the same IP address can be detected by auction platforms.

Consider first the case when $S > R_1$. For this to be the seller's optimal response at time $t$, it must be that given no bids arrive until time $t$, the seller's belief puts higher weight on higher valuation distributions compared to the initial belief.[19] We now show that such beliefs cannot be part of any equilibrium.

For such beliefs to be part of an equilibrium, they must be consistent with the bidders' strategies. Specifically, it must be true that in equilibrium lower valuation bidders bid earlier than higher valuation bidders so that non-arrival of any bid by time $t$ signals a higher probability of higher values compared to the prior.

However, in that case the strategies of the higher types are not best responses. Note that since bidders arrive randomly up to time 0, for any $t \in [-T, 1)$ a history with no bids arriving until $t$ can occur with strictly positive probability.[20] A profitable deviation for higher types is to pool with lower types, since, by following the strategy of the lower types the higher types can reduce the probability of facing the shill bid $S$. This constitutes a profitable deviation from their supposed equilibrium strategies. Therefore the seller bidding $S > R_1$ when no bids have arrived cannot be part of an equilibrium.

Consider now the situation where no bid has arrived till time $t$ (and hence the reserve price has not been met) and the seller chooses to submit a shill bid $S = R_1$. For this to be optimal, the seller's updated beliefs must be the same upon observing arrival of a bid as also upon observing non-arrival of any bid .[21] However, if it is optimal to submit shill bid $S$ – in effect, increase the reserve price from $R_0$ – irrespective of history the seller observes, then that contradicts the fact that the original reserve price $R_0$ was optimally chosen.

The argument that the seller does not shill bid at any inactive time is similar and we skip the details. Essentially, if it is optimal for the seller to submit a shill bid at any inactive time $t$, that would imply either that the strategies of some bidder types are not best responses or that the action chosen by the seller at the last active time before time $t$ was not optimally chosen.‖

[19]In other words, the seller attaches a higher probability to higher values being drawn compared to the prior belief.

[20]No bidder with a value above $R_0$ might arrive. This happens with strictly positive probability. Further, for any $t < 1$, any bidder with type above $R_0$ might arrive in $(t-1, 0)$ and so even if such a type bids immediately after arrival the bid does not reach by $t$ with strictly positive probability.

[21]Note that for the seller's updated beliefs to be the same for these two events, it is a necessary condition that the high and the low bidders do not pool.

The result shows that the seller submits shill bids only at active times, and for the purpose of shill bidding we can ignore inactive times. Note that the seller's strategy has a "Markovian" property that comes naturally from repeated round of Bayesian updating: the mapping from seller's private history to actions could alternatively be characterized as mapping from the (cartesian product of the) seller's updated beliefs and current price to actions. The updated belief incorporates all the useful information from past price changes and past shill bids.

## 3 Monotonicity

For some of the results, we need to introduce a further restriction on strategies. This is introduced below.

**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 $v' \leqslant v$ at time $t$ if the auction price at $t$ is $p < v'$, then the bidder also submits the bid for any other auction price $p' \in (p, v']$.*

This essentially rules out strategies that involve bidding at some point of time $t$ if price does not exceed $p$, but refraining from bidding if it does exceed $p$. Monotonicity has the following useful implication. Suppose in some proposed equilibrium, a bidder (say bidder 1) is supposed to submit a bid at time $t$. Suppose bidder 1 deviates and submits the bid at $t' > t$. Monotonicity implies that this deviation cannot strictly increase the chance of a bid (by some other bidder) being triggered at any future point of time. The later bid 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. Assuming monotonicity gives us the following results .

**Proposition 2.** *Suppose bidders other than 1 use monotonic strategies. In any equilibrium, for bidder 1 of any type $v > R_0$, bidding at $t < 0$ is suboptimal.*

**Proof:** From Proposition 1 we know that in any equilibrium the seller's strategy involves submitting a shill bid only 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 bidders submit bids before time 0. Suitably rename bidders so that bidder 1 of type $v > R_0$ is among these, and submits a serious bid of $v'$ (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 discussed above, 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.||

The result shows that if there is any equilibrium in which some types bid early, it must be that bidders use non-monotonic strategies. As discussed in the introduction, this essentially involves using strategies that threaten to bid early if the auction price is lower than some threshold at some point $t < 0$ (this is a threat because early bids trigger shill bids early with greater probability) but delay bidding if auction price crosses the threshold at $t$. In effect others must say to each early-bidding-type: bid early or face a higher probability that we will bid early. Such threats would work so long the shill bidding consequence of the punishment is worse than that from conforming. Note that, as discussed before, the seller cannot do anything to induce early bidding. The bidders themselves must use non-monotonic strategies to threaten each other to sustain early bidding. Once we rule out non-monotonic strategies, such threats are removed, which is sufficient to remove the possibility of early bidding.

Let us now show that if others use monotonic strategies, a bidder's best response cannot involve incremental bidding. Proposition 2 rules out bidding before time 0 in this case. However, this still leaves open the possibility that a bidder submits a bid at time 0 and another inside the "last minute." So a bidder with value $v$ (say) can submit a bid $v_1$ at time 0 and $v_2$ at time $q \in (0, 1)$ where $v_1 < v_2 \leqslant v$. The next result shows that such incremental bidding is suboptimal. The bidder should bid either only at 0 or only at some $q \in (0, 1)$.

**Proposition 3.** *Suppose bidders other than 1 use monotonic strategies. In any equilibrium, for bidder 1 of 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 formal proof is in the appendix. The idea is quite simple: 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, the second bid arrives before the first bid with strictly positive probability, 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. 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.

## 4 The main results

First, we rule out the possibility of an equilibrium in which bidders wait beyond time 0 to submit bids. This implies that in any equilibrium, bidders do not sacrifice probability-of-bid-arrival in order to snipe the shill bidder. Further, standard weak dominance arguments imply that a bidder must submit a truthful bid (i.e., a bid equal to true valuation). It follows then that despite the presence of shill bidding, the auction is efficient "at the top": whenever the object is allocated to a genuine bidder, it goes to the highest value bidder. Further, since the shill bids might not arrive, the auction might even have greater efficiency "at the bottom": fewer types may be excluded compared to the auction without shill bidding. It follows that under private values, the case against shill-bidding is weaker than one might expect.

**Proposition 4.** *There is no equilibrium in which any type of any bidder bids after 0.*

**Proof:** 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 $t'$.[22] Let $\pi(s)$ be the expected payoff conditional on the bid reaching at $s$.[23] Also, note that a bid at $t$ arrives with probability $(1-t)$. 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.[24] In that case, the worst possible punishment is that other bidders all bid $\overline{v}$ at $s$.[25] Since bidder 1 bids at most own value $v$, the expected payoff $\pi(s)$ cannot be negative for any value of $s$. 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 $\overline{v}$ 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.‖

---

[22] Bidding 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.

[23] The 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.

[24] For 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.

[25] Such a punishment might not be credible, but we are simply showing that even under the worst possible punishment the payoff exceeds zero. Therefore, for any other punishment the payoff exceeds zero as well.

We now prove the main results of the paper.

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

To prove this, we consider the best bid-time of bidder 1 of type $v$ when some subset of other bidders arrive. As we show, the optimal bid-time is invariant across all such subsets and therefore the result does not rely on any knowledge of the precise number of bidders arriving in the auction by any bidder or the seller.

**Proof:** Consider the following bidder 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 valuation $v$ submits a bid equal to $v$ if the auction price at time 0 is less than $v$; otherwise the bidder remains inactive. For any $t \in (0, 1]$, any bidder with valuation $v$ remain inactive if the history of the bidder is such that the bidder has submitted bid equal to $v$ 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 $v$.

It is clear that if the above is an equilibrium, the resulting outcome would be that all arriving bidders would bid for the first time their value of the object at time $t = 0$. Hence, the remaining task is to check that the above is indeed an equilibrium. It is obvious that deviating and bidding at some time $t < 0$ is not a profitable deviation. Since the strategies are monotonic by construction, this follows directly from Proposition 2. Consider now a deviation where a bidder submits an additional bid at some $t > 0$. Again, since the strategies are monotonic by construction, it follows from Proposition 3 that this is unprofitable. Finally, It follows from Proposition 4 that deviating and bidding at some time $t > 0$ is unprofitable. This completes the proof.||

Next, we consider uniqueness. Let us refer to the time at which bids are submitted in equilibrium as the "equilibrium outcome." We now show the following.

**Theorem 2.** *If we restrict attention to monotonic strategies, all bidders of all types above $R_0$ bidding at time 0 is the unique equilibrium outcome.*

**Proof:** Proposition 2 rules out bid times before time 0. From Proposition 3, we know that bidders submit truthful bids once either at 0 or at some point of time in $(0, 1)$. Finally, Proposition 4 rules out the latter. Therefore all types above $R_0$ of all bidders bidding at time 0 is the unique equilibrium outcome.‖

The general properties below follow immediately from the results above.

**Proposition 5.** *Any equilibrium that involves bidding earlier than time 0 must involve non-monotonic strategies by some types (above $R_0$) of some bidders. For the bidders, any such equilibrium is Pareto dominated compared to an equilibrium where all bidders of all types (above $R_0$) bid at time 0.*

**Proof:** The first part follows directly from Proposition 4. Next, suppose there is an equilibrium in non-monotonic strategies in which some types above $R_0$ of some bidders bid before time 0 in equilibrium. This results in the seller's shill bids being triggered earlier than an equilibrium in which all types above $R_0$ bid at time 0. Since all bids at time 0 or earlier reach with certainty, bidding before 0 confers no benefits to the bidders, but since shill bids arrive with strictly higher probability when bids reach early, the equilibrium in non-monotonic strategies generates a strictly lower expected payoff for all serious types of all bidders. This completes the proof.‖

## 5 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 bidding behavior that does not rely on any discontinuity of the bid arrival process. We allow a continuous choice of bid times and a continuous arrival process for submitted bids. The framework can also be useful for analyzing online auctions under common values or other richer valuation environments. Indeed, some of the results in the literature that depend on an assumption of discontinuity unravel in a continuous framework.

In our model, bidders' private values are correlated, implying that shill bidding is useful since it essentially allows the seller to adjust the reserve price mid-auction.[26] As we show,

[26] In our model shill bids are used to update reserve prices. As noted before, under independent private values, shill bids are irrelevant and there is no reason for bidders to delay their bids either. Other reasons for late bidding can arise even under independent private values drawn from a known distribution if bids can

the seller can only react when the auction registers activity. This in turn implies that bidders have an incentive to delay events that lead to the seller learning about the value distribution, reducing the chance of successful shill bids. In other words, bidders bid late not because they want to snipe other bidders, but the shill bidder.

Our result shows that all bidders bidding at the last point of time at which their bid still reaches with probability 1 constitutes an equilibrium. The equilibrium is Pareto dominant for bidders. Further, we show that the equilibrium bid time is unique if we restrict attention to monotonic strategies. Unlike much of the literature, knowledge of actual number of bidders is immaterial for our results. We can therefore allow for random entry, so that neither bidders nor the seller know the precise number of bidders - a setting natural for online auctions.

As noted 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 reduce the seller's payoff in equilibrium by impeding information aggregation. However, our results suggest that the case against shill bidding is much weaker under private values. We show that there are no equilibria that involve bidding later than time 0. An implication of this is that in all equilibria, when the object is sold to a genuine bidder, it goes to the bidder with the highest value. Therefore there is no efficiency loss "at the top." Further, the option of shill bidding might lead the seller to post a lower initial reserve price compared to the case where such an option is not available.[27] Subsequently, given sniping by bidders, some or all updates might not arrive, making the auction with shill bids more efficient ex post compared to auctions without shill bidding.

The auction we consider has a fixed end time. We can also consider efficiency comparisons with auctions that have a flexible end time. In the latter, all shill bids necessarily

---

be cancelled after submission. For example the seller might attempt the following: submit a high shill bid to discover the highest genuine bid, cancel the high shill bid and submit instead a shill bid just below the highest genuine bid (essentially trying to extract the first price in a second-price auction). In such situations, it is reasonable to suppose that the bidders will, in turn, delay submitting their bids in order to delay value-discovery and snipe the seller's shill bids. While a full blown analysis of such alternative motives for late bidding is beyond the scope of the paper, it is possible that the resulting outcome would be similar to the equilibrium outcome from our model, namely that the bidders would delay bid-submission to lower the chance of successful shill bids, but not delay to an extent that would involve sacrificing any probability of their own bid arriving before the auction closes.

[27] Such cases are likely when, for example, the prior over distributions puts sufficient weight on both lower (in stochastic dominance ranking) and higher distributions.

reach. On the other hand, the initial reserve price would be lower than that under a fixed end-time auction. Thus there are two conflicting effects. A flexible end-time auction offers a lower initial reserve price but all shill bids reach, so it is possible that the seller "overshoots," i.e. submits a shill bid that is too high relative to the optimal reserve price if the true distribution were known. A fixed-end time auction may have a higher reserve price, but the shill bids might not reach, so it suffers less from the overshooting problem. One might conjecture that given distributions ranked by first-order stochastic dominance, if the bidders are from a distribution with low mean, they would gain from a flexible end-time auction relative to a fixed end-time one, as the lower initial reserve price effect is likely to outweigh the overshooting effect. Similarly, bidders from high mean distributions are likely to prefer fixed end-time auctions. The highest mean distribution is likely to face a higher expected reserve price under a flexible end-time auction compared to a fixed end-time one.

Which auction is better for the seller? While this is beyond the scope of our analysis, we believe that the key to this lies in entry incentives. Since, ceteris paribus, buyers from a higher distribution are likely to obtain more surplus from auctions with a fixed end-time, such an auction would be more attractive to them, and less attractive to the seller. However, with endogenous buyer entry, buyers from a high distribution might self-select into a fixed-end-time auction, and then it is not clear that a flexible end-time auction necessarily generates more expected revenue for the seller. Studying competition amongst online auction sites using different auction mechanisms in environments characterized by endogenous entry by buyers and sellers should be an interesting and important topic for future research.

# Appendix

## A.1 Proof of Proposition 3

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 $v_1 < v_2 \leqslant v$.

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, 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 $q$:

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

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).

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.

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

**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

$$EP_q(v_1, v_2) \leqslant EP_q(v_1, v). \tag{A.3}$$

**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),

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

Next, let

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

Note that as $v_1 \to v$, $\alpha \to 1$. 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

$$\overline{\pi} = \Pr(v \text{ wins})\Big[\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v)\Big] \tag{A.4}$$

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, $\alpha$ 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:

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

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 $\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \geqslant \max[P_0(v), EP_q(v)]$, or writing it out, we have,

$$\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \geqslant P_0(v), \tag{A.6}$$

$$\alpha(v_1)P_0(v) + (1 - \alpha(v_1))EP_q(v_1, v) \geqslant EP_q(v). \tag{A.7}$$

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

$$P_0(v) \leqslant EP_q(v_1, v). \tag{A.8}$$

From (A.7), we have

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

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.$\|$

## References

**Bag, Parimal Kanti, Ruqu Wang, and Emin Murat Dinlersoz**, "More on phantom bidding," *Economic Theory*, 2000, *15* (3), 701–707. 8

**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, 7

**Chakraborty, Indranil and Georgia Kosmopoulou**, "Auctions with shill bidding," *Economic Theory*, 2004, *24*, 271–287. 9

**Deltas, George**, "When Does Cheating on Mail-in Bids Pay ? A Guide for the Dishonest Auctioneer," *Journal of Economic Behavior and Organization*, 1999, *40*, 313–323. 8

**Engelberg, Joseph and Jared Williams**, "eBays proxy bidding: A license to shill," *Journal of Economic Behavior and Organisation*, 2009, *72*, 509–526. 8

**Graham, Daniel A., Robert C. Marshall, and Jean-Francois Richard**, "Phantom bidding against heterogeneous bidders," *Economics Letters*, 1990, *32* (1), 13–17. 8

**Kosmopoulou, Georgia**, "Bidding off the wall in common value auctions," *International Journal of Economic Theory*, 2011, *7* (2), 217–230. 9

**Lamy, Laurent**, "The Shill Bidding Effect versus the Linkage Principle," *Journal of Economic Theory*, 2009, *144*, 390–413. 9

**Myerson, Roger**, "Optimal Auctions," *Mathematics of Operations Research*, 1981, *6*, 58–63. 3

**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, 6, 8

**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). 8

**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, 7

**The Sunday Times**, "Revealed: how eBay sellers fix auctions," 2007. January 28. 3

**Walton, Kenneth**, *Fake: Forgery, Lies, & eBay*, Gallery Books, 2006. 3

**Whitworth, Dan**, "Man fined over fake eBay auctions (BBC Newsbeat)," 2010. July 5. 3

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