Department of
Economics

University of Pittsburgh
Kenneth P. Dietrich School of Arts and Sciences

WORKING PAPER SERIES

16/015

Sniping in Proxy Auctions with Deadlines

Adam Kapor and Sofia Moroni

May 8, 2016

University of Pittsburgh • Dietrich School of Arts and Sciences • Department of Economics
4901 Wesley W. Posvar Hall • 230 South Bouquet Street • Pittsburgh, PA 15260

Sniping in Proxy Auctions with Deadlines

Adam Kapor and Sofia Moroni*

May 8, 2016

Abstract

In most online auctions with deadlines, bidders submit multiple bids and wait until late in the auction to submit high bids, a practice popularly known as "sniping." This paper shows that whenever there is a nonzero probability that an auction contains "shill bidders", who attempt to raise the sale price without winning, equilibrium play must exhibit sniping. We present a model of online auctions with stochastic bidding opportunities and incomplete information regarding competing players' valuations. We characterize perfect Bayesian equilibria under a trembling hand refinement, allowing for a positive probability that a player is a bidder who plays a fixed strategy that raises the sale price. In doing so, we develop a one-shot deviation principle for a class of continuous-time games with stochastic opportunities to move. We find that in all equilibria, players wait until a late time threshold to place their bids, even as the probability that a shill bidder is present becomes arbitrarily small. Using data from eBay auctions, we show that observed behavior is consistent with equilibrium play and that comparative statics of the timing of bids match the model's predictions. In contrast, when there is no possibility of a shill bidder, the unique equilibrium outcome is that players bid their valuation as soon as a bidding opportunity arrives. The results extend to other types of heuristic bidders who place incremental bids.

*Princeton University and University of Pittsburgh. akapor@princeton.edu, moroni@pitt.edu. We are grateful to Dirk Bergemann, Johannes Hörner, and Larry Samuelson for their insightful comments. We would also like to thank the seminar audiences at Yale, Minnesota, and the Game Theory conference at Stony Brook.

1 Introduction

Many goods are sold via internet auctions with proxy bidding and deadlines. Although these auctions allocate the good to the highest bidder at the second highest bidder's maximum bid, players do not behave like participants in a standard second-price auction. Empirically, on sites such as eBay it is common for bidders to "snipe", that is, to submit high bids at the last instant before the auction ends. Bajari and Hortaçsu (2003) document that the median winning bid arrives after 98.3% of the time in an auction has elapsed. Roth and Ockenfels (2002, 2003) describe similar patterns.[1] As a result of sniping, participants may fail to submit bids before the deadline, and therefore auction outcomes may be inefficient.[2] At the same time, bidders have reason to think that a small but nonzero share of bids in are submitted by "shill bidders" who are colluded with the seller.[3]

The main result of this paper is that the possibility of shill bidders leads players to engage in sniping. We present a model of online private-value auctions with a proxy bidding system and a deadline in which players get stochastic opportunities to place a bid. We allow for a small probability that one bidder is privately a shill bidder who raises the sale price without winning the auction.[4] We characterize the unique equilibrium strategies under a trembling hand refinement. The presence of incomplete information about players' types changes equilibrium outcomes even as the probability of a shill bidder tends to zero.[5] Equilibrium play in the possible presence of shill bidders involves late submission of bids. This equilibrium play is in line with commonly observed behavior in the online auctions market. In contrast, when players' valuations are the only source of incomplete information, we show that players place

[1] The distribution of bids is bimodal, with a small peak at the beginning and many bids at the end of the auction (Bajari and Hortacsu (2003) and Roth and Ockenfels (2002)). If players submit their true valuations, they should bid at most once, but Roth and Ockenfels (2006) document that only 62% of players bid only once, with the mean number of bids per player being 1.89.

[2] In addition, Backus et al. (2015) show that players are less likely to bid in future eBay auctions if they lose to a sniping bid.

[3] Engelberg and Williams (2009) estimate that 1.4% of bids on eBay are submitted by shill bidders who are attempting to push up the price without winning the auction.

[4] We model a shill bidder as playing a discover-and-stop strategy as described in Engelberg and Williams (2009).The shill bidder bids one increment at each opportunity until his proxy is at most one increment below the highest proxy. Because of the way the price updates when bids are less than an increment below the highest bid, the shill-bidder can infer the highest bid with high probability and avoid outbidding it (see footnote 18 for an example).

[5] For this argument, the probability of trembles must vanish faster than the probability of heuristic bidders. An interpretation is that, when a bidder observes an unexpected action that is consistent with heuristic bidding, he assumes that it is more likely to have come from a heuristic bidder than from a mistake. Our order of limits implicitly imposes this assumption.

a bid near their valuations as soon as possible.

Our model captures the two key features of online auctions that distinguish them from ascending auctions and from second price sealed bid auctions. First, on sites such as eBay, auctions last for several days until a deadline. Given the auctions' extended duration, players cannot pay full attention for the entirety of each auction. Furthermore, they may not be able to check on the auction at will. For example, players may find the time or remember to check on their auction at times which they are not able to predict ex-ante. When a player attempts to submit a bid in the final seconds of an auction, internet traffic may prevent his bid from arriving in time. We model bidding opportunities as arriving according to an independent exogenous Poisson process. In the model, a player can place at most one bid at each bidding opportunity. Bids can only be updated upwards and cannot be withdrawn.

Second, we model the proxy bidding system that is used in online auctions. On eBay, when a bidder acts, he places a proxy bid which is not observed by competing bidders. The system automatically outbids players who place proxies below the highest proxy. In the model, as in practice, when a player has an opportunity to act, he observes the value of the second-highest proxy at the time of his bid, and chooses whether, and at what value, to set his proxy. At the deadline, the player with the highest proxy wins the object, and pays an increment above the value of the second-highest proxy.

In a setting in which there is no possibility of a shill bidder, in contrast, we find that under a "trembling hand" refinement the unique equilibrium outcome is that players submit the highest proxy that guarantees a positive payoff conditional on winning, and they do so at the first bidding opportunity. The trembling hand refinement is necessary to obtain unique predictions. When we consider the weaker equilibrium notion of Perfect Bayesian Equilibrium we find many equilibria that rely on one player being indifferent among bids. Some of these equilibria involve implicit collusion by the players, giving a high payoff to the players and as a consequence a small payoff to the seller.[6]

Additionally, we show that the equilibria must exhibit sniping in the presence of a more general class of commitment types that make incremental bids and may outbid the highest bidder. There must be sniping, for example, if there is a bidder who always outbids the current price by an increment.

In order to solve the model, we establish a one shot deviation principle for a class of continuous-time games with stochastic opportunities to move. This technical result simplifies

[6]This equilibria are similar to the ones characterized in Ambrus et al. (2014) and Ockenfels and Roth (2006).

finding equilibria in this class of games, and may be useful for other models with stochastic opportunities to move.

We show that empirical patterns in online auctions are consistent with the model's predictions. Using data on eBay auctions for a video game, a product which is likely to have private values and is relatively homogeneous across auctions, we show that bidding times exhibit a bimodal distribution consistent with the model, and that incremental bids are more likely to occur early in auctions. Moreover, we find that bids by a player "trigger" bidding by the other players in a manner consistent with our model's predictions, and that the comparative statics of the timing of bids with respect to prices and revealed valuations are consistent with the model.

Our paper shows that the belief that other players may not behave as truthful rational bidders can explain empirical features of online auctions. We follow a literature that explains behavior in online auctions. Roth and Ockenfels (2002, 2003, 2006) argue that sniping might arise from "implicitly collusive" equilibria by the bidders. They also argue that sniping might be a best response against naïve incremental players, that is players who bid one increment above the current price, who do not fully understand the proxy system. Roth and Ockenfels (2002) conjecture that the possibility of facing a shill bidder could cause players to snipe in order to avoid price wars. Their model allows for many possible equilibria. Our model gives a *unique prediction* that requires sniping and relates sniping to the presence of players who do not behave rationally. Uniqueness of equilibrium behavior is consistent with the prevalence of sniping in eBay auctions. Bajari and Hortaçsu (2003) provide an explanation for sniping in a common values setting. Hopenhayn and Saeedi (2015) introduce a model with stochastic bidding opportunities, exhibiting a unique equilibrium, in which players' valuations evolve stochastically as the auction proceeds. The present paper focuses on the private values case in which players do not learn about their valuation as the auction evolves.

This paper also adds to a growing literature on games in which players recieve random opportunities to take an action. Calcagno et al. (2013) and Kamada and Muto (2014) consider settings in which players have stochastic opportunities to publicly change their actions before playing a normal-form game. Ambrus and Lu (2014) explore multilateral bargaining with a deadline. Moroni (2014) proves existence of a trembling hand equilibrium in games in which players receive stochastic opportunities that arrive at a Poisson rate, allowing for incomplete information, and for action spaces that depend on the state, as in this paper. Lovo and Tomala (2015) prove existence of markov-perfect equilibria in revision games with stochastic state transitions and complete information.

Our model is closest to Ambrus et al. (2014), who also present a model of online auctions in which bidders receive stochastic opportunities to place their bids.Ambrus et al. do not allow for the possibility of heuristic bidders who play a fixed strategy, but they allow each player to place arbitrarily many bids at each bidding opportunity, rather than just one. Under our assumptions, when there is half a second left in the auction, a player will be much more likely to be able to place only one bid than to be able to place a first bid and, depending on the price update, have enough time to place a second one. This difference has consequences for equilibrium outcomes when all players are rational. When there are no heuristic bidders, in our model there is a unique trembling hand perfect equilibrium in which all players place a truthful bid at the first opportunity. When players can place multiple bids at each opportunity–as in Ambrus et al. (2014)–there exist trembling hand perfect equilibria in which players bid incrementally late in the auction. In their model, because players can place more than one bid at an opportunity they are better able to punish a player that has not bid as expected. The ability to punish makes it possible to sustain equilibria that feature incremental bidding. When players are not rational with probability one and there is any arbitrarily small probability that a player is a "commitment type", our model predicts that all equilibria exhibit sniping.

Finally, our paper builds on a large body of work on "reputational" or "commitment type" players in games, beginning with seminal work by Kreps and Wilson (1982). Mailath and Samuelson (2006) provide a textbook treatment of this literature. In a different application, Abreu and Gul (2000) consider bilateral bargaining in the possible presence of commitment types.

2 Model

A single object is offered for sale in an auction. The auction has length T. There are n players. Time is continuous, but a bidder can place a bid at time $t \in [0,T]$ only if he gets a bidding opportunity at that time. Bidding opportunities arrive at a common rate λ independently across players. Time counts down from T at the start so that time 0 corresponds to the end of the auction.

When a bidding opportunity arrives a player observes the *current price* p and can choose to place a bid or pass up the bidding opportunity. Bids can take values in a finite grid $B =$

$\{m\delta\}_{m=0}^{M}$. If a player places a bid, it must be greater than the current price by at least an increment Δ. That is, a bid b placed by a player must satisfy $b \geq p+\Delta$, where Δ is a multiple of δ. The initial bids of all players are normalized to zero and the initial price is equal to the reserve price. When a player places a bid b the current price p updates as follows: if the bid placed by the player is above the highest current bid placed by any other player, b^H, the current price updates to $p = b^H + \Delta$. If b is less than b^H, the current price updates to $p = b+\Delta$. The winner of the auction is the player who has the highest bid when the auction ends, at time zero, and the price paid by the winner is the current price at that time.[7,8] Thus, at every time there is at most one bid above the current price. There is no discounting.

There are two types of players: commitment type players and payoff type players. Commitment type players are non-strategic and play fixed strategies which we describe in section 5. A payoff type player has a private valuation v for the object. Valuations are independent draws from a probability distribution $q(\cdot)$ with full support in $V \subseteq B$. The price paid by the player who wins the auction is denoted p^F. It is equal to the price at time zero. Payoff type players are risk neutral and the winning player obtains payoff $v - p^F$. In what follows we discuss payoff type players.

A player's information consists of his value for the object, the history of his own bid and the history of the price as well as the bidding opportunities that he has had. In particular, a player neither observes when other players get bidding opportunities nor what the highest bid is.

A complete history up to time t consists of the history of all bids up to period t as well as the times of all opportunities. A player's private history consists of times in which bids were placed, the identity of the players that placed the bids, the price updates and the identity of the player who held the highest proxy at each one of these times. Formally, suppose at time t there have been k bids placed. Let $\{t_m\}_{m=1}^{k}$ denote the times of bids placed by players $\{j_m\}_{m=1}^{k}$ resulting in price updates $\{p_m\}_{m=1}^{k}$ with the identity of the players who hold the winning bid given by $\{j_m^w\}_{m=1}^{k}$ up to time t. Let $\{t_{i,m}\}_{m=1}^{r}$ denote the times of i's bidding opportunities that occur before t and $\{b_{i,m}\}_{m=1}^{r}$ i's proxy after each of these opportunities. A history for player i is a vector $h_i^t = \left(\left((t_1, j_1, p_1, j_m^w), \ldots, (t_k, j_k, p_k, j_k^w)\right), (\{t_{i,m}\}_{m=1}^{r}, \{b_{i,m}\}_{m=1}^{r})\right)$. Let $\mathscr{H}_i^{t,(k,r)}$ be the set of histories with k bids and r player i arrivals. Let $\mathscr{H}_i^t = \bigcup_{k,r\geq 0} \mathscr{H}_i^{t,(k,r)}$ denote the set of player i's private histories up to time t. Let $\bar{\mathscr{H}}_i^t$ denote the set of private histories of

[7] In online auctions such as e-Bay the current price updates to $\max\{b^H + \Delta, b\}$. Our assumption simplifies the analysis.

[8] If the two highest bids are the same we assume that the player who placed the bid earlier wins the good.

player i in which i has a bidding opportunity at time t. The private history process defines a filtration $\{\mathscr{F}_i^t\}_{t\in[0,T]}$ for each player, where $\mathscr{F}_i^t$ denotes the σ-algebra induced by player i's information $V \times \mathscr{H}_i^t$. A behavior strategy σ_i for player i is a process adapted to $\{\mathscr{F}_i^t\}_{t\in[0,T]}$ mapping private histories in $\bar{\mathscr{H}}_i^t$ into probability distributions over $\{b \in B | b \geq p+\Delta\} \cup b^{\emptyset}$, where $b^{\emptyset}$ denotes not placing a bid. In particular, the process must be measurable with respect to time.

At time t each player has a belief over the complete histories as well as the valuations of other players given his private history. Let μ_i^t denote the process of the beliefs of player i. More precisely, μ_i^t is a process of probability distributions over $\mathscr{H}^t \equiv \times_{j\neq i} V \times \mathscr{H}_j^t$ that is adapted to the filtration $\mathscr{F}_i^t$. Beliefs about the unobservable proxy and types of other players are updated according to Bayes' rule.[9] Let $\mu = (\mu_i^t)_{i\in\{1,\dots,n\},t}$ denote the belief system of all players and let $\sigma = (\sigma_1, \dots \sigma_n)$ denote the strategy profile.

Let $\mathbb{E}_{\mu_i^t,\sigma}\left(v_i - p^H | i \text{ wins }, h\right)$ denote i's expected payoff conditional on winning, given that i's belief at time t is μ_i^t, the history at time is $h \in \mathscr{H}_i^t$ and players play according to the strategy profile σ. Let $\mathbb{P}_{\mu_i^t,\sigma}(i \text{ wins}, h)$ be the probability that i wins the auction, given the belief, private history and the strategy profile. Player i's expected payoff at history $h \in \bar{\mathscr{H}}_i^t$, given the assessment (σ, μ) is

$$U_i\left((\sigma,\mu),h\right) = \mathbb{E}_{\mu_i^t,\sigma^*}(v_i - p^H | i \text{ wins}, h)\mathbb{P}_{\mu_i^t,\sigma^*}(i \text{ wins}, h).$$

We next define a Perfect Bayesian Equilibrium in a standard way.

Definition 1. A *Perfect Bayesian Equilibrium* (PBE) is defined as a strategy profile and beliefs (σ^*, μ^*) such that

a) Beliefs are updated according to Bayes' rule.

b) For $i \in \{1,\dots,n\}$, σ_i^* is optimal given the belief system μ_i^*, that is for every history $h \in \mathscr{H}_i^t$ and for every t

$$\sigma_i^* \in \text{argmax}_{\sigma_i} U_i\left((\sigma_i, \sigma_{-i}^*, \mu^*), h\right).$$

As is often the case in auction models, there are equilibria in dominated strategies. To rule out such equilibria, we will consider a refinement in the spirit of extensive form trembling hand

[9] Details on how beliefs update according to Bayes' rule are given in the appendix.

perfection as introduced by Selten (1975).[10] For each player $i \in \{1, \dots, n\}$, let S_i^ε denote the set of totally mixed behavioral strategies of player i that put weight at least ε on every bid that is available as well as $b^\emptyset$ at every private history in $\bar{\mathscr{H}}_i^t$ for every time t. We refer to a strategy $\sigma_i^\varepsilon \in S_i^\varepsilon$as an ε-constrained strategy.

Definition 2. An *ε-constrained equilibrium* σ^ε is a strategy profile such that for each i $\sigma_i^\varepsilon \in S_i^\varepsilon$ and σ_i^ε maximizes i's utility over strategies in S_i^ε.

Next we define a trembling hand perfect equilibrium as the limit of ε-constrained equilibria as $\varepsilon \to 0$.

Definition 3. A strategy profile σ is a *trembling hand perfect equilibrium (THPE)* if there exists a sequence $(\varepsilon^m)_{m=1,2,\dots}$, with $\varepsilon^m > 0$ and $\lim_{m\to 0} \varepsilon^m = 0$, and ε^m-constrained equilibria σ^{ε^m} such that $\sigma_i^{\varepsilon^m}(h) \to \sigma_i(h)$ for every player i, history $h \in \bar{\mathscr{H}}_i^t$ and time t.

3 Preliminaries: one-shot deviation principle

In this section we show that a notion analogous to the one-shot deviation principle holds in the present stochastic moves setting. The one-shot deviation principle (OSDP) is useful to characterize PBE and THPE. In the deterministic moves setting, the OSDP tells us that, in order to verify that a strategy and belief assessment constitutes an equilibrium, we only need to check that players do not wish to deviate at any given information set while acting according to the equilibrium strategies at every other information set (See Fudenberg and Tirole (1991), Hendon et al. (1996) and Selten (1975)). Analogously, in the stochastic moves setting we show that in order to verify that an assessment is an equilibrium we only need to check that players do not have profitable deviations at any given opportunity while acting according to the equilibrium strategies at every other opportunity.

For any ε-constrained strategy we can compute each player's belief over histories using Bayes' rule. Given a ε-constrained strategy σ^ε, let $\mu(\sigma^\varepsilon)$ denote the belief system associated to it.

For a given private history $h \in \bar{\mathscr{H}}_i^t$, let $S_{i,h}$ denote the set of probability distributions over the actions that are available to player i at history h. Let $S_{i,h}^\varepsilon \subseteq S_{i,h}$ denote the set of probability distributions that put weight at least ε on all available actions at history h. Let $(\tilde{\sigma}_{i,h}, \sigma_{i,-h})$

[10]The refinement is in the spirit of extensive form trembling hand perfection. Note however that the space of histories is infinite. An analogous notion of equilibrium is introduced by Kamada and Muto (2014).

denote the strategy of player i that corresponds to $\tilde{\sigma}_{i,h}$ at history h and coincides with σ_i at all histories other than h.

Definition 4. An assessment (σ,μ) is a local best response for player i at history $h\in\bar{\mathscr{H}}_i^t$ over S if

$$\sigma_i(h)\in\operatorname{argmax}_{\tilde{\sigma}_{i,h}\in S}U_i\left(\left((\tilde{\sigma}_{i,h},\sigma_{i,-h},\sigma_{-i}),\mu\right),h\right).$$

Thus, an assessment (σ,μ) is a local best response for player i at a history h in which player i has an opportunity to move if player i's behavioral strategy at that history is optimal given that all players behave according σ after history h.

Proposition 1 (One-shot deviation principle)**.** *1. An assessment (σ,μ) is a PBE if and only if for every player i and history $h\in\bar{\mathscr{H}}_i^t$, (σ,μ) is a local best response over $S_{i,h}$.*

2. A strategy profile σ^ε is an ε-constrained equilibrium if, for every player i and history $h\in\bar{\mathscr{H}}_i^t$, the assessment $(\sigma^\varepsilon,\mu(\sigma^\varepsilon))$ is a local best response over $S^\varepsilon_{i,h}$.[11]

Proof. Assume that the strategy profile σ does not have any profitable local deviations but it is not a PBE (the argument for ε-constrained equilibrium is analogous). There must be a player i, a history $h\in\bar{\mathscr{H}}_i^t$ and a strategy profile $\tilde{\sigma}_i$ that is a profitable deviation for player i at history h. Consider the strategy $\tilde{\tilde{\sigma}}_i$ that coincides with $\tilde{\sigma}_i$ in all histories in which player i has received at most k opportunities after history h and coincides with σ after all histories in which i has received more than k opportunities. Let d be the difference in i's expected payoff at history h from following strategy $\tilde{\sigma}_i$ instead of σ_i. Because the probability of there being more than k arrivals converges to zero as k approaches infinity, we can find k large enough that $\tilde{\tilde{\sigma}}_i$ is a profitable deviation that gives player i a payoff of at least $d/2$ above the payoff of σ_i.[12] Thus, $\tilde{\tilde{\sigma}}_i$ is a profitable deviation that differs from σ_i only in i's first k opportunities after history h. Consider a history $\tilde{h}\in\bar{\mathscr{H}}_i^{t-\tau}$ that is a continuation of h and features k opportunities after h. Because $\tilde{\tilde{\sigma}}_i$ coincides with σ_i after the k'th arrival following h, the payoff from $\tilde{h}$ under $\tilde{\tilde{\sigma}}_i$ is identical to the payoff from $(\tilde{\tilde{\sigma}}_i(\tilde{h}),\sigma_{i,-\tilde{h}},\sigma_{-i})$. Now, because σ_i admits no profitable local deviations, σ_i must give a weakly higher payoff than $\tilde{\tilde{\sigma}}_i$ at $\tilde{h}$. Thus, we can assume $\tilde{\tilde{\sigma}}_i$ differs from σ_i only on histories with at most $k-1$ arrivals after h. Reasoning by induction

[11] No local deviations is sufficient but not necessary for an ε-constrained equilibrium because any strategy that differs from σ^ε in a zero measure set of histories gives the same payoff to each agent. Thus, there may be ε-constrained equilibria that are not a local best response in a zero measure set of histories.

[12] Recall that there are finitely many actions, so that payoffs are bounded.

we conclude that σ must be optimal for each player at every opportunity given the belief system.

An analogous argument establishes the result for THPE in 2. $\square$

4 Equilibrium without commitment players

In this section we consider the game in which all players are payoff types and choose their strategies knowing that all players are rational. We show that every trembling hand perfect equilibrium (THPE) is such that players bid their valuation minus an increment as soon as they get a bidding opportunity regardless of the history of play.[13] We refer to bidding the valuation minus an increment as *truthful bidding*.

The argument consists of two steps. First we show that late enough in the auction, regardless of the history of play, players will place the highest bid that guarantees a non-negative payoff. This first step is established in Lemma 1. The second step is an unravelling argument. Since play at the end of the auction does not depend on the players' actions, at a slightly earlier time, players can only gain from placing the highest profitable bid regardless of the history of play. As a result, there is no last time at which players bid less than the highest profitable bid and players must place a truthful bid as soon as they arrive at the auction (Proposition 2).

Lemma 1. *There exists $t_0(n) > 0$ such that in every ε-constrained equilibrium each player i puts weight ε on all bids other than bid $b = v_i - \Delta$ if it is above the current price, as soon as an opportunity arrives after $t_0(n)$. Therefore, in every THPE players place a truthful bid at the first opportunity after time $t_0(n)$.*

The proof of Lemma 1 is given in the appendix, and shows that a lower bound for the payoff of truthful bidding exceeds an upper bound for the payoff of any other bid. Let us give an intuition for the results. Late enough in the auction, a player who gets an opportunity might not get a subsequent chance to outbid the current price. Thus if there is a positive probability that the unobservable highest bid is below the player's valuation—which is always true in an ε-constrained equilibrium—it is better to place a bid above the current price than to pass up a bidding opportunity. Furthermore, no sequence of rewards and punishments could

[13]Note that because of the price update assumption the valuation minus the increment is the highest profitable bid for each player.

incentivize a player to place a bid below his truthful bid. For instance, suppose a player is punished for bidding higher than expected. He can only be punished once it is revealed that he has deviated. Moreover, a bid is only revealed once it is outbid. Thus, in order to punish, opposing players need at least two opportunities: one to discover the deviation, and a second to implement a punishment. Sufficiently late in the auction, the probability of opponents receiving two opportunities to act is vanishingly small. Thus, a player cannot be dissuaded from placing a truthful bid.[14]

Our next result proves that all trembling hand perfect equilibria must have the players bid truthfully at the first opportunity. The result is stated in the following Proposition.

Proposition 2. *In every trembling hand perfect equilibrium, players bid their valuation minus an increment, if it is above the current price, as soon as a bidding opportunity arrives.*

The proof is provided in the appendix, and proceeds by contradiction. We provide the following intuition for the argument. From Lemma 1, there is a time threshold after which players bid truthfully at the first opportunity regardless of history. We show that there cannot be an earliest such time. Hence, players must bid truthfully from the beginning of the game. We establish that if there were such a time, a player who receives an opportunity slightly earlier than this earliest time would have incentives to bid truthfully as well. To show this "unraveling" result, we note that in order to be dissuaded from bidding truthfully, a player needs to expect other players to behave differently depending on his bid. If truthful bidding starts τ seconds after a player's arrival, then any punishments and rewards for not outbidding can only occur in in a time window of length τ. For small τ, it is unlikely that players would have an opportunity to detect a deviation, and a subsequent opportunity to implement a punishment, in a window of length τ. Hence for small τ it is unlikely that any punishment can be implemented before the threshold. Thus, every player prefers to bid truthfully τ seconds before the supposed earliest time.[15]

[14]This argument illustrates the key difference between the present paper and Ambrus et al. (2014). In their setting, because players can place more than one bid at an opportunity, there is no loss from placing a small bid. If the opponent had trembled and placed a higher bid than expected, a player can bid incrementally until obtaining the highest bid. Further, players are better able to punish deviations, as they can discover a deviation and punish it at the same opportunity.

[15]The arguments in Lemma 1 and Proposition 2 do not rely on the players having the same arrival rate. If they were to have heterogenous rates of arrival the unique equilibrium play would also involve truthful bidding at the first opportunity.

5 Equilibrium with commitment types

In this section we assume that there is a small probability that a player is non-strategic. We refer to the non-strategic players as "commitment types" and we assume that they play a fixed heuristic strategy. In most of the following results we will assume that the commitment type is a shill bidder[16] who plays a discover-and-stop strategy as described by Engelberg and Williams (2009). The shill bidder commitment type bids one increment at the time until his bid is less than one increment below the highest bid. In section 5.3 we consider more general commitment types who also bid incrementally. Bidders who bid one increment above the current price whenever they are outbid or shill bidders that sometimes mistakenly outbid the highest bid are included in this more general set of commitment types.

The introduction of commitment types is inspired by empirical observations made by several authors. Roth and Ockenfels (2002) find in a survey of eBay participants that 10% of respondents were not aware of the proxy system and thought the price paid when winning was their own bid. Ely and Hossain (2009) find that observed sniping behavior can be explained by the presence of "naive bidders". Engelberg and Williams (2009) find that over 1% of bids in eBay auctions correspond to a strategy which they call "Discover-and-Stop" which involves bidding incrementally in a manner that minimizes the probability of outbidding the highest bid. They find by a natural experiment that bidding using this strategy is profitable for a seller for re-posting costs below 19% of the sale price of the item.

In the following sections we show that whenever there is a small probability that a player is a commitment type, in every THPE there is a time threshold such that truthful bids are placed, on path, only after this threshold. The time threshold is decreasing in the rate of arrival of the players. Thus, if the arrival rate of bidding opportunities is high, final bids are placed in the last instants of the auction. We interpret placing these late bids as "sniping".

In deriving these results, the key observation is that players prefer to place a late bid when facing a commitment type. For intuition, consider an auction with two bidders who have the same valuation, and suppose there is a small probability that one of them (player 1) is the shill bidder. Once player 1 reveals that he is not the commitment type, play must follow the unique trembling hand equilibrium when all players are rational. Thus, player 1 has incentives to imitate the commitment type in order to avoid the "rational players" equilibrium early in the auction, and waits until the last instants of the auction to snipe. Therefore, player 1 imitates

[16] A shill bidder is a bidder who is colluded with the seller with the purpose of increasing the final price paid by the winning player.

the commitment type, and the opposing player behaves as if best-responding to a commitment type. This imitation is beneficial to both players because the later that competing players plan to place their high bids, the higher is the probability these bids will not be realized. The expected payoff of the seller is lower than it would be if the players were to place high bids at the first opportunity.

5.1 Equilibrium against a potential shill bidder

In this section we describe equilibrium play when there is a small probability that one player is a shill bidder that is colluded with the seller. Consider the game with two players, players 1 and 2. There is a prior probability $\gamma_0 > 0$ that player 1 is a shill bidder who plays a heuristic strategy as described below. Player 2 is known to be a rational payoff type. We first consider the case with known valuations, and then extend to incomplete information about valuations.

Definition 5. A *shill bidder* is a player who places a bid equal to one increment above the current price at each opportunity but never exceeds the highest bid.[17]

The strategy of a shill bidder is to bid incrementally up to an increment below the highest bid in order to increase the price paid by the highest bidder. Our shill bidder plays an ideal version of the strategy that a shill bidder may play in eBay. In practice a shill bidder cannot observe the highest bid but can discover it with high probability as explained in Engelberg and Williams (2009).[18] In section 5.3 we consider more general commitment types, including a shill bidder who outbids the highest bid with some probability. All THPE feature sniping when there is a small probability that one of these more general commitment types is present.

We refer to a player that is not a shill bidder as a *rational player*. A *sniping bid* is defined as a truthful bid that is placed after a time threshold $\bar{t}(\lambda)$, such that $\bar{t}(\lambda)$ converges to zero as $\lambda \to \infty$. Thus, if opportunities arise very often, sniping bids arrive just before the auction ends. We say that a player "snipes" if she places a sniping bid.

[17] We make the assumption that shill bidder types do not tremble in an ε-constrained equilibrium to simplify the analysis.

[18] For example, suppose the current highest proxy is 5 dollars and the increment is one dollar. A shill-bidder who places a bid of $4.01 will see the current price increase to 5 instead of 5.01 and thus learn the highest bid. In practice, buyers tend to place their bids in dollar increments which makes this strategy likely to succeed.

Players with known valuations

Let $v_1 = v_2 = v > p + \Delta$ where p is the reserve or initial price set by the seller. We now show in this simpler setting that all THPE must involve sniping.

The first useful result is that a player is best off placing a small bid when facing a shill bidder. We state this result in the following Lemma.

Lemma 2. *If player 1 is known to be a shill bidder, and player 2 is rational with valuation $v \geq p + 2\Delta$, then player 2 bids $b = p + 2\Delta$ at the first opportunity.*

Proof. Let player 1 be a shill bidder, and let player 2 be rational. We will argue that player 2's optimal strategy involves bidding $p + 2\Delta$ at her first opportunity, and never bidding at opportunities in which she currently holds the highest bid.[19] Suppose that player 2 gets an opportunity at time t and has not had previous opportunities. Player 2's payoff from placing bid $p + 2\Delta$ and not placing other bids is given by

$$U_2^{SB}(t, p+2\Delta) = (v - p - \Delta)e^{-\lambda t} + (v - p - 2\Delta)\left(1 - e^{-\lambda t}\right).$$

To understand this expression, note that if the shill bidder has no arrivals, an event which occurs with probability $e^{-\lambda t}$, the payoff is $(v - p - \Delta)$. If the shill-bidder has one or more arrivals, the payoff is $(v - p - 2\Delta)$. At the first bid the shill-bidder discovers that the highest bid is $p + 2\Delta$ and stops bidding.[20]

At the same time, player 2's payoff from submitting a larger bid b at time t and never bidding again, conditional on player 1 being a shill bidder is given by

$$U_2^{SB}(t, b) = \sum_{i=1}^{(b-p-1)/\Delta} (v - p - i\Delta)e^{-\lambda t}\frac{(\lambda t)^i}{i!} + \sum_{(b-p)/\Delta}^{\infty} (v - b - \Delta)e^{-\lambda t}\frac{(\lambda t)^i}{i!}.$$

The shill bidder bids incrementally up to the highest bid. The probability that the shill bidder has i arrivals is $\frac{(\lambda t)^i}{i!}e^{-\lambda t}$ and the payoff for 2 in this case is $\max\{(v - (p + i\Delta)), (v - (b + \Delta))$.

For every t, $U_2^B(t, p + 2\Delta) > U_2^B(t, b)$. Thus, if player 2 knows that she is facing the shill-bidder, she prefers to bid $p + 2\Delta$ rather than $b > v + 2\Delta$. Moreover, if player 2 gets

[19] Note that this is a partial description of a strategy, which does not specify what player 2 does off the path of play.

[20] Note that a bid of $p + \Delta$ would be outbid by the shill bidder while trying to discover the highest bid. Therefore, a bid of $p + 2\Delta$ is more convenient than a bid of $p + \Delta$.

an opportunity while currently holding the winning bid, then she will win with probability 1, regardless of her action, but increasing her bid will increase the price paid with positive probability. By the OSDP, in order to verify that player 2's optimal strategy involves bidding $p+2\Delta$ at the first opportunity and never bidding when winning, it suffices to rule out deviations of the above forms. $\square$

Let's now return to the case in which player 2 is known to be rational, but there is a prior probability $\gamma_0 \in (0,1)$ that player 1 is a shill bidder. We say that a player *reveals rationality* if he takes an action that the commitment type would not take and that leads to a history of an opposing player which could not occur when the player is a shill bidder. Thus, if player 1 places a non-incremental bid, he reveals that he is not the shill bidder.

In order to describe equilibrium play when player 1 may be a shill bidder, it is useful to make an observation regarding continuation play once player 1 reveals rationality.

Lemma 3. *Equilibrium play once player 1 reveals rationality is as follows:*

1. *Once player 1 reveals rationality, both players bid truthfully at the first opportunity regardless of history.*

2. *If player 1 were to reveal rationality, he would do so by bidding truthfully.*

The proof is given in the appendix. We will now show that player 1 does not bid truthfully early in the auction.

Suppose the rational type of player 1 bids truthfully at the first opportunity (as in the no commitment types case). Player 2's payoff from bidding $p+2\Delta$ at the first opportunity and bidding $v-\Delta$ if outbid (as required by Lemma 3) is given by,

$$\gamma_t U_2^{SB}(t, p+2\Delta) + (1-\gamma_t)e^{-\lambda t}(v-p-\Delta), \tag{1}$$

where γ_t denotes 2's belief of player 1 being a shill bidder at time t. With probability $1-\gamma_t$ player 1 is a rational type, in which case, player 2 earns positive profits only if player 1 does not have another arrival.

Player 2's payoff from bidding $v-\Delta$ at time t is given by

$$\gamma_t U_2^{SB}(t, v-\Delta) + (1-\gamma_t)e^{-\lambda t}(v-p-\Delta). \tag{2}$$

Because $U_2^{SB}(t, p+2\Delta) > U_2^{SB}(t, v-\Delta)$ as long as $\gamma_t > 0$—that is as long as player 1 has not revealed rationality—player 2 prefers to bid $p+2\Delta$ instead of $v-\Delta$. In the appendix we show that player 2 would not want to place any bid above $p+2\Delta$ either.

If player 1 is expected to bid truthfully, player 2 best responds to the shill bidder. As a result, early in the auction player 1 would prefer to not bid truthfully and imitate the shill bidder: if he reveals rationality play reverts to the rational players' equilibrium which gives player 1 a lower payoff.[21]

We have seen that truthful bidding at the first opportunity does not constitute a THPE as it does in the absence of uncertainty about shill bidders. We now describe a THPE of the game with a potential shill bidder. Let (σ_1, σ_2) be a strategy profile satisfying the following: player 1 takes actions that are consistent with imitating the shill bidder until time threshold $\bar{t}_1 = 1/\lambda$. After $\bar{t}_1$ player 1 bids truthfully at the first opportunity. As long as player 1 has not revealed rationality, player 2 places a bid of $p+2\Delta$ at the first opportunity after a time threshold $\bar{t}_2$ (to be determined) and does not bid again.

To see that (σ_1, σ_2) are the limit of ε-constrained equilibria we use the OSDP (Proposition 1). Recall that player 1 bids truthfully when he reveals rationality (Lemma 3), and thus we only need to check histories before player 1 reveals rationality.[22] Let $(\sigma_1^\varepsilon, \sigma_2^\varepsilon)$ denote ε-constrained strategies. Suppose σ_2^ε puts as much probability weight as possible on the actions recommended in σ_2 at each history and weight ε on all other actions. Player 1's payoff from bidding truthfully at time t if player 2 behaves according to σ_2^ε is given by

$$e^{-\lambda t}(v - p - \beta\Delta)\mathbb{P}(\text{players do not tremble}) + M\mathbb{P}(\text{players tremble}), \tag{3}$$

where $\beta = 3$ if player 2 has placed a bid and $\beta = 1$ otherwise, and $M > 0$ is a constant bounded by v. The payoff from not revealing rationality at time t and bidding truthfully at the next opportunity is given by

$$e^{-\lambda t}\lambda t(v - p - \beta\Delta)\mathbb{P}(\text{players do not tremble}) + M\mathbb{P}(\text{players tremble}). \tag{4}$$

If the strategy σ_1^ε puts as much weight as possible on placing a bid when (3)>(4) and puts

[21] Player 1's payoff from bidding truthfully at time t is $e^{-\lambda t}(v-p-\Delta)$. The payoff from waiting until after time threshold $\bar{t}$ to bid truthfully is $\lambda\bar{t}e^{-\lambda\bar{t}}(v-p-\Delta)$. Thus, imitating the shill bidder is profitable as long as $\lambda t > 1$.

[22] Note that, when imitating the shill bidder, player 1 only places losing bids. When player 1 outbids player 2, she reveals rationality immediately.

as much weight as possible on not placing a bid for t such that (3)>(4), σ_1^ε does not have any profitable local deviations. Moreover, $\mathbb{P}(\text{players tremble}) = O(\varepsilon)$ as $\varepsilon \to 0$. Thus in the limit as $\varepsilon \to 0$ the condition for player 1 to place a bid is $\lambda t \leq 1$.

Now assuming player 1 plays according to σ_1^ε, player 2's payoff from bidding $p+2\Delta$ at time t is,

$$\gamma_t(e^{-\lambda t}(v-p-\Delta)+(1-e^{-\lambda t})(v-p-2\Delta))+(1-\gamma_t)(v-\mathbb{E}_{\mu_2^t}(b_1)-\Delta)e^{-1}+O(\varepsilon), \quad (5)$$

where $\mathbb{E}_{\mu_2^t}(b_1)$ denotes the expected bid that player 1 places while he only takes actions consistent with the shill bidder's behavior. The first term is the payoff player 2 receives when facing the shill bidder. If the shill bidder does not get an opportunity to bid, player 2's payoff is $v-p-\Delta$. If the shill bidder receives at least one opportunity, he would place a bid and discover player 2's bid and stop bidding. Player 2's payoff in the latter case is $v-p-2\Delta$. The second term is the payoff when facing a rational player 1. Starting at time approximately $1/\lambda$ player 1 places a bid at the first opportunity. Thus, if the rational player 1 receives any opportunity after time $1/\lambda$, player 2's payoff is zero.

Player 2's payoff from waiting until next opportunity, and bidding $p+2\Delta$ at that opportunity, is given by,

$$\gamma_t(\lambda t e^{-\lambda t}(v-p-\Delta)+(1-\lambda t e^{-\lambda t}-e^{-\lambda t})(v-p-2\Delta))+(1-\gamma_t)(v-\mathbb{E}_{\mu_2^t}(b_1)-\Delta)e^{-1}(1-e^{-\lambda t}) \quad (6)$$
$$+O(\varepsilon).$$

As before, the first term is the payoff originating from facing the shill bidder and the second term is the payoff is the payoff when player 1 is rational. If strategy σ_2^ε puts the maximum possible weight on not bidding whenever (6)>(5) and puts the maximum possible weight on bidding $p+2\Delta$ whenever (6)<(5) then σ_2^ε does not have any local profitable deviations.[23] $\bar{t}_2$ is defined to be the value at which (6) = (5). Taking the limit when $\varepsilon \to 0$, we see that (σ_1, σ_2) must be a THPE.

Note that the timing at which player 2 places the first bid depends on how accurately player 1 imitates the shill bidder in equilibrium. If player 1 behaves exactly like the shill

[23] Note that placing a bid higher than $p+2\Delta$ can only be worse for player 2 because it is worse against the shill bidder and against the rational player who imitates the shill bidder, while it makes no difference after the rational player 1 starts bidding truthfully.

bidder then $\gamma_t = \gamma_0$ but if the rational player does not bid incrementally at the first opportunity then γ_t decreases over time. The value of $\mathbb{E}_{\mu_2^t}(b_1)$ is also affected by the rational player 1's imitation. There may be multiple THPE that vary in player 1's behavior before revealing rationality. Note also that as $\gamma_0 \to 0$, the threshold at which player 2 places the first bid $\bar{t}_2$ converges to ∞.

We have shown that when there is a positive probability that player 1 is a shill bidder there are equilibria that feature sniping by players 1 and 2 and low early bids by player 2 when the probability that player 1 is a shill bidder is low. Early bids are referred to as *squatting bids* and are also a common behavior in eBay auctions. In fact, all THPE must be of the form of (σ_1, σ_2) in this setting. The result is stated in the following Proposition.

Proposition 3. *Consider an auction with two players such that $v_1 = v_2 = v$, $p + 2\Delta < v$ and suppose there is a prior probability $\gamma_0 > 0$ that player 1 is a shill bidder. In all THPE:*

1. *Player 1 does not place any bids at opportunities before time threshold $\bar{t}_1 = 1/\lambda$, and places a truthful bid at the first opportunity after $\bar{t}_1$.*

2. *As long as player 1 has not revealed rationality, player 2 does not place any bids until the first opportunity after time threshold $\bar{t}_2$. If player 1 reveals rationality, player 2 bids truthfully at the first opportunity.*

The proofs of this section are in the appendix.

Note that the player who is known to be rational, player 2, receives a higher payoff than player 1.[24] Now, consider the case in which there is a prior probability that player 2 is also a shill bidder. In equilibrium, the first player to have an opportunity would choose to place a small (squatting) bid and, thus, reveal rationality. After one player is known to be rational the equilibrium play that follows must be as described in Proposition 3. This discussion is formalized in the following Corollary.

Corollary 1. *If there is prior probability $\tilde{\gamma}_0 > 0$ that player 2 is a shill bidder, then the first player to receive an opportunity places a bid of $p + 2\Delta$ and equilibrium play follows as in Proposition 3.*

[24]The payoff from revealing rationality is $e^{-1}(v - p - \Delta)$ while the payoff from imitating the behavioral player is $e^{-1}(v - p - 3\Delta)$ when the opposing player receives an opportunity and $(v - p - \Delta)$ otherwise.

5.2 Private information about players' valuations

We now allow players have private information about their own valuation for the good in addition to incomplete information about behavioral types. We show that, as in the known valuations case, all THPE feature sniping bids. The threshold times at which players start to place their truthful bids have a "ladder structure," in which players with higher valuations begin to submit sniping bids, conditional on getting bidding opportunities, sooner than their lower valuation counterparts.

5.2.1 Two-valuation case

To illustrate the results, we first consider the case in which the two players may have two possible valuations for a good, v^L and v^H. Let p^L denote the prior probability that a player's valuation is v^L. For simplicity, we assume that the players' valuations are independently and identically distributed whenever neither player is a shill bidder. As before, there is a prior probability γ_0 that player 1 is a shill bidder.

We show that equilibrium play in this setting is as follows, in all THPE. As long as player 1 has not revealed rationality, player 2 places a bid of $p+2\Delta$ at the first opportunity. Player 2 of valuation v^L bids truthfully only after being outbid. Player 2 of valuation v^H bids truthfully after a time threshold $\bar{t}_2(v^H)$. Player 1 bids truthfully at the first opportunity after time threshold $\bar{t}_1(v_1)$ or at the first opportunity after player 2 bids after threshold $\bar{t}_2(v^H)$. The sniping thresholds satisfy $\bar{t}_1(v^L) \leq \bar{t}_2(v^H) \leq \bar{t}_1(v^H)$.

We now sketch the argument.

As in the known valuation case, the proofs rely on the one-shot deviation principle. To simplify notation, we omit the $O(\varepsilon)$ terms in this part. Lemma 3 applies. Therefore, player 1 reveals rationality by bidding truthfully, and once player 1 reveals rationality player 2 bids truthfully at the first opportunity.

Threshold of player 1 of type v^L Let $t \leq \bar{t}_2(v^H)$. Player 1 of valuation v^L's payoff of bidding at time t and revealing rationality is

$$e^{-\lambda t}(v^L - p - \Delta),$$

since player 2 bids truthfully as soon as player 1 bids.

The payoff of waiting until the next opportunity to reveal rationality is given by

$$\lambda t e^{-\lambda t}(v^L - p - \Delta)p^L + (1 - e^{-\lambda t})e^{-\lambda t}(v^L - p - \Delta)(1 - p^L).$$

Player 2 of type v^L only bids after player 1 places a bid. Thus, if facing player 2 of valuation v^L, player 1 gets a positive payoff if once her bid is placed player 2 does not get another arrival. This event has probability $\lambda t e^{-\lambda t}$. If player 1 is facing player 2 of valuation v^H, then player 1 gets a positive payoff if he gets another opportunity to bid and player 2 does not .

If for example $\lambda t \geq \max\{2, \log\left(\frac{1-p^L}{p^L}\right)\}$ then player 1 of type v^L prefers to not bid at time t. At t close to zero, player 1 prefers to bid. Thus, in equilibrium there must be a threshold $\bar{t}_1(v^L)$ from which player 1 starts bidding truthfully.

Threshold of player 2 of type v^H Assume that player 2 has already placed a bid of $p + 2\Delta$. Let's now see that there is a time threshold $\bar{t}_2(v^H)$ after which player 2 of type v^H bids truthfully at the first opportunity. Once player 2 snipes, a rational player 1 must bid $v_1 - \Delta$ at the first opportunity in a trembling hand perfect equilibrium. Assume $\bar{t}_1(v^H) \geq t \geq \bar{t}_1(v^L)$. In the limit when $\gamma_t = 0$, if player 2 is of type v^H and the current price is $p + \Delta$, the payoff to player 2 of bidding truthfully at time t is given by

$$e^{-\lambda t}(v^H - p - \Delta) + (1 - e^{-\lambda t})(v^H - v^L)p^L, \tag{7}$$

since player 1 bids truthfully at the first opportunity after player 2 places a bid. The payoff of bidding at the next bidding opportunity is given by

$$\begin{aligned} e^{-\lambda t}(1 - e^{-\lambda t})(v^H - p - \Delta)(1 - p^L) + p^L\Big(\lambda(t - \bar{t}_1(v^L))e^{-\lambda t}(v^H - p - \Delta) + \\ \left(1 - e^{-\lambda(t - \bar{t}_1(v^L))} - \lambda(t - \bar{t}_1(v^L))e^{-\lambda(t - \bar{t}_1(v^L))}\right)e^{-\lambda \bar{t}_1(v^L)}(v^H - v^L) + \\ \left(1 - e^{-\lambda t}\right)(1 - e^{-\lambda \bar{t}_1(v^L)})(v^H - v^L)\Big). \end{aligned} \tag{8}$$

To understand the previous expression note that if player 1 is of type v^H, which occurs with probability $(1 - p^L)$, 2 gets a positive payoff only if player 1 does not get an opportunity. If player 1 is of type v^L, which occurs with probability p^L, between times t and $\bar{t}_1(v^L)$ player 1 bids the valuation minus an increment at the first bidding opportunity after 2 bids and between times $\bar{t}_1(v^L)$ and zero player 1 bids $v^L - \Delta$ regardless of 2's play. The term

$$\lambda(t-\bar{t}_1(v^L))e^{-\lambda t}$$

is the probability that player 1 gets no opportunities after $\bar{t}_1(v^L)$ and that player 2 receives an opportunity to bid and 1 does not get an opportunity to place a bid as a response between t and $\bar{t}_1(v^L)$. The term

$$\left(1-e^{-\lambda t}\right)(1-e^{-\lambda \bar{t}_1(v^L)})$$

is the probability that player 2 gets one opportunity before the end of the auction and player 1 gets at least one opportunity after $\bar{t}_1(v^L)$. The term

$$\left(1-e^{-\lambda(t-\bar{t}_1(v^L))}-\lambda(t-\bar{t}_1(v^L))e^{-\lambda(t-\bar{t}_1(v^L))}\right)e^{-\lambda \bar{t}_1(v^L)}$$

is the probability that player 1 gets an opportunity to place a bid after 2 places a bid between times t and $\bar{t}_1(v^L)$ and does not get an opportunity after time $\bar{t}_1(v^L)$.[25] Rearranging equation (8) shows that player 2's payoff from waiting until the next opportunity to bid is given by:

$$e^{-\lambda t}(1-e^{-\lambda t})(v^H-p-\Delta)(1-p^L)+p^L\left(\lambda(t-\bar{t}_1(v^L))e^{-\lambda t}(v^L-p-\Delta)+(1-2e^{-\lambda t}+e^{-\lambda t-\lambda \bar{t}})(v^H-v^L)\right). \tag{9}$$

Thus, in the limit as $\gamma_t \to 0$ player 2 of valuation type v^H is indifferent between bidding truthfully and waiting until the next opportunity to place that bid if

$$\lambda(t-\bar{t}_1(v^L))=\frac{\left(\rho e^{-\lambda t}+1\right)(v^H-p-\Delta)+(v^H-v^L)(1-e^{-\lambda \bar{t}_1(v^L)})}{v^L-p-\Delta}.$$

where $\rho \equiv \frac{1-p^L}{p^L}$. The time threshold $\bar{t}_2(v^H)$ at which player 2 bids truthfully at the first opportunity is the time t that satisfies the previous equation.

The time threshold of player 2, $\bar{t}_2(v^H)$, is increasing in the current price. That is, if the current price is higher, player 2 snipes earlier. Intuitively, player 2 trades-off a bigger surplus

[25] $e^{-\lambda \bar{t}_1(v^L)}$ is the probability that player 1 has no arrivals after $\bar{t}_1(v^L)$. The term $\left(1-e^{-\lambda(t-\bar{t}_1(v^L)}-\lambda(t-\bar{t}_1(v^L))e^{-\lambda(t-\bar{t}_1(v^L))}\right)$ is the probability of at least two Poisson arrivals in the interval $(t,\bar{t}_1(v^L))$, one for player 2 after which we need pay attention only to player 1's arrival process.

when winning the auction against a lower probability of success. For lower p the value of winning is higher when player 1 does not receive an opportunity. Thus, player 2 is willing to risk not getting another opportunity in order to delay the bidding of the low valuation player 1 and obtain a greater surplus when winning the auction.

In an analogous manner, we find the threshold $\bar{t}_1(v^H)$ at which rational player 1 of type v^H bids.

By determining the unique equilibrium play towards the end of the auction and arguing recursively for earlier times, we prove in the appendix that this "ladder structure" for sniping thresholds is the unique equilibrium play in all THPE.

5.2.2 General case

The analysis from the two-valuation case generalizes to a valuation distribution that is on the grid $\{v^0, v^1, v^2, \ldots, v^n\}$ with $v^i < v^j$ if $i < j$ and $v^0 > p + 2\Delta$. We show that the type of player 2 with the lowest valuation prefers to place a low bid as long as player 1 has not revealed rationality.[26] Thus, the lowest type of player 1 prefers to not reveal rationality early in order to prevent the lowest type of player 2 from responding with a high bid. Now, the second to lowest type of player 2 has incentives to bid earlier than the lowest type. Thus if player 1 sees that player 2 places a higher bid, player 1 knows that player 2 has submitted a high bid, and thus player 1 has incentives to bid truthfully in a THPE. Thus, the second to lowest type of player 2 prefers to delay his truthful bids until a time threshold. Following this argument we obtain time thresholds, increasing in a player's valuation, that determine when players place their truthful bids in equilibrium. Formally, equilibrium play is described in the following Proposition.

Proposition 4. *Consider an auction with two players, player 1 and player 2, who have privately-known iid valuations in a grid $\{v^0, v^1, v^2, \ldots, v^n\}$ with $v^i < v^j$ if $i < j$. Let there be an initial probability $\gamma_0 > 0$ that player 1 is a shill-bidder type. Suppose the reservation price p^0 satisfies $p^0 < v^0 - 2\Delta$. There is $\bar{\gamma}_0 > 0$ such that if $\gamma_0 < \bar{\gamma}_0$, then in every THPE there are functions $\{t_1^k(p)\}_{k=1}^n$ and $\{t_2^j(p)\}_{j=1}^n$ which map the current price to positive numbers, such that*

[26] The lowest type of player 2 knows that if her opponent is not a shill bidder, then the opponent will always outbid her toward the end of the auction if the opponent gets an opportunity. As a result, there is nothing to be gained against a rational player 1 by bidding high. In contrast, if player 1 is a shill bidder, it is strictly better for player 2 to submit as low a bid as possible.

– *Player 2 with valuation v^0 bids $p^0 + 2\Delta$ at the first opportunity, and bids $v^0 - \Delta$ at the first opportunity after player 1 reveals rationality.*

– *If player 1 has not revealed rationality, player 2 holds the winning bid, player 2 has valuation v^j, and the current price is p, then player 2 bids $v^j - \Delta$ at the first arrival after $t_2^j(p)$.*

– *Player 1 of valuation v^k does not outbid player 2 before time $t_1^k(p)$, and bids $v^k - \Delta$ after time $t_1^k(p)$.*

– *$t_1^k(p)$ is increasing in k, and $t_2^j(p)$ is increasing in j. For all j and k, $t_1^k(p)$ and $t_2^j(p)$ are increasing in p.*

The proof is in the appendix. In the appendix, in addition we derive upper bounds on the sniping thresholds $t_i^k(p)$ in Proposition 4. Let p_0 be the prior probability of player 2 being of type v^0. The timing at which players 1 and 2 places a bid $v^k - \Delta$ and $v^k - \Delta$, t_1^k, t_2^j and $\hat{t}_2^j$ are bounded above by $\bar{t}_1^k$, $\bar{t}_2^j$ and $\bar{\hat{t}}_2^j$ which are defined defined recursively. We have:

$$t_2^0 = \bar{t}_2^0 = \bar{\hat{t}}_2^j = 0, \quad \lambda t_1^0 \leq \max\left\{\ln\left(\frac{1}{p_0}\right), 2\right\} = \lambda \bar{t}_1^0$$

$$\lambda t_2^j \leq \lambda \bar{t}_1^{j-1} + \frac{(1 - e^{-\lambda \bar{t}_1^{j-1}})(v^j - v^0)}{v^0 - p - \Delta} = \lambda \bar{t}_2^j \tag{10}$$

$$\lambda t_1^k \leq \lambda \bar{t}_2^k + \frac{3(v^k - p - \Delta)}{v^0 - p - \Delta} = \lambda \bar{t}_1^k \tag{11}$$

$$\lambda \hat{t}_2^j \leq \lambda \bar{t}_1^{j-1} + \left(\frac{2(1 - e^{-\lambda \bar{t}_1^{j-1}})(v^j - v^0)}{v^0 - p - \Delta}\right)^{\frac{1}{2}} = \lambda \bar{\hat{t}}_2^j. \tag{12}$$

Each $\bar{t}_i^k$ converges to zero as λ converges to ∞. Thus, as λ grows large the players are placing sniping bids in the last instants of the auction.

5.3 Other commitment types

The results about sniping bids extend to settings where there may be commitment players playing other fixed heuristic strategies. In this section we consider heuristic bidders who place low bids. This more general set of commitment players includes players who place a bid at one increment above the current price at every opportunity and shill bidders who may outbid the highest bid by mistake.

Suppose players 1 and 2 both have known valuation v. Let $U_2^{\xi}(t,v)$ denote player 2's payoff from bidding $v-\Delta$ at time t if player 1 is a commitment type ξ. Let σ_2 denote a strategy of player 2, and let $U_2^{\xi}(t,\sigma_2)$ denote the expected payoff of player 2 when using strategy σ_2 against non-strategic type ξ. $u_1^{\xi}(t)$ denotes the expected payoff at time t of player 1 if player 2 believes 1 is type ξ player with probability 1.

Definition 6 (Commitment type)**.** We say a non-strategic type ξ is a *commitment type* if a player of type ξ bids at most one increment above the current price at each opportunity and for each λ, there is a strategy $\sigma_2^{\xi}(\lambda)$ of player 2 such that for every t $\lim_{\lambda\to\infty}\frac{U_2^{\xi}(t,v)}{U_2^{\xi}(t,\sigma_2^{\xi}(\lambda))}=0$.[27]

Commitment types are players that only place incremental bids and against whom opposing players do not want to place truthful bids early. For example, a shill bidder is a commitment type, as is a bidder that raises his bid by one increment at each arrival until holding the lead. By placing a truthful bid after a threshold $\bar{t}=\frac{t_0}{\lambda}$ for some $t_0>0$, player 2 can guarantee a payoff that is bounded away from zero. Thus, there is a strategy $\sigma_2(\lambda)$ that satisfies $\lim_{\lambda\to\infty}\frac{U_2^{\xi}(t,v)}{U_2^{\xi}(t,\sigma_2^{\xi}(\lambda))}=0$.

Both players are thought to be a commitment type with initial probability γ_0. We now show that all equilibria must exhibit sniping for other possible commitment types.[28]

Let's now see that player 1 must place a truthful bid after a threshold that goes to zero as $\lambda\to\infty$.

Suppose the commitment type does not outbid player 2 as in the shill-bidding example. If σ_2^{ξ} is player 2's best response to the commitment type of player 1 then the strategy σ_2^{ξ} lets player 2 have the winning bid at time t. The payoff from bidding $v-\Delta$ is

$$\gamma_t U_2^{\xi}(t,v)+(1-\gamma_t)\left(e^{-\lambda t}(v-p-\Delta)\right). \tag{13}$$

If player 2 places a small bid, she has the opportunity to learn whether 1 behaves as the commitment type or the rational player. The payoff from playing according to σ_2^{ξ} and best responding to the rational and commitment player is

$$\gamma_t U_2^{\xi}(t,\sigma_2^{\xi})+(1-\gamma_t)\left(e^{-\lambda t}(v-p-\Delta)\right). \tag{14}$$

[27] We assume that commitment types do not tremble to simplify the analysis.

[28] This statement is not vacuous because this game must have equilibria by Moroni (2014).

Equation (13) is strictly above (14) whenever $U_2^{\xi}(t,\sigma_2^{\xi}) > U_2^{\xi}(t,v)$. Thus if player 1 has not placed a bid, player 2 would choose to bid according to σ_2^{ξ} in order to find out 1's type.

Now, suppose player 1 gets an opportunity to place a bid at time t and in equilibrium he is expected to bid truthfully. If player 1 imitates the commitment player, player 2 will continue to best respond to the commitment type. Since player 1 gets a strictly positive payoff when player 2 best-responds to the commitment type and the payoff of revealing rationality converges to zero as $\lambda \to \infty$, imitating the commitment type is a profitable deviation from bidding $v-\Delta$ at time t. Thus, in equilibrium player 1 cannot place a bid of $v-\Delta$ early in the auction at every opportunity. Player 1 would bid $v-\Delta$ starting at time $t = 1/\lambda$ because player 2 acts as if best responding to the commitment type and revealing rationality leads to equilibrium play in which 2 bids truthfully at the first opportunity.

Now, let's consider the case in which the commitment type may outbid player 2. Assume that the commitment type is an incremental bidder that bids the current price plus an increment at each opportunity or outbids with an exogenous fixed probability.[29] Player 2 will prefer to place a "test bid" instead of a truthful bid. The payoff from bidding $v-\Delta$ is

$$\gamma_t U_2^{\xi}(t,v) + (1-\gamma_t)e^{-\lambda t}(v-p-\Delta). \tag{15}$$

If player 2 places a small bid, she can find out whether 1 behaves as the behavioral type or the rational player. Let $\tilde{U}_2^{\xi}(t,\sigma_2^{\xi})$ be the payoff of player 2 from placing a bid that is two increments above the current price at time t and best responding against the behavioral type afterwards. The payoff from bidding two increments and best responding to the rational and commitment player is

$$\gamma_t \tilde{U}_2^{\xi}(t,\sigma_2^{\xi}) + (1-\gamma_t)e^{-\lambda t}(v-p-\Delta). \tag{16}$$

Equation (15) is strictly above (16) whenever $\tilde{U}_2^{\xi}(t,\sigma_2^{\xi}) > U_2^{\xi}(t,v)$ thus if player 1 has not placed a bid, player 2 would choose to bid two increments rather than bidding truthfully. Now if player 2 does not place a bid at time t and bids two increments at the next opportunity her payoff is

$$\gamma_t \int_0^t \tilde{U}_2^{\xi}(\tau,\sigma_2^{\xi})\lambda e^{-\lambda\tau}\, d\tau + (1-\gamma_t)e^{-\lambda t}(1-e^{-\lambda t})(v-p-\Delta). \tag{17}$$

[29] An incremental bidder satisfies the conditions stated in the definition of a commitment type.

Thus, if γ_t is small or $\tilde{U}_2^{\xi}(t,\sigma_2^{\xi}) > \int_0^t \tilde{U}_2^{\xi}(\tau,\sigma_2^{\xi})\lambda e^{-\lambda\tau}d\tau$ player 2 prefers to place a test bid instead of not placing any bids.

From the discussion above, it follows that the first player to get an opportunity must place a test bid. After one player has placed a test bid, the other player who is expected to bid truthfully can pretend to be a behavioral type and obtain a high payoff. To determine how early player 1 is willing to place a bid equal to $v-\Delta$ we use the fact that in equilibrium one-shot deviations must be unprofitable. The earliest time t at which player 1 is willing to bid truthfully is such that the payoff of waiting until the next opportunity to bid $v-\Delta$ is less than the payoff of bidding $v-\Delta$ at time t. For small enough γ_0 unprofitable one-shot deviations implies

$$\int_0^t u_1^{\xi}(t)\lambda e^{-2\lambda\tau}d\tau + \int_0^t (v-p-\Delta)\lambda e^{-3\lambda\tau}d\tau \le e^{-\lambda t}(v-p-\Delta). \tag{18}$$

The left hand side is a lower bound on player 1's payoff from imitating the commitment type. $\lambda e^{-2\lambda(t-\tau)}$ is the probability density of the event that player 2 arrives first. If player 2 moves first she would best respond to the commitment type of player 1 and at player 1's next opportunity 1 bids $v-\Delta$ and obtains a positive payoff only if player 2 does not get another opportunity which happens with probability $e^{-\lambda\tau}$ if 1's next opportunity arrives at time τ. The first integral is greater or equal than

$$u_1^{\xi}(t)\frac{1}{2}\left(1-e^{-2\lambda t}\right). \tag{19}$$

Since the right-hand side of (18) converges to zero as $\lambda \to \infty$, the time t at which the player prefers to bid truthfully converges to zero as λ converges to ∞.

Therefore, for λ large enough that the time t that satisfies equation (18) is small, and γ_0 small enough that players are willing to place a "test bid", we can conclude that player 1 waits until a late time threshold to place a truthful bid.

The following Proposition summarizes the previous discussion.

Proposition 5. *If the commitment type is either a player who never outbids the current price, or a type that outbids the current price by one increment with some fixed probability, there exists $\bar{\gamma}_0$ such that for every $\gamma_0 < \bar{\gamma}_0$ equilibrium play in THPE must exhibit sniping.*

5.4 Perfect Bayesian Equilibrium.

The uniqueness of equilibrium does not hold when we consider a weaker equilibrium notion. In this section we show that when we consider Perfect Bayesian Equilibria there may be many equilibria. Some of these equilibria present features similar to the "implicitly collusive" equilibria discussed in Roth and Ockenfels (2002).

The following are examples of equilibrium play in Perfect Bayesian Equilibria in the auction with no commitment types.

Example 1 (Perfect Information)**.** Suppose $v_1 = v_2 = v$ almost surely, and suppose the initial price is $p = 0$. Some equilibria give a high payoff to the players and are supported by indifferences which are ruled out by the refinement. The following are examples of equilibrium play:

- Upon the first arrival,the arriving player bids $v - \Delta$. Once a bid has been placed, the opposing player does not outbid. If the initial price is zero, this equilibrium gives payoff $v \cdot (1 - e^{-2\lambda T})/2$.

- The first player to arrive bids $v - \Delta$, and the second player bids $b \in [\Delta, v]$ at her first arrival.

- One player bids Δ upon arrival and the other player bids his valuation upon arrival.

- A player bids Δ upon the first arrival, and does not bid as long as he holds the lead, but bids $v - \Delta$ as soon as the other player places a bid before $t = 1/\lambda$. The player who hasn't bid does not place a bid until time $1/\lambda$, at which time she places a bid equal to v.

It is easy to verify that the previous examples of play are part of Perfect Bayesian Equilibria by noting that they do not have any local profitable deviations (see Proposition 1).

6 Empirical Analysis

In order to test the model's predictions, we gathered data from online auctions for a good that is likely to have private values and minimal differences in quality across auctions. The data consists of 1977 eBay auctions that took place from 31/12/11 until 01/15/12 for a video game called Modern Warfare 3 for Xbox 360. After removing auctions which ended early via "buy it now", we are left with 1473 auctions, of which 1456 ended with a sale by auction. These data were gathered from the eBay website using a Python script. The site displays prices, bids, sellers, ratings and duration of each auction finished within the previous two weeks. These data include the timings of all bids, whether placed manually or by the proxy system.[30] Table 1 displays summary statistics for observed variables. There is little variation in product and observable seller characteristics. Most of the products are of "brand new" or "like new" quality. Most sellers have very high percentages of positive feedback. The typical sale price, including shipping, is around $43, but there is considerable variation.

Table 1: Summary statistics

Variable	**Mean**	**Std. Dev.**	**Min.**	**Max.**	**N**
final price	39.178	5.082	2.25	65	1473
shipping price	3.697	1.229	0.01	30	1473
shipping price not stated	0.037	0.19	0	1	1473
price including shipping	42.873	4.998	6.25	69	1418
sold via auction or buy-it-now	0.99	0.097	0	1	1473
number of bids	10.547	8.893	1	44	1473
condition==Acceptable	0.009	0.094	0	1	1473
condition==Good	0.018	0.132	0	1	1473
condition==Very Good	0.09	0.286	0	1	1473
condition==Like New	0.49	0.5	0	1	1473
condition==Brand New	0.394	0.489	0	1	1473
reserve price	17.371	15.319	0	49.99	1473
reserve + shipping	21.068	15.342	1	53.99	1473
seller score	1974.48	20218.094	-1	439177	1473
duration (days)	4.009	2.419	1	10	1473
seller percentage	99.015	4.247	33.3	100	1421

We find that empirical patterns in this data are consistent with the model's predictions. First, as predicted we find a bimodal distribution of bid timings, with the majority of bids arriving either early or late in the auction. Second, consistent with players imitating a shill bidder before eventually sniping, we show that incremental bidding is more likely in players'

[30]If a player sets a proxy bidder at a value below the highest bid, it is revealed once it is outbid by another player.

earlier bids than in their final bids. Third, consistent with 4, we show that the arrival of a bid "triggers" bidding by competing players. In particular, we test whether the presence of a bid in a time window is correlated with bids in successive time windows, controlling for the timing of the bid. We find there is a positive correlation between the two. The effect is stronger for the window immediately after a bid. We control for auction covariates, including the number of bidders who participated, to try to account for auction heterogeneity. This finding is consistent with our result that, in equilibrium, players are waiting for either a time threshold or another player's bid, whichever comes first, in order to begin sniping at their next opportunity. Fourth, consistent with the comparative statics predictions in theorem 4, we show that sniping begins earlier when the price is higher, and that bids occur earlier for players with higher valuations.

In the appendix, table 5 shows summary statistics for the timing at which the winning bid arrives at the auction for sold items. Note that the median time left before the end of the auction before the last bid arrives is 0.0003 of the total duration of the auction. That is, in an auction that lasts for 7 days the median final bid arrives 2.82 minutes before the deadline. Table 6 shows statistics for the final price including shipping costs for the auctions that end with a sale. Tables 7 and 8 show that the number of bidders is below the number of bids for each percentile, indicating that there are players who place more than one bid in each auction.

6.1 Timing of Bids

The empirical literature on online auctions with hard deadlines has shown that many bidders place proxies at the very end in the auction. (Bajari and Hortacsu (2003), Roth and Ockenfels (2002)) The same pattern occurs in our data. Figure 6.1 shows histograms of the time left before the end of the auction in which bids arrive.[31] In our model, the first rational player to arrive at the auction places a bid. Thus, one should expect a spike in the number of bids at the beginning of the auction and a second spike at the end when players snipe or reveal rationality. From Figure 6.1 we can see that bids are bimodal with a small peak at the beginning of the auction and a much bigger peak at the end. Figures 6.1 and 6.1 demonstrate that many bids come within the final seconds of their respective auctions.

[31] Our data records timings only up to the second.

Figure 1: When are bids placed? Proportion of time remaining when players submit bids

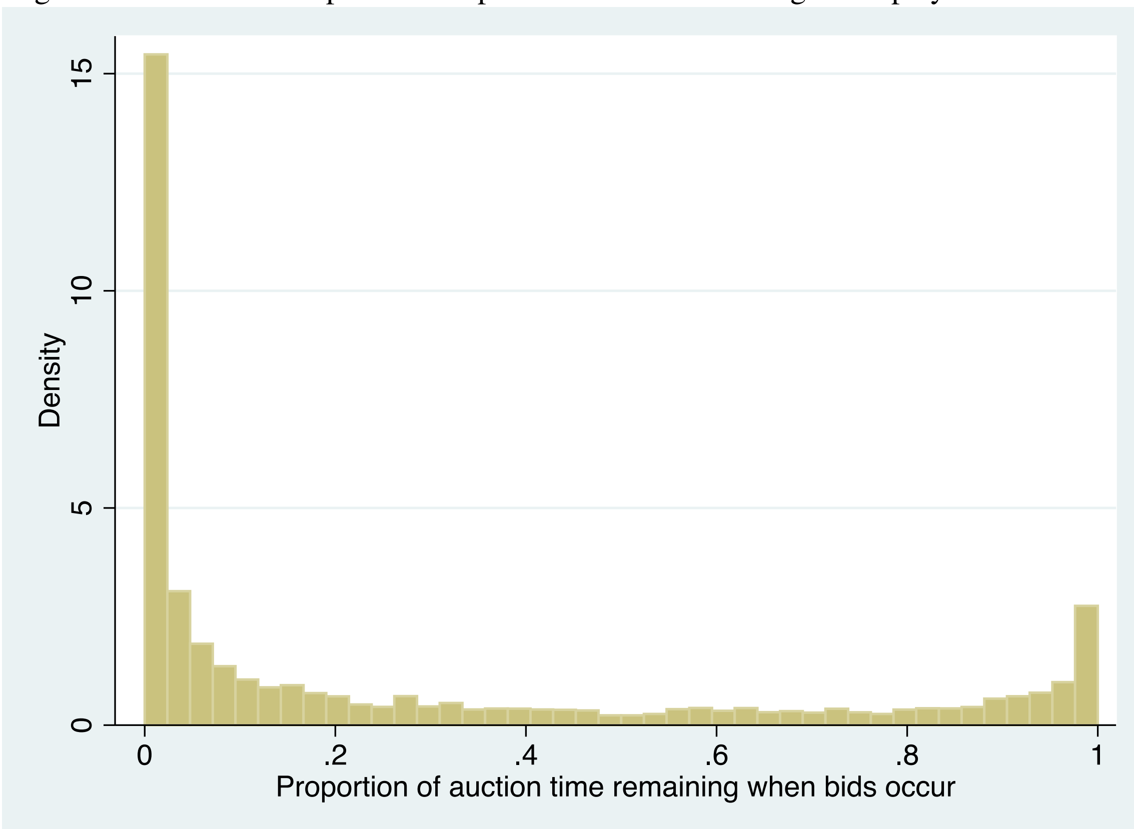

Figure 2: Minutes remaining when players submit bids, within final hour of each auction

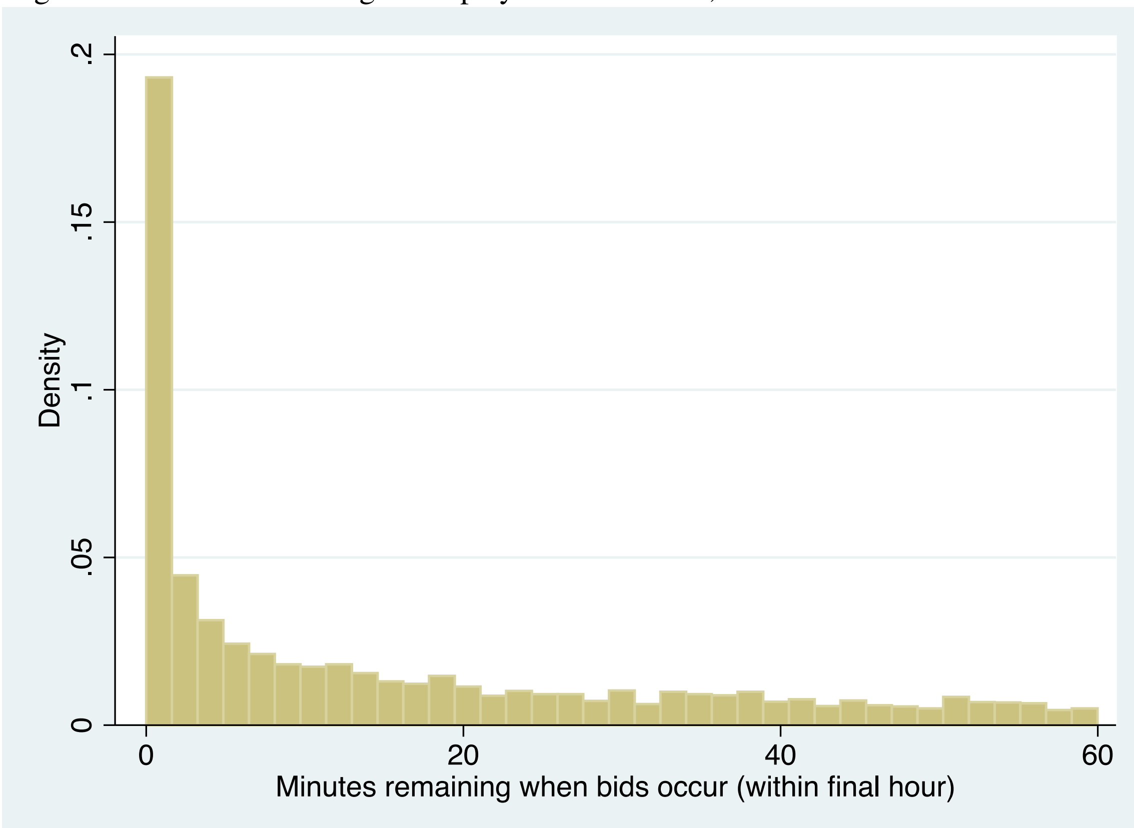

Figure 3: Seconds remaining when players submit bids, within final 100 seconds of each auction

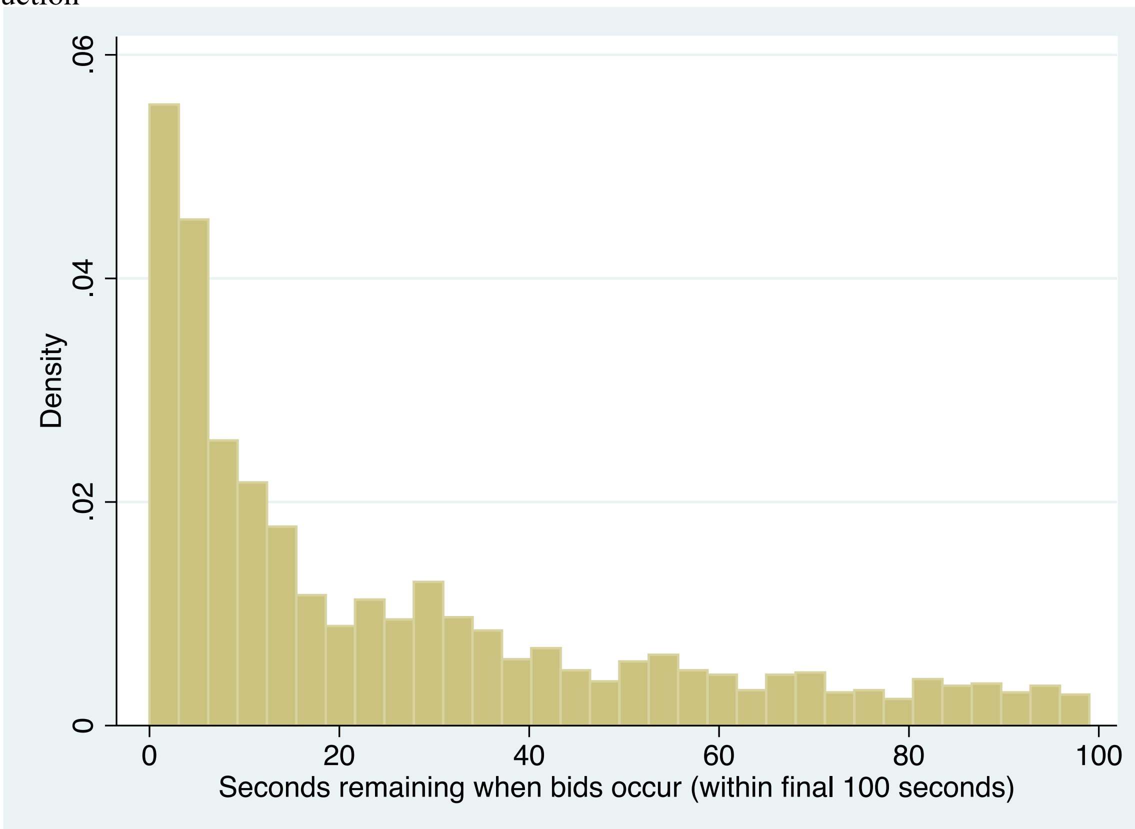

6.2 Incremental Bidding

We say that an incremental bid occurs when a player bids the minimum possible bid that can be placed at that time, which corresponds to the current price plus an increment. The majority of auctions have at least one instance in which a player places an incremental bid. In fact, 63.8% of auctions with at least one bidder, and 84.5% of auctions with 2 bidders or more, present at least one incremental bid. In our sample, incremental bids account for 35.9% of submitted bids.

Our model predicts that incremental bidding occurs before the end of auctions, when a player is imitating a behavioral type. Because a player who imitates a behavioral type will eventually want to reveal rationality by bidding his valuation (theorem 4), players' final bids are relatively less likely to be incremental. Additionally, with more than two players, we conjecture that incremental bids can be placed by players who are known to be rational as long as at least one of the participating players has not revealed rationality. We find that these predictions are consistent with the data.

Let $incr_i$ be an indicator for bid i being an incremental bid, let $last_i$ be an indicator that equals 1 if bid i is a player's final bid, and let rep_i equal 1 if the player who placed bid i placed more than one bid. Consider the following probit regression:

$$incr_i = \Phi\left(\beta_1 last^i + \beta_2 rep^i + \gamma X\right)$$

where X are auction covariates. Table 2 displays the results. The coefficient on the dummy for the last bid by a player is negative and significant. This means that the fact that a bid is the last bid by a player makes it less likely that it be an incremental bid. The coefficient on "repeat bidder" is positive and significant.

6.3 Bidding "triggers" bidding of competing players.

In the model (theorem 4), if no player has yet bid, then players wait until time thresholds that depend on their valuation, after which they bid their valuation at the first opportunity. Once a player has placed a bid, however, the other player places a bid equal to the valuation at the first opportunity. Additionally, with more than two players, late enough in the auction, we conjecture that bidding by a player triggers responses by other bidders as well. Therefore, at the end of the auction an earlier bid (be it because of the sniper's realization of an opportunity or his valuation) must make it more likely that other players bid earlier.

Let $\{t_j\}_{j=1}^{N}$ be a set of timings with $t_j^N = 0$. we define a variable $b_j^i = 1$ if in auction i there was at least one bid in the time interval $[t_j, t_{j+1}]$ of the auction. Consider the following probit regression:

$$b_j^{i*} = \Phi\left(\sum_{t=1}^{T} \beta_t b_{j-t}^i + \gamma X\right)$$

where X are auction covariates including the time. If bidding triggers more bidding, we would expect that $\beta_t > 0$. Table shows the results of this estimation using time windows of two and five minutes. The coefficients on b_{j-1}, b_{j-2}, are significant and positive. As some auctions may have a higher probability of bidding activity than others in each period (e.g. because some auctions have more participants), which may lead to positive coefficients on $b_{j-1}, b_{j-2}, \ldots$, we control for the total number of bids in the auction. Table 3 displays

EQUATION	VARIABLES	(1) Incremental bid	(2) Incremental bid
Incremental bid	Player's last bid	-0.330***	-0.324***
		(0.0244)	(0.0243)
	Bidder bids more than once	0.579***	0.698***
		(0.0367)	(0.0354)
	Current price	0.0178***	0.0162***
		(0.00112)	(0.00109)
	Shipping price	-0.00231	-0.000804
		(0.0112)	(0.0111)
	Last 5 seconds	0.172	0.206
		(0.127)	(0.127)
	Last 15 seconds	0.308**	0.337**
		(0.152)	(0.153)
	Last 30 seconds	0.327***	0.357***
		(0.106)	(0.106)
	Last minute	0.309***	0.338***
		(0.0964)	(0.0962)
	Last five minutes	0.412***	0.431***
		(0.0590)	(0.0585)
	Last 20 minutes	0.137***	0.147***
		(0.0465)	(0.0457)
	Last 45 minutes	0.110**	0.122***
		(0.0472)	(0.0465)
	Last hour	0.110*	0.122*
		(0.0661)	(0.0655)
	Earlier than 3 days	-0.106***	-0.113***
		(0.0372)	(0.0369)
	Constant	-0.979***	-1.072***
		(0.0615)	(0.0609)
	Observations	14,503	14,870

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2: Probit: What predicts incremental bidding? (1) auctions with 2 or more players. (2) auctions with 1 or more players.

the results. In the first column, we divide the sample into 15-second periods and consider the final two hours of each auction. t-statistics are shown in parentheses. Bidding in the previous 15-second period, b_{j-1}, and in the period before that, b_{j-2}, are strongly positively associated with the presence of a bid in the current window. We allow a flexible function of time as a control, including a cubic in time remaining in addition to indicators for the final and penultimate periods. In table 9 in the appendix we report results from linear models of the number of bids in each 15-second window. The results are qualitatively the same.

6.4 Comparative statics

Finally, we test comparative statics predictions of the model. Theorem 4 establishes that bidders of valuation v_k begin sniping at their first opportunity after time thresholds $t_i^k(p)$ which are increasing in the current price p and in the valuation v_i^k. Moreover, theorem 4 establishes that when bidders submit bids after their time thresholds, they bid truthfully. As a result, we expect that (1) bids will occur earlier when the current price is higher, and (2) bids of higher amounts will occur earlier conditional on the current price. These predictions are especially informative tests of the model because they emerge from equilibrium play in the model but were not facts used to motivate the model or choose among models.

To test these predictions, we restrict attention to the final minutes of each auction. We define a bid as sniping if it raises the current price by \$5 or more, or is a winning bid, and is the highest bid submitted by a particular bidder. We regress the timing of each sniping bid on the current price, the bid amount, and auction-level covariates. Some subtlety arises from the fact that winning bids are censored. If a bid does not win the auction, we observe the amount, but for the highest bid in each auction we observe only that it is above the second-highest. To allow for censoring, we allow winning bids to have a separate coefficient on bid amount as well as a separate intercept, so that the relationship between bid amount and timing is estimated off of non-winning bids only.[32]

Table 4 displays the results. t-statistics are shown in parentheses. The columns restrict attention to the final 15, 10, 5, 3, and 5 minutes of the auction respectively. Our preferred specifications restrict attention to the final 5 or 3 minutes. The dependent variable is time remaining, measured in minutes.[33] The first two rows show that the current price before the

[32]Conditional on bid amount, whether or not a bid is winning is random from the bidder's point of view as it depends on the highest opposing proxy, which is not observed.

[33]We observe time to the nearest second.

Table 3: Probit: probability of bids as a function of bids in previous time intervals. Sample: all time windows in the final 2 hours of auctions.

	b_j^i	b_j^i	b_j^i
b_{j-1}^i	0.875	1.072	0.959
	(26.76)**	(32.92)**	(27.20)**
b_{j-2}^i	0.535		
	(13.22)**		
time remaining	-0.033	-0.011	-0.033
	(11.80)**	(28.52)**	(11.47)**
$time^2$	0.000		0.000
	(8.02)**		(7.67)**
$time^3$	-0.000		-0.000
	(6.68)**		(6.31)**
final period	1.294		1.308
	(24.96)**		(25.70)**
penultimate period	0.467		0.484
	(8.43)**		(8.84)**
current price	0.002	0.003	0.002
	(2.93)**	(3.78)**	(2.57)*
seller score	-0.000	-0.000	-0.000
	(1.55)	(1.76)+	(1.65)+
shipping price	-0.006		-0.007
	(0.95)		(0.98)
num. bidders	0.041	0.043	0.042
	(12.52)**	(12.80)**	(12.39)**
Constant	-1.751	-1.859	-1.716
	(32.01)**	(48.82)**	(29.78)**
N	127,558	128,639	128,639

+ $p < 0.1$; * $p < 0.05$; ** $p < 0.01$

Probit regressions: dependent variable is 1(bid). Window = 15-second period, sample = final 2 hours of each auction.

Table 4: Comparative statics: timing of bids

minutes remaining	<15	<10	<5	<3	<5
current price	0.055	0.063	0.028	0.017	0.016
	(2.71)**	(5.02)**	(6.02)**	(5.46)**	(2.78)**
bid amount	0.136	0.168	0.431	0.327	
	(0.90)	(0.95)	(3.07)**	(2.94)**	
num. bidders	-0.072	-0.109	-0.025	-0.005	-0.023
	(1.95)+	(4.14)**	(2.24)*	(0.63)	(2.03)*
seller score	-0.000	-0.000	-0.000	-0.000	-0.000
	(5.58)**	(3.77)**	(5.26)**	(5.30)**	(4.94)**
winner	2.845	3.686	13.832	10.501	
	(0.44)	(0.50)	(3.12)**	(2.76)**	
winner*bid amount	-0.211	-0.217	-0.451	-0.333	
	(1.39)	(1.24)	(3.19)**	(2.99)**	
shipping price	-0.142	0.023	-0.002	0.028	0.000
	(1.08)	(0.27)	(0.04)	(1.06)	(0.01)
Constant	1.129	-2.114	-13.294	-10.512	0.217
	(0.18)	(0.29)	(3.00)**	(2.77)**	(0.83)
F statistic	18.7	23.3	12.1	11.9	12.2
R-squared	0.03	0.06	0.03	0.03	0.01
N	929	865	762	715	762

+ $p < 0.1$; * $p < 0.05$; ** $p < 0.01$

dependent variable is minutes remaining. Sample = sniping bids. A bid is sniping if it is the bidder's highest bid and either wins the auction or is at least 5 dollars above the current price. Errors clustered at auction level.

bid, and the amount of the bid, are both positively associated with time remaining, consistent with theorem 4. When we restrict attention to the final 5 or 3 minutes, both coefficients are significant at the 1% level.

7 Conclusions

There are two key empirical patterns present in online proxy auctions with deadlines. First, there is evidence that some players follow heuristic bidding strategies, such as shill bidding. Second, sniping is pervasive across auctions. We present a model that relates these two facts. When there is a possibility that some players are shill bidders, or other heuristic bidding types, we find that sniping must arise. If there is an arbitrarily small but nonzero probability that a player uses a fixed strategy involving small bids, then bids will arrive late. In contrast, when all players are known to behave as rational bidders who value the object and are not colluded with the seller, then the unique prediction is that players bid truthfully as soon as possible. Revision games with reputational types provide a tractable setting that gives these results as unique predictions. Patterns of bid timing and bid amounts in eBay auctions are consistent with our model. In particular, we find support for the comparative statics of bid timing with respect to price and valuation that our model predicts, and for the prediction that bidding triggers other bids, patterns which emerge from the model but were not obvious ex ante.

A Appendix

A.1 Belief updates according to Bayes' rule.

The are three possible cases: either player i placed a proxy at time t, another player placed a bid at that time or no bid was placed.

i) If player i placed a bid, player i's belief regarding histories does not change.

ii) Suppose at time t a bid was placed by player j, the price updated to p and j^w is the winning bidder and let $a = (t, p, j, j^w)$. If $j = j^w$ (that is the player who placed a proxy at time t becomes the leading bidder), the set of possible proxies placed at time t is given by $B_j(a) = \{b \geq p - \Delta\}$. If the proxy placed is below the highest proxy, that is $j \neq j^w$, its value is revealed to all players. The set of possible proxies placed at time t is a singleton that contains this revealed bid $B(a) = \{b^j(t) = p - \Delta\}$. Let $\mu_i(v_j|(h_i^t, a))$ denote the probability that player i's belief assigns to the event that player j's valuation type is v_j. We have,

$$\mu_i(v_j|(h_i^t, a)) = \frac{\mu_i(v_j|h_i^t) \sum_{b \in B(a)} \sigma_j(b|h_i^t, v_j)}{\sum_{v_j \in V} \mu_i(v_j|h_i^t) \sum_{b \in B(a)} \sigma_j(b|h_i^t, v_j)}$$

if the denominator is positive.

iii) If no action is observed at time t, the belief updates as follows. Suppose t_k is the last time before t in which a player placed a bid. If no bid has been placed, let $t_k = T$. We need to condition on the unobservable arrivals the player had in the interval $[t_k, t]$. Suppose player j's bidding opportunities between times $s_0 = t_k$ and t arrived at times $\{s_m\}_{m=1}^{k}$. Let $\nu_i(v_j|h^{s_0}, s_{-1}) = \mu_i(v_j \times V^{n-1} \times \mathscr{H}_{t_k}|h_i^{t_k})$, be the belief that player j is of type v_j at time t_k. The marginal of player i's belief regarding j's type at time s_{m+1} conditional on player j having opportunities at times $\{s_1, \dots, s_{m+1}\}$ is defined recursively by,

$$\nu_i(v_j|h_i^{s_{m+1}}, s_m) := \frac{\nu_i(v_j|h_i^{s_m}, s_{m-1})\sigma_j(b^{\emptyset}|h_i^{s_{m+1}}, v_j)}{\sum_{v \in V^n} \nu_i(v|h_i^{s_m}, s_{m-1})\sigma_j(b^{\emptyset}|h_i^{s_{m+1}}, v)},$$

if the denominator is positive.

Player i's belief that j is of type v_j at time t given the private history h_i^t

$$\mu_i(v_j|h_i^t) = \sum_{j \geq 1} e^{-\lambda(t-t_k)} \int_{t_k \leq s_1 \leq \dots \leq s_j \leq t} \nu_i(v_j|h^{s_0}, s_{-1}) \left(\prod_{m=1}^{j} \frac{\lambda \sigma_j(b^{\emptyset}|h_i^{s_m}, v_j)}{\sum_{v \in V^n} \nu_i(v|h_i^{s_{m-1}}, s)\sigma_j(b^{\emptyset}|h_i^{s_m}, v)} \right) ds_1 \cdots ds_j.$$
$$+ e^{-\lambda(t-t_k)} \nu_i(v_j|h^{s_0}, s_{-1}) \tag{20}$$

Suppose t_k was the time the highest proxy at time was placed by player j. Let $\mu_i(b^H|h_i^t)$ denote the belief regarding the highest proxy at time t. We have,

$$\mu_i(b^H|h_i^t) = \frac{\sum_{v_k \in V} \sigma_k(b^H|h_i^{t_k}, v_k)\mu_i(v_k|h_i^t)}{\sum_{b \geq p-\Delta, b \in B} \sum_{v \in V^n} \sigma_k(b^H|h_i^{t_k}, v_k)\mu_i(v_k|h_i^t)},$$

If the denominator is positive.

A.2 Proofs

A.2.1 Proof of Lemma 1

Proof. Fix $\varepsilon > 0$. First note that an ε-constrained equilibrium puts weight at most ε on a bid above the valuation.[34] Let μ_i denote player i's belief regarding the value of current highest bid and the other players' type at time t. Let b be a random variable that is distributed according

[34] The expected payoff of bidding the valuation is strictly above the payoff from placing a bid strictly above the valuation, since, as all players may tremble, the latter bid results in paying a price for the object strictly above the valuation with positive probability.

to player i's belief about the value of the highest bid at time t. If player i holds the highest bid then i's belief, μ_i, puts probability 1 on i's bid. We assume that if i holds the highest bid, his bid is less than $v_i - \Delta$. We will see that for small t player i prefers to place a truthful bid if an opportunity arises at time t.

Suppose that i does not hold the highest bid at time t. In what follows, we provide a lower bound on the payoff of bidding at time t and an upper bound on the payoff of not bidding at time t. We then show that the former must be above the latter for small t. Thus, if player i does not hold the highest bid, he prefers to place a bid at time t, for small enough t.

Let p denote the current price at time t. Suppose $v_i > p + \Delta$. Let $\mathbb{P}_{\mu_i}(b < v_i - \Delta)$ denote the probability that highest bid is less than $v_i - \Delta$ according to i's belief. The payoff from placing a bid $v_i - \Delta$ is at least

$$e^{-\lambda(n-1)t}\mathbb{P}_{\mu_i}(b < v_i - \Delta) \cdot (v_i - \mathbb{E}_{\mu_i}(b | b < v_i - \Delta) - \Delta). \tag{21}$$

To understand the expression in (21) note that $e^{-\lambda(n-1)t}$ is the probability that players other than i don't get another opportunity to bid by the end of the auction, in which event i obtains expected payoff $(v_i - \mathbb{E}_{\mu_i}(b|b < v_i - \Delta) - \Delta)$ when $b < v_i - \Delta$.

The payoff of not bidding is at most

$$(1 - e^{-\lambda t})\mathbb{P}_{\mu_i}(b < v_i - \Delta) \cdot (v_i - \mathbb{E}_{\mu_i}(b | b < v_i - \Delta) - \Delta). \tag{22}$$

$(1 - e^{-\lambda t})$ is the probability that i gets another opportunity. If i places a bid at time t he obtains at most (assuming opposing players do not place another bid) payoff $(v_i - \mathbb{E}_{\mu_i}(b|b < v_i - \Delta) - \Delta)$ when $b < v_i - \Delta$.

First note that $\mathbb{P}_{\mu_i}(b < v_i - \Delta) > 0$ since in an ε-constrained equilibrium all bids above $p + \Delta$ could have been placed as a "tremble" by the player who placed the current highest bid. Second, for small enough t, $e^{-\lambda(n-1)t} + e^{-\lambda t} \geq 1$. Thus, for t small enough (21)>(22) and i prefers to outbid competing players and puts probability weight of ε on not placing a bid ($b^{\emptyset}$).

Next we show that player i will not place a bid $b' < v_i - \Delta$. Consider the event in which b is between b' and $v_i - \Delta$, i's payoff from bidding $v_i - \Delta$ is greater than the payoff from bidding

$b' < v_i - \Delta$ by at least

$$\left[e^{-\lambda(n-1)t}(v_i - \mathbb{E}_{\mu_i}(b|b \in [b', v_i - \Delta)) - \Delta) - (1 - e^{-\lambda t})(v_i - \mathbb{E}_{\mu_i}(b|b \in [b', v_i - \Delta)) - \Delta))\right]. \tag{23}$$

To understand equation (23) note that, in the worst case, bidding $v_i - \Delta$ only gives a positive payoff in the event that no other player gets an arrival. In the event that $b > b'$, after bidding b' at time t, i does not hold the highest bid. Therefore, in the best case, after placing bid b' i wins with probability one if he gets another bidding opportunity before the auction ends. Note that for small enough t (23) is strictly positive if $e^{-\lambda(n-1)t} + e^{-\lambda t} \geq 1$.

Consider the event in which the highest bid by opposing players, b, is at or below b'. If players other than i do not get another arrival, bids b' and $v_i - \Delta$ give the same payoff. For b' to be a better choice than $v_i - \Delta$, competing players must play differently upon realizing that i placed the higher bid. However, since a player $j \neq i$ does not observe whether i bid b' or $v_i - \Delta$ unless some player in $-i$ outbids b', equilibrium play after bid b' can only differ from play after bid $v_i - \Delta$ at bidding opportunities after b' has been outbid. Now, the moment the highest bid is raised above $v_i - \Delta$, b' and $v_i - \Delta$ give the same payoff since i cannot outbid again. Therefore, in order to compare the expected payoff from bids $v_i - \Delta$ and b' we condition on the first time the highest opposing bid falls between b' and $v_i - \Delta$ (after which players could punish a higher bid).

As before, we construct a bound for the difference in payoff between bidding $v_i - \Delta$ and b' and show that for small t this difference must be positive. The worst case for bid $v_i - \Delta$ is as follows: after the first time a bid in $[b', v_i)$ arrives, the opposing player who obtains the next bidding opportunity outbids i with probability one. On the other hand, in the best case scenario, bidding b' induces other players to not place any further bids and i wins conditional on obtaining another opportunity to bid.

Let b^k denote a random variable that is distributed as the k'th next highest bid by opposing players given i's beliefs after placing bid b' at time t. The difference in i's expected payoff from bidding $v_i - \Delta$ and b' in the event that $b' < b$ is at least

$$\sum_{k \geq 1} \mathbb{P}(b^k \in [b', v_i], b^{k-1} < b') \Big[e^{-\lambda(n-1)t}(v_i - \mathbb{E}_{\mu_i}(b^k | b^k \in [b', v_i], b^{k-1} < b') - \Delta) \tag{24}$$
$$- \left(1 - e^{-\lambda t}\right)(v_i - \mathbb{E}_{\mu_i}(b^k | b^k \in [b', v_i], b^{k-1} < b') - \Delta) \Big].$$

The previous expression is derived as follows. The aggregate arrival rate of all of opposing players is $\lambda(n-1)$. From the previous discussion, the expected highest bid of the opposing players does not depend on i's bid as long as it hasn't been raised above b'. For this reason in (24) we condition on the time that highest bid by players in $-i$ falls in $[b', v_i]$ and assume that a soon as a bid of $v_i - \Delta$ by i is "found out," he is punished at the next opportunity of any opposing player. For b' to be as convenient as possible, we assume that after b' is outbid other players do not outbid i's bid. Thus, player i gets a higher payoff by bidding $v_i - \Delta$ in the event that the k'th next highest bid is the first to be between $[b', v_i]$ and other players do not get more arrivals, which occurs with probability at least $e^{-\lambda(n-1)t}$.[35] In this case, a b' bid would have given i a zero payoff. On the other hand, if i has at least one more arrival, which occurs with probability of at most $(1-e^{-\lambda t})$, after bid b' is outbid i's opponents abstain from placing any other bids and i outbids b at the next opportunity. The expression in equation (24) is strictly positive for small enough t such that $e^{-\lambda(n-1)t} + e^{-\lambda t} \geq 1$.

Now, the expected difference between bidding v_i and b' is

$$\mathbb{P}_{\mu_i}(b \in [b', v_i - \Delta)) \cdot (23) + \mathbb{P}_{\mu_i}(b < b') \cdot (24)$$

which, by the previous discussion, and since $\mathbb{P}_{\mu_i}(b \in [b', v_i - \Delta)) > 0$ and $\mathbb{P}_{\mu_i}(b < b') > 0$, is strictly positive for small enough t. Thus, in an ε-constrained equilibrium player i places the least possible probability weight, ε, on bids $b' < v_i$ for $t < t_0(n)$ with $e^{-\lambda(n-1)t_0(n)} + e^{-\lambda t_0(n)} = 1$. $\square$

A.2.2 Proof of Proposition 2

Proof. By contradiction, suppose that $\bar{t} > 0$ is the supremum time from which players put as much weight as possible on bidding truthfully in every ε-constrained equilibrium. We will show that in a ε-constrained equilibrium players must bid their valuation when possible from a time threshold $\bar{t} + \tau$ for some $\tau > 0$ that does not depend on ε. Equilibrium play after $\bar{t}$ does not depend on players' actions in previous opportunities. Let b denote the highest bid by players other than i and suppose $b < v_i - \Delta$.[36] Let $V_i(b', b, \bar{t})$ denote the payoff from time $\bar{t}$, if i holds bid b'. Let $\mu_i^{\bar{t}+\tau}$ denote i's belief at time $\bar{t} + \tau$ over histories and opponents' types. Suppose that player i does not hold the winning bid. We will now show that the

[35] Note that the k'th highest bid arrives at some time $\tau < t$. The probability that there is no additional arrival by players in $-i$ after time τ is $e^{-\lambda(n-1)\tau}$. Thus, $e^{-\lambda(n-1)t}$ is a lower bound on the probability that players in $-i$ do not receive another opportunity after the k'th next highest bid is placed.

[36] If $b > v_i - \Delta$ i's payoff is zero regardless of his actions.

expected payoff from holding bid $v_i - \Delta$ at time $\bar{t}$ is strictly above the payoff of holding a bid $b' < v_i - \Delta$ in any ε-constrained equilibrium. In fact, we have,

$$\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(v_i - \Delta, b, \bar{t})|b < v_i - \Delta) = (1 - e^{-\lambda \bar{t}})\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(b', b, \bar{t})|b < v_i - \Delta). \tag{25}$$

That is, the expected value of holding a losing bid at time $\bar{t}$ from the perspective of the belief at time $t + \tau$ is equal to the expected value of holding the winning bid at time $\bar{t}$ times the probability that i receives another opportunity. This follows from the fact that equilibrium play does not depend on i's actions after time $\bar{t}$ and that i gets zero payoff if he does not have an opportunity to place another bid. Now, if player i were to bid $v_i - \Delta$ at time $\bar{t} + \tau$, in the worst case for bid $v_i - \Delta$, opposing players punish i with their actions between times $\bar{t} + \tau$ and $\bar{t}$, and place a bid above $v_i - \Delta$. By assumption, after $\bar{t}$ play does not depend on i's actions. Thus, the payoff from bidding $v_i - \Delta$ is at least,

$$e^{-\lambda(n-1)\tau}\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(v_i - \Delta, b, \bar{t})|b < v_i - \Delta)\mathbb{P}(b < v_i - \Delta), \tag{26}$$

where $e^{-\lambda(n-1)\tau}$ is the probability that other players don't get an opportunity from $\bar{t} + \tau$ until $\bar{t}$. On the other hand, in the best case for bid b', opposing players do not place any further bids between times $\bar{t} + \tau$ and $\bar{t}$ and i holds bid $v_i - \Delta$ at time $\bar{t}$ if he gets and additional bidding opportunity between $\bar{t} + \tau$ and $\bar{t}$. Thus, the payoff of placing bid b' is at most

$$\left[(1 - e^{-\lambda\tau})\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(v_i - \Delta, b, \bar{t})|b < v_i - \Delta) + e^{-\lambda\tau}\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(b', b, \bar{t})|b < v_i - \Delta)\right]\mathbb{P}(b < v_i - \Delta). \tag{27}$$

The first term is the payoff arising from the event in which i gets an opportunity in $[\bar{t}, \bar{t} + \tau]$. The second term is the payoff from the event that i does not get an opportunity. From equation (25) equation (26) is greater than equation (27) if

$$e^{-\lambda(n-1)\tau}\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(v_i - \Delta, b, \bar{t})|b < v_i - \Delta) > \left((1 - e^{-\lambda\tau}e^{-\lambda\bar{t}}\right)\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(v_i - \Delta, b, \bar{t})|b < v_i - \Delta).$$

Note that the expected payoff $\mathbb{E}_{\mu_i^{\bar{t}+\tau}}(V_i(v_i - \Delta, b, \bar{t})|b < v_i - \Delta)$ is strictly greater than zero in an ε-constrained equilibrium. Thus, for small enough τ, independent of ε, player i prefers to place bid $v_i - \Delta$ instead of b'.

Now, if i holds the highest bid at time $\tau + \bar{t}$ at $b' < v_i - \Delta$, he would also prefer to update his bid to $v_i - \Delta$. He would do so because all the opposing players j are bidding $v_j - \Delta$ at the first opportunity. For any belief that player i may have about the other players valuations there is the possibility that a player may tremble and place a bid between b' and $v_i - \Delta$. Thus, while bids b' and $v_i - \Delta$ give the same payoff if i receives an opportunity after $\bar{t} + \tau$, i is strictly worse off with bid b' in the event that he does not get another opportunity before the end of the auction.

In conclusion, for τ small enough, there is $\bar{\varepsilon} > 0$, such that for every $\varepsilon < \bar{\varepsilon}$, the payoff of bidding $v_i - \Delta$ is strictly above the payoff of not holding a winning bid. Thus, player i puts the least possible weight, ε, on bids other than $v_i - \Delta$ at every opportunity between $\bar{t} + \tau$ and $\bar{t}$. Because this is true for every ε and every player i in a trembling hand perfect equilibrium players must bid their valuation minus the increment from time $\bar{t} + \tau$ on. This observation contradicts that $\bar{t}$ is the infimum time at which players start placing truthful bids. □

A.2.3 Proof of Lemma 3

1. Follows from the results in section 4 since a shill bidder does not tremble.

2. Once player 1 reveals rationality equilibrium play reverts to the unique THPE that exists when players are rational. Thus, from the results in section 4 player 1 bids truthfully at the first opportunity.

A.2.4 Proof of Proposition 3

The existence of equilibrium follows from Moroni (2014). Let $\tilde{\sigma}_1$ denote a strategy of player 1 in which she bids $v - \Delta$ at the first arrival after time $t_1 = 1/\lambda$ and does not outbid player 2 before time t_1. Let $\tilde{\sigma}_2$ denote the strategy of player 2 described in Proposition 3.

Since Player 1 is bidding $v - \Delta$, player 2 cannot gain from outbidding the rational player 1 after 1 has placed her bid. For a fixed $\gamma_t > 0$, as the probability that player 1 trembles tends to zero, Player 2 behaves as if best responding to a shill bidder commitment type.

Now, suppose that player 2 has placed a bid of $p + 2\Delta$, and equilibrium play is such that player 2 does not bid $v - \Delta$ unless player 1 outbids. The payoff of bidding v at time t for player 1 is

$$e^{-\lambda t}(v - p - \Delta) \tag{28}$$

The payoff of waiting until the next opportunity to bid $v-\Delta$ for Player 1 is given by,

$$\lambda t e^{-\lambda t}(v-p-\Delta) \tag{29}$$

Thus, player 1 will prefer to bid $v-\Delta$ after time $t_1 = 1/\lambda$ and not before, conditional on the proposed play for player 2 so that the proposed equilibrium strategies do not have a one-shot deviation.

Let's see that the equilibrium play in every trembling hand perfect equilibrium is as in $(\tilde{\sigma}_1, \tilde{\sigma}_2)$. In fact, we know that after a sufficiently late time threshold player 1 will bid $v-\Delta$ at the first bidding opportunity. By contradiction, suppose there is a supremum time $\bar{t}$ such that Player 1 plays according to $\tilde{\sigma}_1$. Note that slightly earlier than $\bar{t}$ player 2's best response would also be to not raise a bid of $p+2\Delta$. The probability of facing a "punishment" from player 1 is small for times near $\bar{t}$. After $\bar{t}$, 2 is strictly better off best responding to the shill bidder type given 1's strategy after $\bar{t}$. Thus, there is $\bar{\tau}$ such that player 2 behaves according to $\tilde{\sigma}_2$ at times $\bar{t}+\tau$ with $\tau \leq \bar{\tau}$. If $\tau+\bar{t} < t_1$, then by (28) and (29), at an opportunity that arrives at a time in $[\bar{t}, \bar{t}+\tau)$ player 1 would optimally follow $\tilde{\sigma}_1$. This is a contradiction. Player 1 must behave according to $\tilde{\sigma}_1$ to which player 2 best responds with $\tilde{\sigma}_2$.

Now, the time at which Player 2 places the first bid can be determined noting that the payoff of bidding $p+2\Delta$ at time t is

$$\gamma_t(e^{-\lambda t}(v-p-\Delta)+(1-e^{-\lambda t})(v-p-2\Delta))+(1-\gamma_t)(v-\mathbb{E}_{\mu_2}(b|b\leq p+\Delta)-\Delta)e^{-\lambda \bar{t}} \tag{30}$$

Payoff of waiting until next opportunity to bid is

$$\gamma_t(\lambda t e^{-\lambda t}(v-p-\Delta)+(1-\lambda t e^{-\lambda t}-e^{-\lambda t})(v-p-2\Delta))+(1-\gamma_t)(v-\mathbb{E}_{\mu_2}(b|b\leq p+\Delta)-\Delta)e^{-\lambda \bar{t}}(1-e^{-\lambda t}) \tag{31}$$

where $\mathbb{E}_{\mu_2}(b|b<p+\Delta)$ denotes the expected bid of player 1 which must be less than or equal to $p+\Delta$.
For small enough γ_t (30) is greater than (31).

Since an equilibrium exists and in order for the strategies to not have a profitable one shot deviation, given how play must be when the end of the auction is sufficiently close, equilibrium play in a trembling hand perfect equilibrium must be as in $(\tilde{\sigma}_1, \tilde{\sigma}_2)$.

A.2.5 Proof of Proposition 4

Proof. By an argument similar to the one in the same valuation equal to v case a player with valuation v^0 would not want to bid the valuation unless outbid. This implies that the latest player 1 of type v^0 bids is if all types of player 2 other than v^0 would place a bid as soon as they get a chance. The payoff of bidding at time t is given by the probability that player 1 does not get another arrival times the valuation minus the price paid.

$$e^{-\lambda t}(v^0 - p - \Delta)$$

The payoff of bidding of waiting until the next opportunity to bid is given by

$$\mathbb{P}_{\mu_1(t)}(v_1 = v^0)\lambda t e^{-\lambda t}(v^0 - p - \Delta) + (1 - \mathbb{P}_{\mu_1(t)}(v_1 = v^0))e^{-\lambda t}(1 - e^{-\lambda t})(v^0 - p - \Delta)$$

That is, if facing the type v^0 player, player 1 only gets a positive payoff if once his bid is placed player 1 does not get another arrival. This event has probability $\lambda t e^{-\lambda t}$. In the case in which player 1 is facing a player of valuation higher than v^0, player 1 will get a positive payoff if he gets another arrival to bid and player 2 does not get another arrival. Whenever t is such that $e^{-\lambda t} \leq p_0 \leq \mathbb{P}_{\mu_1(t)}(v_2 = v_0)$ and $\lambda t \geq 2$ the payoff of waiting until the next opportunity to bid is higher than the payoff of bidding at time t. Thus, player 1 of type v^0 will place a bid starting at a time later than $\bar{t}_1^0$.

Recursively, for player 1 of type v^k an upper bound on the timing of bidding $v^k - \Delta$ is obtained by assuming that all types v^k or lower of player 2 are bidding their valuation starting at the time $\bar{t}_2^k$ unless they are outbid at which point they would outbid right away.

The payoff of bidding in this case is

$$e^{-\lambda t}(v^k - p - \Delta) + (1 - e^{-\lambda t})(v^k - \mathbb{E}_{\mu_1(t)}(v_2 | v_2 \leq v^k))\mathbb{P}_{\mu_1(t)}(v_2 \leq v^k). \tag{32}$$

The payoff of waiting for the next opportunity to bid is given by

$$e^{-\lambda t}(1-e^{-\lambda t})(v^k-p-\Delta)\mathbb{P}_{\mu_1(t)}(v_2>v^k)+\mathbb{P}_{\mu_1(t)}(v_2\leq v^k)\times$$
$$\Big(\lambda(t-\bar{t}_2^k)e^{-\lambda t}(\mathbb{E}_{\mu_1(t)}(v_2|v_2\leq v^k)-p-\Delta)+(1+e^{-\lambda t}e^{-\lambda\bar{t}_2^k}-2e^{-\lambda t})(v^k-\mathbb{E}_{\mu_1(t)}(v_2|v_2\leq v^k))\Big). \tag{33}$$

Equating both equations and re-arranging one obtains

$$\lambda(t-\bar{t}_2^k)(\mathbb{E}_{\mu_1(t)}(v_2|v_2 \leq v^k)-p-\Delta)\mathbb{P}_{\mu_1(t)}(v_2\leq v^k)=\mathbb{P}_{\mu_1(t)}(v_2\leq v^k)\big((1-e^{-\lambda t})(v^k-p-\Delta)+$$
$$(1-e^{-\lambda\bar{t}_2^k})(v^k-\mathbb{E}_{\mu_1(t)}(v_2|v_2\leq v^k))\big)+e^{-\lambda t}(v^k-p-\Delta).$$

Which noting that if $e^{-\lambda t}\leq p_0\leq\mathbb{P}(v_2\leq v^k)$ this equations implies that the time $\bar{t}$ that equates (32) and (33) is such that

$$\lambda(\bar{t}-\bar{t}_2^k)\leq\frac{(2-e^{-\lambda\bar{t}})(v^k-p-\Delta)+(1-e^{-\lambda t_2^k})(v^k-\mathbb{E}_{\mu_1(\bar{t})}(v_2|v_2<v^k))}{\mathbb{E}_{\mu_1(\bar{t})}(v_2|v_2<v^k)-p-\Delta}.$$

Therefore $\lambda\bar{t}\leq\lambda\bar{t}_k$ and the time player 1 of type v^k starts bidding is such that $t_1^k\leq\bar{t}$ since other players may start to bid later than $\bar{t}_1^k$ in equilibrium, which creates incentives for player 1 to bid later.

For player 2 of type v^j, if current price is $p+\Delta$, the payoff of bidding at time t is given by

$$e^{-\lambda t}(v^j-p-\Delta)+(1-e^{-\lambda t})(v^j-\mathbb{E}_{\mu_2(t)}(v_1|v_1<v^j))\mathbb{P}_{\mu_2(t)}(v_1<v^j). \tag{34}$$

The payoff of bidding at the next bidding opportunity is at least

$$e^{-\lambda t}(v^j-p-\Delta)+\mathbb{P}_{\mu_2(t)}(v_1<v^j)\Big(\lambda(t-\bar{t}_1^j)e^{-\lambda t}(\mathbb{E}_{\mu_2(t)}(v_1|v_1 < v^j)-p-\Delta)+$$
$$(1+e^{-\lambda t}e^{-\lambda\bar{t}_1^j}-2e^{-\lambda t})(v^j-\mathbb{E}_{\mu_2(t)}(v_1|v_1\leq v^j))\Big).$$

Equating both equations we obtain

$$\lambda(t-\bar{t}_1^k)=\frac{(v^j-\mathbb{E}_{\mu_2(t)}(v_1|v_1\leq v^j))(1-e^{-\lambda\bar{t}_1^j})}{\mathbb{E}_{\mu_2(t)}(v_1|v_1\leq v^j)-p-\Delta}.$$

If the current price is p (that is player 1 has not bid yet), the payoff of bidding at the next bidding opportunity is given by

$$e^{-\lambda t}(v^j - p) + \mathbb{P}_{\mu_2(t)}(v_1 < v^j)\Big(\lambda(t-\bar{t}_1^j)e^{-\lambda t}(\mathbb{E}_{\mu_2(t)}(v_1|v_1 < v^j) - p) + \frac{\left(\lambda(t-\bar{t}_1^j)\right)^2}{2!}e^{-\lambda t}(\mathbb{E}_{\mu_2(t)}(v_1|v_1 \le v^j) - p - \Delta) + (1 + e^{-\lambda t}e^{-\lambda \bar{t}_1^j} - 2e^{-\lambda t})(v^j - \mathbb{E}_{\mu_2(t)}(v_1|v_1 \le v^j))\Big). \tag{35}$$

The payoff of bidding $v^j - \Delta$ at time t is given by

$$e^{-\lambda t}(v^j - p) + \Big(\lambda(t-\bar{t}_1^j)e^{-\lambda t}(v - p - \Delta) + (1 - e^{-\lambda(t-\bar{t}_1^j)} - e^{-\lambda(t-\bar{t}_1^j)}\lambda(t-\bar{t}_1^j)) + \lambda(t-\bar{t}_1^j)(1-e^{-\lambda\bar{t}_1^j}) + e^{-\lambda(t-\bar{t}_1^j)}(1-e^{-\lambda\bar{t}_1^j})\Big)(v^j - \mathbb{E}_{\mu_2(t)}(v_1|v_1 \le v^j))\Big)\mathbb{P}_{\mu_2(t)}(v_1 < v^j) =$$

$$e^{-\lambda t}(v^j - p) + \left(\lambda(t-\bar{t}_1^j)e^{-\lambda t}(\mathbb{E}_{\mu_2(t)}(v_1|v_1 < v^j) - p - \Delta) + (1-e^{-\lambda t})(v^j - \mathbb{E}_{\mu_2(t)}(v_1|v_1 < v^j))\right)\mathbb{P}_{\mu_2(t)}(v_1 < v^j). \tag{36}$$

The time that equates equations (35) and (36) is such that

$$\lambda(t-\bar{t}_1^j)\Delta + \frac{\left(\lambda(t-\bar{t}_1^j)\right)^2}{2!}(\mathbb{E}_{\mu_2(t)}(v_1|v_1 < v^j) - p - \Delta) = (v^j - \mathbb{E}_{\mu_2(t)}(v_1|v_1 \le v^j))(1-e^{-\lambda\bar{t}_1^j})$$

which implies (12).

Thus, we see that for any type of player 2 with valuation above v^0 for γ_t sufficiently close to zero the player would prefer to bid $v_2 - \Delta$ than to bid something lower sufficiently late in the auction.

Let's see now that the time thresholds are increasing in the valuation of the player and decreasing in the price of the auction. The reason why the time thresholds are increasing in the valuation of the player and in the price is that for any given time the payoff from bidding earlier is greater the greater the valuation of the player. In fact, in equilibrium, suppose valuation types $v^m \leq v^k$ of player 2 are starting to bid later than time t, while types of player 2 such that $v^m > v^k$ are already bidding.

The payoff from bidding at time t is given by,

$$e^{-\lambda t}(v^k - p - \Delta) + (1 - e^{-\lambda t})\left(v^k - \mathbb{E}_{\mu_1(t)}(v_2 | v_2 \leq v^k)\right)\mathbb{P}(v_2 \leq v^k) \tag{37}$$

The payoff of waiting until the next opportunity to bid is given by

$$\begin{aligned} e^{-\lambda t}(1 - e^{-\lambda t})(v^k - p - \Delta)\mathbb{P}(v_2 > \quad & v^k) + \sum_{m \leq k} \mathbb{P}_{\mu_1(t)}(v_2 = v^m) \times \\ \left[\lambda(t - t_2^m)e^{-\lambda t}(v^m - p - \Delta) + \quad \right. & \left. (1 + e^{-\lambda t_2^m} e^{-\lambda t} - 2e^{-\lambda t})(v^k - v^m)\right] \end{aligned} \tag{38}$$

Equating both equations we find that the timing t_1^k that makes player 1 indifferent between bidding and not bidding is such that

$$\begin{gathered} \lambda t_1^k \mathbb{E}_{\mu_1(t_1^k)}(v_2 - p - \Delta | v_2 \leq v^k)\mathbb{P}_{\mu_1(t_1^k)}(v_2 \leq v^k) = \\ \sum_{m \leq k}\left(\lambda t_2^m(v^m - p - \Delta) + (1 - e^{-\lambda t_2^m})(v^k - v^m)\right)\mathbb{P}_{\mu_1(t_1^k)}(v_2 = v^m) + \\ (e^{-\lambda t_1^k}(v^k - p - \Delta) + \mathbb{P}_{\mu_1(t_1^k)}(v_2 \leq v^k)(1 - e^{-\lambda t_1^k})(v^k - p - \Delta)) \end{gathered} \tag{39}$$

Note that t_1^k is increasing in p. Let t_1^k denote the timing that makes equation (39) true. For a player 1 of type v^{k+1} the payoff of bidding at time t is given by

$$e^{-\lambda t}(v^{k+1} - p - \Delta) + (1 - e^{-\lambda t})\left(\left(v^{k+1} - \mathbb{E}_{\mu_{1(t)}}(v_2 | v_2 \leq v^k)\right)\mathbb{P}_{\mu_1(t)}(v_2 \leq v^k) + \left(v^{k+1} - v^k\right)\mathbb{P}_{\mu_1(t)}(v_2 = v^k)\right)$$

The payoff from waiting until the next bidding opportunity is given by

$$\begin{aligned} e^{-\lambda t}(1 - e^{-\lambda t})(v^{k+1} - p - \Delta)\mathbb{P}_{\mu_1(t)}(v_2 \quad > \quad & v^k) + \sum_{m \leq k} \mathbb{P}_{\mu_1(t)}(v_2 = v^m)\left[\lambda(t - t_2^m)e^{-\lambda t}(v^m - p - \Delta)\right. \\ + \quad & \left.(1 + e^{-\lambda t_2^m} e^{-\lambda t} - 2e^{-\lambda t})(v^{k+1} - v^m)\right] \end{aligned}$$

Equating both equations we obtain

$$\begin{aligned}\lambda t\mathbb{E}_{\mu_1(t)}(v_2-p-\Delta|v_2\leq v^k)\mathbb{P}_{\mu_1(t)}(v_2\leq v^k) =\\ \sum_{m\leq k}\left(\lambda t_2^m(v^m-p-\Delta)+(1-e^{-\lambda t_2^m})(v^{k+1}-v^m)\right)\mathbb{P}_{\mu_1(t)}(v_2=v^m)+\\ (e^{-\lambda t}(v^{k+1}-p-\Delta)+\mathbb{P}_{\mu_1(t)}(v_2<v^k)(1-e^{-\lambda t})(v^{k+1}-p-\Delta))+\\ \left(v^{k+1}-v^k\right)\mathbb{P}_{\mu_1(t)}(v_2=v^k)\end{aligned} \tag{40}$$

At t_1^k the right hand side is strictly above the left hand side of equation (40) thus, the timing at which a player 1 of type v^{k+1} bids must be earlier in the auction than the timing at which a player 1 of type v^k bids. Now for player 2, if at time t types of player 1 with valuation $v^m<v^j$ are not bidding until after times t_1^m in equilibrium and types of player 1 with valuation $v^m\geq v^j$ the payoff from bidding at time t is given by (34). The payoff from waiting until the next opportunity to bid is given by,

$$\begin{aligned}e^{-\lambda t}(v^j-p-\Delta)+\sum_{m<j}\mathbb{P}_{\mu_2(t)}(v_1 &= v^m)\left[\lambda(t-t_1^m)e^{-\lambda t}(v^m-p-\Delta)+\right.\\ &\left.(1+e^{-\lambda t_1^m}e^{-\lambda t}-2e^{-\lambda t})(v^j-v^m)\right]\end{aligned}$$

The payoff of bidding at time t_2^j and the payoff of waiting until the next bidding opportunity are equal if

$$\lambda t_2^j\mathbb{E}_{\mu_2(t_2^j)}(v_1-p-\Delta|v_1<v^j)\mathbb{P}_{\mu_2(t_2^j)}(v_1<v^j)=\sum_{v^m<v^j}\mathbb{P}_{\mu_2(t_2^j)}(v_1=v^m)\left[\lambda t_1^m(v^m-p-\Delta)+(1-e^{-\lambda t_1^m})(v^j-v^m)\right] \tag{41}$$

Let t_2^j denote the time that makes equation (41). If player 2 was of type v^{j+1} the difference between bidding at the next arrival and bidding at time t is given by

$$\lambda t\mathbb{E}_{\mu_2(t)}(v_1-p-\Delta|v_1<v^j)\mathbb{P}(v_1<v^j)-\sum_{v^m<v^j}\mathbb{P}_{\mu_2(t)}(v_1=v^m)\left[\lambda t_1^m(v^m-p-\Delta)+(1-e^{-\lambda t_1^m})(v^{j+1}-v^m)\right]$$

The last expression is negative for $t=t_2^j$ and therefore, the timing at which a player 2 of type v^{j+1} bids must be earlier than t_2^j.

Assuming that for each player i, players whose valuation is below i's valuation start bidding later, while players whose valuation is above i are bidding truthfully at the first opportunity, we have established that the timing at which players bid is increasing with their own valuation type if at the time t in which they bid. Let's see now that the bidding time thresholds must be such that the assumptions hold. Let's first see that $0 = t_2^0 < t_1^0 < t_2^1 < t_1^1$.

Suppose $t_2^m < t_1^m$ for $m \in \{1, \dots, r\}$ and $t_1^{r+1} \leq t_2^{r+1}$. In analogy to (41), the time threshold t_2^{r+1} at which player 2 of type v^{r+1} bids is such that

$$\lambda t_2^{r+1} \mathbb{E}_{\mu_2(t_2^{r+1})}(v_1 - p - \Delta | v_1 \leq v^{r+1}) \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 \leq v^{r+1}) = \sum_{v^m \leq v^{r+1}} \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^m) \left[\lambda t_1^m (v^m - p - \Delta) + (1 - e^{-\lambda t_1^m})(v^{r+1} - v^m) \right] \tag{42}$$

t_1^{r+1} must satisfy

$$\begin{aligned} \lambda t_1^{r+1} \quad & \mathbb{E}_{\mu_1(t_1^{r+1})}(v_2 - p - \Delta | v_2 < v^{r+1}) \mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^{r+1}) = \\ & \sum_{v^m < v^{r+1}} \left(\lambda t_2^m (v^m - p - \Delta) + (1 - e^{-\lambda t_2^m})(v^{r+1} - v^m) \right) \mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^m) + \\ & (e^{-\lambda t_1^{r+1}}(v^{r+1} - p - \Delta) + (1 - e^{-\lambda t_1^{r+1}}) \mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^{r+1})(v^{r+1} - p - \Delta)) \end{aligned} \tag{43}$$

Since $t_2^r < t_1^r$ we know, using equations (39) and (41), that

$$\begin{aligned} & \Bigg(\sum_{m \leq r} \left(\lambda t_2^m (v^m - p - \Delta) + (1 - e^{-\lambda t_2^m})(v^r - v^m) \right) \mathbb{P}_{\mu_1(t_1^r)}(v_2 = v^m) + \\ & e^{-\lambda t_1^r}(v^r - p - \Delta) + \mathbb{P}_{\mu_1(t_1^r)}(v_2 \leq v^r)(v^r - p - \Delta) - \lambda t_1^r \mathbb{P}_{\mu_1(t_1^r)}(v_2 = v^r)(v^r - p - \Delta) \Bigg) \times \\ & \frac{1}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} \\ & \geq \left(\sum_{m < r} \mathbb{P}_{\mu_2(t_2^r)}(v_1 = v^m) \left(\lambda t_1^m (v^m - p - \Delta) + (1 - e^{-\lambda t_1^r})(v^r - v^m) \right) \right) \frac{1}{\mathbb{P}_{\mu_2(t_2^r)}(v_1 < v^r)} \end{aligned} \tag{44}$$

Let

$$\star = \lambda t_1^{r+1} \mathbb{E}_{\mu_1(t_1^{r+1})}(v_2 - p - \Delta | v_2 < v^r)$$

and let

$$\star\star = \lambda t_2^{r+1}\mathbb{E}_{\mu_2(t_2^{r+1})}(v_1 - p - \Delta | v_1 < v^r)$$

If we prove that $\star > \star\star$, we would obtain a contradiction with $t_1^{r+1} > t_2^{r+1}$ since $\mathbb{E}_{\mu_1(t_1^{r+1})}(v_2 - p - \Delta | v_2 < v^{r+1}) = \mathbb{E}_{\mu_2(t_2^{r+1})}(v_1 - p - \Delta | v_1 < v^{r+1})$ by symmetry.

We have

$$\begin{aligned}
\star \quad = \quad & \Bigg(\sum_{m\le r}\Big(\lambda t_2^m(v^m - p - \Delta) + (1 - e^{-\lambda t_2^m})(v^r - v^m)\Big)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^m) + \\
& \sum_{m\le r}(1 - e^{-\lambda t_2^m})(v^{r+1} - v^r)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^m) \\
& + e^{-\lambda t_1^{r+1}}(v^{r+1} - p - \Delta) + \mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 \le v^r)(v^{r+1} - p - \Delta)(1 - e^{-\lambda t_1^{r+1}}) \\
& - \lambda t_1^{r+1}(v^r - p - \Delta)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^r)\Bigg)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
\ge \quad & \sum_{m<r}\mathbb{P}_{\mu_2(t_2^r)}(v_1 = v^m)(\lambda t_1^m(v^m - p - \Delta) + (1 - e^{-\lambda t_1^m})(v^r - v^m))\frac{1}{\mathbb{P}_{\mu_2(t_2^r)}(v_1 < v^r)} + \\
& - e^{-\lambda t_1^r}(v^r - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} - (v^r - p - \Delta)(1 - e^{-\lambda t_1^r})\frac{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^{r+1})}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} \\
& + e^{-\lambda t_1^{r+1}}(v^{r+1} - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} + (v^{r+1} - p - \Delta)(1 - e^{-\lambda t_1^{r+1}})\frac{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^{r+1})}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& + \sum_{m\le r}(1 - e^{-\lambda t_2^m})(v^{r+1} - v^r)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^m)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& - \lambda t_1^{r+1}(v^r - p - \Delta)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^r)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& + \lambda t_1^r\mathbb{P}_{\mu_1(t_1^r)}(v_2 = v^r)(v^r - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)}
\end{aligned} \tag{45}$$

Now, rewriting we have,

$$\begin{aligned}\star\star \;=\; & \Bigg(\sum_{v^m \le v^{r+1}} \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^m)\Big[\lambda t_1^m (v^m - p - \Delta) + (1 - e^{-\lambda t_1^m})(v^{r+1} - v^m)\Big] - \\ & \lambda t_2^{r+1}(v^{r+1} - p - \Delta)\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^{r+1}) - \lambda t_2^{r+1}(v^r - p - \Delta)\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^r)\Bigg) \times \\ & \frac{1}{\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 < v^r)}\end{aligned}$$

And from (45) we have

$$
\begin{aligned}
\star \geq\ & \star\star + \Bigg(-\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^r)\lambda t_1^r (v^r - p - \Delta) - \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^{r+1})\lambda t_1^{r+1}(v^{r+1} - p - \Delta) \\
& -\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^r)(v^{r+1} - v^r)(1 - e^{-\lambda t_1^r}) + \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^{r+1})\lambda t_2^{r+1}(v^{r+1} - p - \Delta) + \\
& +\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^r)\lambda t_2^{r+1}(v^r - p - \Delta) \\
& - \sum_{m<r}(1 - e^{-\lambda t_1^m})(v^{r+1} - v^r)\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^m)\Bigg)\frac{1}{\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 < v^r)} + \\
& -e^{-\lambda t_1^r}(v^r - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} - (v^r - p - \Delta)(1 - e^{-\lambda t_1^r})\frac{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^{r+1})}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} \\
& +e^{-\lambda t_1^{r+1}}(v^{r+1} - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} + (v^{r+1} - p - \Delta)(1 - e^{-\lambda t_1^{r+1}})\frac{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^{r+1})}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& + \sum_{m \leq r}(1 - e^{-\lambda t_2^m})(v^{r+1} - v^r)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^m)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& -\lambda t_1^{r+1}(v^r - p - \Delta)\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^r)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& +\lambda t_1^r \mathbb{P}_{\mu_1(t_1^r)}(v_2 = v^r)(v^r - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} \\
\geq\ & \star\star + \sum_{m \leq r}(e^{-\lambda t_1^m} - e^{-\lambda t_2^m})(v^{r+1} - v^r)\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^m)\frac{1}{\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 < v^r)} \\
& +(v^{r+1} - v^r)\frac{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^{r+1})}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} + \Big(e^{-\lambda t_1^r}(v^r - p - \Delta) - e^{-\lambda t_1^{r+1}}(v^{r+1} - p - \Delta)\Big) \\
& \times \frac{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^{r+1})}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} \\
& +e^{-\lambda t_1^{r+1}}(v^{r+1} - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 < v^r)} - e^{-\lambda t_1^r}(v^r - p - \Delta)\frac{1}{\mathbb{P}_{\mu_1(t_1^r)}(v_2 < v^r)} \\
\geq\ & \star\star
\end{aligned}
$$

The first inequality comes from (44). The second inequality is derived from the strategies from time t_1^r on since $\mathbb{P}_{\mu_2(t_2^r)}(v_1 = v^m) \geq \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^m)$ for $m < r$. The third inequality comes from $p + \Delta > 0$, from symmetry and equilibrium strategies since $t_2^{r+1} > t_1^r > t_2^r$, $\mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^r) \leq \mathbb{P}_{\mu_1(t_1^r)}(v_2 = v^r)$, $\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 = v^m) \geq \mathbb{P}_{\mu_2(t_2^{r+1})}(v_1 = v^m)$ for $m \leq r$ since $t_1^{r+1} \leq t_2^{r+1}$. The third and last inequality comes from the fact that the first sum in the right hand side after the second inequality is below the term $\mathbb{P}_{\mu_1(t_1^{r+1})}(v_2 \leq v^r)(v^{r+1} - v^r)$. The sum of the four remaining terms is positive because $e^{-\lambda t^r} v^r \leq e^{-\lambda t^{r+1}} v^{r+1}$. This can be seen from equations (37) and (38). The multiplicative terms involving time in the second term of equation (37) increases in t more slowly than multiplicative terms in the second term of (38). This implies that the difference between the second terms of equations (37) and (38) is more negative for $t = t^{r+1}$ and player 1 type $v^k = v^{r+1}$ than for time $t = t^r$ and player 1 type $v^k = v^r$. If it were the case that $e^{-\lambda t^r} v^r > e^{-\lambda t^{r+1}} v^{r+1}$ the difference between the first terms of equations (37) and (38) would be smaller at time $t = t^{r+1}$ and player 1's type $v^k = v^{r+1}$ than for time $t = t^r$ and player 1 type $v^k = v^r$. This contradicts the fact that at time $t = t^{r+1}$ and player 1's type $v^k = v^{r+1}$ (37) and (38) must be equal.

This concludes the proof that the timings of bidding are increasing on each player's own valuation type. From equations (39) and (41) we can conclude that the bidding timings decrease in the reserve price p.

Uniqueness of equilibrium play:

Furthermore, the equilibrium is unique. In fact, by arguments similar to the non-reputation case player 1 must bid the valuation before after a time threshold sufficiently late in the auction. Let $\tilde{\sigma}$ denote the strategies described above. Suppose there is an equilibrium with strategies different than $\tilde{\sigma}$. Let $\tilde{t}_1^k$ be the supremum time after which player 1 of type v^k bids in a way different than $\tilde{\sigma}$ and let $\tilde{t}$ be the minimum of these suprema. A little earlier than $\tilde{t}$ all types of player 2 would be playing according to $\tilde{\sigma}$ and the type of player 1 who was not playing according to $\tilde{\sigma}$ would want to deviate. If player 1 were to deviate during imitation phase revelation would ensue and play would revert to the unique THPBE with rational players. If player 2 is of the lowest type the best player 1 could do in response, would be to continue to imitate the commitment player until the time in which it is optimal to bid, in which case the best response of player 2 is the same. If player 1 is not supposed to bid yet and sees a deviation, if there is some type of player 2 that is already bidding player 1 would bid the valuation right away because he would assume player 2 is a high valuation type who bid the valuation. If it is earlier than the earliest time in which any player type of player 2 starts bidding, player 1 could bid the same way as if player 2 had placed the expected bid, the probability of a commitment player can deter player 2 from bidding higher.

A.3 Empirical Appendix

Below we provide summary statistics for our dataset. □

Table 5: Timing of winning bid

(T-Time winning bid)/T				
	Percentiles	Smallest		
1%	1.65E-06	1.16E-06		
5%	3.86E-06	1.16E-06		
10%	7.72E-06	1.16e-06	Obs	1250
25%	0.000027	1.65e-06	Sum of Wgt.	1250
50%	0.0002803		Mean	0.0332026
		Largest	Std. Dev.	0.1271027
75%	0.0053704	0.9693171		
90%	0.0439043	0.9805573	Variance	0.0161551
95%	0.1742774	0.9883102	Skewness	5.452606
99%	0.8354375	0.9913773	Kurtosis	34.65962

Table 6: Final price (US$) including shipping costs

Price plus shipping cost				
	Percentiles	Smallest		
1%	33	20.01		
5%	36	26.5		
10%	38	29	Obs	1242
25%	40	29.05	Sum of Wgt.	1242
50%	43		Mean	43.0027
		Largest	Std. Dev.	4.482932
75%	45.5	59		
90%	48.98	59.09	Variance	20.09668
95%	50	60	Skewness	0.2624133
99%	55.5	62	Kurtosis	4.378569

Table 7: Number of bidders per auction

Number of bidders				
	Percentiles	Smallest		
1%	0	0		
5%	0	0		
10%	0	0	Obs	2008
25%	0	0	Sum of Wgt.	2008
50%	3		Mean	4.178287
		Largest	Std. Dev.	4.198143
75%	8	16		
90%	10	16	Variance	17.6244
95%	12	16	Skewness	0.6821443
99%	14	19	Kurtosis	2.221005

Table 8: Number of bids per auction

Number of bids				
	Percentiles	Smallest		
1%	0	0		
5%	0	0		
10%	0	0	Obs	2008
25%	0	0	Sum of Wgt.	2008
50%	4		Mean	7.791335
		Largest	Std. Dev.	8.968126
75%	14	42		
90%	21	42	Variance	80.42729
95%	25	43	Skewness	1.08752
99%	33	44	Kurtosis	3.406615

Table 9: Linear regressions: number of bids as a function of bids in previous time intervals. Sample: all time windows in the final 2 hours of auctions.

	(1)	(2)
b^i_{j-1}	0.191	0.179
	(14.29)**	(13.96)**
b^i_{j-2}	0.095	0.083
	(9.23)**	(8.05)**
time remaining	-0.003	-0.003
	(11.62)**	(11.61)**
$time^2$	0.000	0.000
	(10.17)**	(10.18)**
$time^3$	-0.000	-0.000
	(9.55)**	(9.56)**
final period	0.607	0.609
	(22.85)**	(22.89)**
penultimate period	0.140	0.141
	(8.92)**	(8.99)**
current price	0.000	0.000
	(1.04)	(0.96)
seller score	-0.000	
	(2.71)**	
shipping price	-0.000	
	(1.44)	
num. bidders	0.002	
	(10.07)**	
Constant	0.061	0.073
	(11.99)**	(12.71)**
F statistic	110.6	145.6
R-squared	0.14	0.14
N	127,558	127,558

+ $p < 0.1$; * $p < 0.05$; ** $p < 0.01$

Linear regressions: dependent variable is number of bids in each window. Window = 15-second period, sample = final 2 hours of each auction. Errors clustered at auction level. Specification 2 includes auction fixed effects.

References

ABREU, D. AND F. GUL (2000): “Bargaining and Reputation,” *Econometrica*, 68, 85–117.

AMBRUS, A., J. BURNS, AND Y. ISHII (2014): “Gradual bidding in ebay-like auctions,” *Working Paper*, 1–50.

AMBRUS, A. AND S. E. LU (2014): “A continuous model of multilateral bargaining with random arrival times,” *American Economics Journal: Microeconomics*.

ARIELY, D., A. OCKENFELS, AND A. E. ROTH (2005): “An Experimental Analysis of Ending Rules in Internet Auctions,” *Rand Journal of Economics*, 36, 891–908.

BACKUS, M., T. BLAKE, D. V. MASTEROV, AND S. TADELIS (2015): “Is Sniping A Problem For Online Auction Markets?” in *Proceedings of the 24th International Conference on World Wide Web*, International World Wide Web Conferences Steering Committee, 88–96.

BAJARI, P. AND A. HORTACSU (2003): “The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay,” *RAND Journal of Economics*, 34, 329–355.

CALCAGNO, R., Y. KAMADA, S. LOVO, AND T. SUGAYA (2013): “Asynchronous Revision Games with Deadline : Unique Equilibrium in Coordination Games,” *Forthcoming Theoretical Economics*, 1–59.

ELY, J. AND T. HOSSAIN (2009): “Sniping and Squatting in Auction Markets,” *American Economic Journal:*, 68–94.

ENGELBERG, J. AND J. WILLIAMS (2009): “eBay’s proxy bidding: A license to shill,” *Journal of Economic Behavior & Organization*, 72, 509–526.

FUDENBERG, D. AND D. K. LEVINE (1989): “Reputation and Equilibrium Selection in Games with a Patient Player,” *Econometrica*, 759–778.

——— (1992): “Maintaining a Reputation when Strategies Are Imperfectly Observed,” *Review of Economic Studies*, 561–579.

FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, MIT Press.

HENDON, E., H. J. R. JACOBSEN, AND B. SLOTH (1996): “The One-Shot-Deviation Principle for Sequential Rationality,” *Games and Economic Behavior*, 12, 274–282.

HOPENHAYN, H. AND M. SAEEDI (2015): "Dynamic Bidding in Second Price Auctions," Working Paper.

KAMADA, Y. AND M. KANDORI (2011): "Revision games," *Working Paper*, 1–23.

KAMADA, Y. AND N. MUTO (2014): "Multi-Agent Search with Deadline," *Working Paper*, 1–77.

KREPS, D. AND R. WILSON (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 253–279.

LOVO, S. AND T. TOMALA (2015): "Markov-Perfect Equilibria in Stochastic Revision Games," *Working Paper*.

MAILATH, G. J. AND L. SAMUELSON (2006): *Repeated games and reputations*, vol. 2, Oxford university press Oxford.

MORONI, S. (2014): "Existence of trembling hand equilibrium in revision games with imperfect information." *Working Paper*.

MYERSON, R. B. (1997): *Game Theory: Analysis of Conflict*, Harvard University Press.

OCKENFELS, A. AND A. E. ROTH (2006): "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction," *Games and Economic Behavior*, 55, 297–320.

ROTH, A. E. AND A. OCKENFELS (2002): "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92, 1093–1103.

SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International journal of game theory*.