Gradual bidding in eBay-like auctions*

Attila Ambrus[†] Yuhta Ishii [‡] James Burns [§]

August 1, 2018

Abstract

This paper shows that in online auctions like eBay, even despite the option to leave proxy bids, if bidders are not continuously participating in the auction but can only place bids at random times, then many different equilibria arise besides truthful bidding. These equilibria can involve bidding behavior commonly observed on eBay such as gradual bidding, periods of inactivity, and waiting to start bidding until the end of the auction. As a result, the revenue of the seller in the latter can be a small fraction of what could be obtained at a sealed-bid second-price auction.

1 Introduction

A distinguishing feature of online auctions, relative to spot auctions, is that they typically last a relatively long time.[1] However, this aspect is often suppressed in the related economics literature. In particular, for bidders with private valuations, online auction mechanisms such as eBay, where bidders can leave a proxy bid and the highest bidder wins the object at a price equal to the second highest bid (plus a minimum bid increment), are commonly argued as strategically equivalent to second-price sealed-bid auctions. Since bidding one's true valuation in the latter context is a weakly dominant strategy, and placing a bid requires some costly effort, such an argument concludes that a rational bidder at an eBay-like auction should only place one bid equal to her true valuation, at her earliest convenience.

In contrast with the above predictions, observed bidding behavior on eBay involves both a substantial amount of gradual bidding and last-minute bidding (commonly referred to as

*We thank Itay Fainmesser, Drew Fudenberg, Tanjim Hossain, Andrew Postlewaite, Alvin Roth, Larry Samuelson, and anonymous referees as well as seminar participants at Duke University, the Harvard-MIT theory seminar, University of Maryland, Washington University and University of Pittsburgh for useful comments.

†Department of Economics, Duke University, Box 90097, Durham NC 27708, aa231@duke.edu

‡Centro de Investigación Económica, ITAM, yuhta.ishii@itam.mx

§Department of Economics, Harvard University, Cambridge, MA 02138, jburns@fas.harvard.edu

1On eBay, sellers can specify durations ranging from 1 to 30 days.

sniping). Ockenfels and Roth (2006) report that the average number of bids per bidder is 1.89 and 38% of bidders submit more than one bid. In the field experiment of Hossain and Morgan (2006), at least one bidder placed multiple bids in 76% of the observed auctions.[2] Regarding sniping, Roth and Ockenfels (2002) report that 18% of auctions in their data received bids in the last minute. Relatedly Bajari and Hortaçsu (2004) find that the median winning bid arrives after 98.3% of the auction time has elapsed, while 25% of winning bids arrive after 99.8% of the auction time has elapsed.

While Roth and Ockenfels (2002) propose a model in which last-minute bidding can be an equilibrium, the existing literature typically considers gradual bidding to be naive or irrational.[3] For example, Ku et al. (2005) explain bidding behavior on online auctions with a model of emotional decision-making and competitive arousal. Ely and Hossain (2009) models gradual bidders as confused, mistaking eBay's proxy system for an ascending auction. Similarly Hossain (2008) explains observed bidding behavior using behavioral buyers who learn about their own valuations through the process of placing bids.[4]

In this paper we show that if bidders are not present for the whole duration of the auction (a clearly unrealistic scenario for online auctions), and instead have periodic random opportunities to check the status of the auction to place bids,[5] then despite the possibility of proxy bids, there can be many different equilibria of the resulting game with perfectly rational bidders, in a private value context. The best equilibrium for the seller in this game still implies truthful bidding by the bidders at the first time they can place a bid. If the time horizon of the auction is long, the seller's revenue in this equilibrium is approximately what he could get in a second-price sealed bid auction. However, there are typically many other equilibria, in weakly conditionally undominated strategies, which involve gradual bidding,

[2]See also Zeithammer and Adams (2010), who find evidence in online auction field data for gradual bidding. For example they find that the frequency of the winning bid being exactly the minimum increment higher than the runner up bid to be too high than what would be implied by truthful bidding (they observe the winning bidder's proxy bid). Moreover, such an event is much more likely if the winning bidder places a bid after the runner up bidder than vice versa.

[3]However, Gonzalez et al. (2009), using ebay data, reject the hypothesis that bidders follow a war of snipe profile as in Roth and Ockenfels (2002). See also Ariely et al. (2005) for a related laboratory experiment.

[4]See also Compte and Jehiel (2004) and Rasmusen (2006) for models in which bidders learn their true valuations over the course of the auction. Other explanations include the presence of multiple overlapping auctions for identical or close substitute objects as in Peters and Severinov (2006), Hendricks et al. (2012), and Fu (2009). However, gradual and last-minute bidding seems to be prevalent also for rare or unique objects, and not only for objects with many close substitutes being auctioned at any time (for example, they occur in the experiments of Ariely et al. (2005) despite the fact that there is no concurrent competing auction). Furthermore, as Hossain (2008) points out, this type of argument also has trouble explaining many bids by the same bidder in a short interval of time, which is quite common in ebay. Bajari and Hortaçsu (2004) raise the possibility that all eBay auctions have some common value component.

[5]It is important for our results that whenever a bidder gets a chance to check the status of the auction, she can place multiple bids. In particular, if she places a bid incrementing the current price but gets notified that this bid was not enough to take over the lead, she can place a subsequent higher bid.

long periods of intentionally not placing bids, and sniping. The expected revenue of the seller from these latter equilibria can be a very small fraction of the expected revenue from the best equilibrium, even when the time horizon for the auction is arbitrarily large so that bidders (with large probability) get many opportunities to place bids.

To understand the intuition for the existence of such equilibria, consider two bidders, each with valuation $v > 2$ for the object being sold who obtain random opportunities to place bids (including potentially proxy bids exceeding what becomes the prevailing price) according to independent and identically distributed Poisson arrival processes. Suppose that the initial price is 0. Clearly, there is an equilibrium in which whenever the current price is below v and a bidder has the opportunity to make a bid, he places a bid of v. However, when the time horizon is sufficiently short, there is another equilibrium in which each bidder increases the price gradually by the minimum bid increment at each bidding opportunity. The key insight here is that gradual bidding is self-enforcing: if other bidders follow such a strategy then it is in the strict best interest of a bidder to do likewise. Increasing the price by more than the minimum increment does not increase the likelihood of winning the object, but rather accelerates the increase in the price, reducing the surplus accruing to the winning bidder. Hence, gradual bidding is a form of implicit collusion among bidders that is consistent with equilibrium in our model.

If the time horizon of the auction is long (relative to the arrival rates of bidding opportunities) then there are equilibria in which bidders both bid gradually at certain histories and also intentionally pass up bidding opportunities at others. More specifically, in Section 4, we construct equilibria characterized by a sequence of threshold times, τ_p, for each price p, which we call *completely gradual equilibria over delay sequences*, in which each bidder (on the equilibrium path) only begin to place bids at a prevailing price of p until times after the corresponding threshold time τ_p. This bidding behavior shares a similarity to sniping behavior studied in Roth and Ockenfels (2002), with a key difference that the snipers in our equilibrium only overbid incrementally rather than truthfully.[6] A noteworthy feature of these completely gradual equilibria over delay sequences is that it can prescribe placing a bid equal to the minimum increment upon the first arrival, followed by a long period of inactivity, which is in turn followed by all bidders bidding incrementally near the end of the auction. This is consistent with findings reported in Roth and Ockenfels (2002) and Bajari and Hortaçsu (2004) that the distribution of bid times is bimodal with a small mass of bidding near the the beginning of the auction, followed by a large mass of activity near

[6]In our framework, one can show that the strategy profiles studied in Roth and Ockenfels (2002) where bidders wait until the end of the auction to bid truthfully cannot be an equilibrium. Such strategy profiles in a continuous-time framework (with no special "last period" as in Roth and Ockenfels (2002)) unravel, as each bidder would prefer to snipe slightly earlier than the others.

the end.

The existence of completely gradual bidding equilibria over delay sequences holds substantial implications for the expected revenue of a seller. Note that in these equilibria, if the time horizon is sufficiently long then bidders are inactive for most of the duration of the auction, and only begin to incrementally bid until times close to the deadline. For this reason, the expected revenue of the seller in these equilibria can be a small fraction of v no matter how long the auction is, or how frequently players get opportunities to bid. This becomes especially problematic for the seller when the size of the minimum increment is small relative to the valuation of the bidders for the object. Formally, Theorem 4.1 demonstrates the existence of completely gradual equilibrium over delay sequences in which the fraction of the object's valuation accruing to the seller as revenue becomes arbitrarily small as the minimum increment becomes small relative to the valuation of the object, holding other parameters (including the number of bidders) fixed.

Conversely, when the number of bidders becomes arbitrarily large holding other parameters (including the minimum bid increment) fixed, we show that in every equilibrium in weakly conditionally undominated strategies, the seller's revenue approximates the bidders' valuation for the object. Therefore competition, in the form of a large enough number of bidders, ultimately erodes how much surplus bidders can extract from the seller in equilibria involving implicit collusion to keep the increase of winning price slow and gradual. As a result, these observations highlight the importance of the relative sizes between the minimum bid increment and competition in determining the expected revenue that the seller can expect to achieve in these auctions.

Gradual bidding equilibria over delay sequences also exist when arrival rates are time dependent, such as when they increase as the end of the auction approaches. Relative to constant arrival rates, players wait longer before placing bids, and incremental bidding might only occur during a short time interval before the end of the auction. For this reason, even when arrival rates reach very high levels before the end of the auction, as long as they are bounded, the expected number of bids, the winning price, and expected revenue can remain very low. In line with this point, Roth and Ockenfels (2002) observe that 90% of bidders reported that even after specifically planning to bid late, frequently some other issue or technical difficulty arose that prevented them from being available at the end of the auction to place a bid as planned.[7]

Finally, we note that our work is part of a recent string of papers examining continuous

[7]In the Supplementary Appendix, we also demonstrate that the existence of gradual bidding equilibria extend beyond the complete information symmetric bidders case, to specifications in which different bidders can have different valuations or are uncertain about other bidders' valuations.

time games with random discrete opportunities to take actions, in different contexts: Ambrus and Lu (2015) investigate multilateral bargaining with a deadline in a similar context, while Kamada and Kandori (2009) and Calcagno et al. (2014) examine situations in which players can publicly modify their action plans before playing a normal-form game. An important difference between these models and the one in this paper, is that in the former models, the actions players can take are unrestricted by previous history. In contrast, in an auction game like in the current paper previous bids restrict the set of feasible bids subsequently, as the leading price can only increase. There is also a recent string of papers in industrial organizations, on structural estimation of continuous-time models in which players can change their actions at discrete random times, but payoffs are accumulated continuously (e.g. Doraszelski and Judd (2012), Arcidiacono et al. (2016)).

2 The Model and Preliminaries

In this section we formally introduce a model of online auctions, conduct a formal investigation of weakly conditionally dominated strategies, and describe a benchmark equilibrium in which all players truthfully bid their valuations whenever they obtain a bidding opportunity.

2.1 The Model

A continuous-time single-good dynamic auction is defined by a set of n potential bidders, $N = \{1, \ldots, n\}$, with a common valuation $v \in \mathbb{R}_{++}$ and a common Poisson arrival rate $\lambda \in \mathbb{R}_{++}$, and a start time $T < 0$.[8] We normalize the end time of the auction to 0. This setting corresponds to an environment in which the good has a known common value.[9] While this is clearly a simplifying assumption, recent research on eBay suggests that a large number of auctions fall in this category. In particular, Einav et al. (2011) find numerous cases on eBay in which the same seller simultaneously sells identical items through both auctions and a traditional posted price mechanism. The latter can be considered as the market price, or commonly known value of the particular good.

Between times T and 0, bidders obtain random opportunities to place bids according to independent Poisson processes with arrival rates specified above. We normalize the starting bid to 0 and denote the minimum bid increment by $\Delta > 0$. For simplicity, we assume that v is an integer multiple of Δ so that $v/\Delta \in \mathbb{Z}_{++}$.[10] Bidders may make multiple bids during a

[8] $\mathbb{R}_{++}$ is the set of strictly positive real numbers.

[9] In the Supplementary Appendix, we explore extensions in which valuations and Poisson arrival rates may be heterogeneous across bidders.

[10] $\mathbb{Z}_{++}$ denotes the set of strictly positive integers.

single arrival and observe the outcome of each bid within this single arrival. For simplicity we assume that all bids made during an arrival are carried out instantaneously.

We assume that bidders can leave proxy bids, and hence we need to distinguish between the current price p and the current highest bid $B \geq p$. Furthermore let $C \in \{\emptyset, 1, \dots, n\}$ denote the identity of the *current winning bidder*, where $C = \emptyset$ denotes the scenario in which no one has placed a bid yet. We refer to all bidders who are not the current winning bidder as *losing bidders*.[11] The game begins with $B = 0$, $p = \emptyset$, and $C = \emptyset$ and B, p, and C adjust according to the rules to be described.

Upon the arrival of a bidding opportunity for bidder i, bidder i has at his disposal the following feasible actions:

$$A_i(B, p, C) := \begin{cases} \{b = k\Delta : k \in \mathbb{Z}_{++}, b \geq B + \Delta\} \cup \{a\} & \text{if } i = C, \\ \{b = k\Delta : k \in \mathbb{Z}_{++}, b \geq p + \Delta\} \cup \{a\} & \text{if } i \neq C, \end{cases}$$

where for expositional simplicity, we establish the convention throughout the paper that $\emptyset + \Delta = 0$. In words, i may either abstain (a) or place a strictly positive bid of at least $B + \Delta$ if $i = C$ or at least $\max\{\Delta, p + \Delta\}$ if $i \neq C$. If a bidder abstains, (B, p, C) remain unchanged and bidding ceases until the next random bidding opportunity arrives for some player. On the other hand, if a bid b is made by player i at time t, the *previous* triple of current highest bid, price, and current winning bidder $(\hat{B}, \hat{p}, \hat{C})$ updates to the *current* corresponding triple (B, p, C) as follows:

1. If $b \geq \hat{B} + \Delta$, then $B = b, p = \hat{B}, C = i$;[12]

2. If $b \leq \hat{B}$, then $B = \hat{B}$, $p = b$, $C = \hat{C}$.

Upon observing this outcome, i *may place further bids* or play the action a and abstain, in which case further bidding at time t stops and bidding ceases until the next random bidding opportunity arrives for some player. The auction ends at time $t = 0$, at which point the leading bidder at that time wins the good and pays the prevailing price. As in eBay auctions, we assume that the time paths of p and C are publicly observed, but that the current B is only known by the current winning bidder (provided that $C \neq \emptyset$). This defines the dynamic auction with common value v, n bidders, bid arrival rate λ, minimal increment Δ, and time horizon T, which we denote by $\Gamma(v, n, \lambda, \Delta, T)$.

[11] Thus the game begins with all bidders being losing bidders.

[12] In a previous version of the model, we defined the rules such that p becomes $B + 1$, as opposed to B. This corresponds better to the eBay mechanism, but makes the analysis more cumbersome. Nevertheless, the qualitative conclusions were the same in the two versions of the model.

The assumption that bidders can place multiple bids at a *single* bidding opportunity allows them to bid incrementally, regardless of the current highest bid. In particular, a losing bidder can always bid the current price plus Δ, until she becomes the winning bidder. This is an important component of our model.

2.2 Strategies and Weak Conditional Undomination

Here we formally define public and private histories, information sets, and strategies of players. We are particularly interested in strategies that are weakly undominated conditional on any information set of a player. It is well known that in a static second price auction with known valuations the only weakly undominated strategy of a player is bidding her true valuation, and the only Nash equilibrium of the game in weakly undominated strategies involves every player bidding her true valuation. Below we show that while imposing weak conditional undomination in our dynamic auction game does impose restrictions on strategies, e.g. it implies that no player ever places a bid higher than the true valuation v, it still leaves a lot of freedom for players to place bids below the true valuation, or to abstain from bidding even when he is a losing bidder. We do not provide a full characterization of weakly conditionally undominated strategies in the game. Rather we provide a set of sufficient conditions (in Claim 2.2) as well as a set of necessary conditions (Claims 2.3 and 2.4) for a strategy to be weakly conditionally undominated. The sufficient conditions that we provide are all satisfied by the strategies of all equilibria we construct in the paper, and hence all equilibria we present involve weakly conditionally undominated strategies. The necessary conditions for a strategy to be weakly conditionally undominated will be used to obtain bounds on the expected revenue of the seller for equilibria in weakly conditionally undominated strategies.

In our game, a *public* history at time t consists of the calendar time t itself and the path of winning prices and winner identities between T and t. Formally the set of all public time-t histories, H^t, with typical element h^t is described as follows: i) if no bidder has yet placed a bid, then $h^t = (\emptyset)$, and ii) if at least one bidder has placed a bid, then h^t is described by a finite sequence of prices, identities of winning bidders, and times: $h^t = (p_\ell, i_\ell, t_\ell)_{\ell=1}^k$, where $T \leq t_1 \leq \cdots \leq t_k \leq t$ and for all ℓ, $(p_\ell, i_\ell) \neq (p_{\ell+1}, i_{\ell+1})$. In words, the sequence $(p_\ell, i_\ell, t_\ell)_{\ell=1}^k$ describes all of the times (possibly including time t itself) and the values to which either the winning bidder or price changed. Note importantly that because a bidder may place multiple bids in a single bidding opportunity, we allow for the possibility that $t_\ell = t_{\ell+1}$ as long as either $p_\ell \neq p_{\ell+1}$ or $i_\ell \neq i_{\ell+1}$. Given a public history, h^t, it is useful to define the current price and current winning bidder:

$$(P(h^t), J(h^t)) = \begin{cases} (\emptyset, \emptyset) & \text{if } h^t = \emptyset, \\ (p_k, i_k) & \text{if } h^t = (p_\ell, i_\ell, t_\ell)_{\ell=0}^k . \end{cases}$$

Player i's *private* history at time t consists of the times in the interval $[T, t]$ (including t itself) at which player i obtained a bidding opportunity and the actions she took at these opportunities. This in particular implies that if player i is the winning bidder at time t, she knows the highest bid B at time t. Let us represent the set of all time t-private histories of player i as H_i^t with typical element h_i^t. The information set of player i at t denoted $I_i^t = (h^t, h_i^t)$ consists of both the public history and player i's private history at t.

Let $\mathcal{I}_i^t$ be the set of possible time t information sets of player i and let $\mathcal{I}_i = \bigcup_{t\in[T,0]} \mathcal{I}_i^t$. Let $\widehat{\mathcal{I}}_i^t$ be the elements $\mathcal{I}_i^t$ at which player i has a bidding opportunity at time t and define $\widehat{\mathcal{I}}_i = \bigcup_{t\in[T,0]} \widehat{\mathcal{I}}_i^t$. Note that for any time t, the full profile of information sets, $I^t = (I_1^t, I_2^t, \ldots, I_n^t)$, uniquely pins down the full history of play until time t in the game. Given a full profile of information sets, let $B(I_i^t, I_{-i}^t)$ denote the true highest bid at time t at the history corresponding to this profile of information sets. Importantly, note that for each player i and an information set $I_i \in \mathcal{I}_i$, not all profile of information sets of opponents may be consistent with I_i.[13] Thus, we define $\mathcal{I}_{-i}(I_i)$ to be the set of all possible profiles of information sets of opponents that are consistent with I_i. Also, for each $I_i \in \mathcal{I}_i$, we let $B(I_i) = \bigcup_{I_{-i}\in\mathcal{I}_{-i}(I_i)} B(I_i, I_{-i})$, which is the set of all possible current highest bids consistent with player i's information set I_i.[14] Finally as is standard in general extensive form games, we define $u_i(s_i, s_{-i} \mid I_i, I_{-i})$ to be the continuation payoff after the full history determined by (I_i, I_{-i}) conditional on all players playing according to the strategy profile (s_i, s_{-i}) subsequently.

A strategy for player i specifies a feasible action $s_i(I_i)$ for every $I_i \in \widehat{\mathcal{I}}_i$.[15] Let S_i be the set of strategies of player i, with generic element s_i. We say that a strategy for player i is a *public* strategy if s_i only depends on the public history: for any time t and $I_i = (h^t, h_i^t)$, $\hat{I}_i = (\hat{h}^t, \hat{h}_i^t) \in \widehat{\mathcal{I}}_i^t$ for which $h^t = \hat{h}^t$, $s_i(I_i) = s_i(\hat{I}_i)$. When s_i is a public strategy, we slightly abuse notation and denote by $s_i(h^t)$ the action prescribed at all information sets in $\widehat{\mathcal{I}}_i^t$ with public history h^t. Additionally we say that s_i is *Markovian* if s_i is a public strategy that only depends on directly payoff relevant information, namely the current price, the current time

[13]For example, if I_i consists of a public history in which no players have bid, then player $j \neq i$'s information set must also consist of the same public history.

[14]Note that if $I_i = (h^t, h_i^t)$ and $i = J(h^t)$, then $B(I_i)$ is a singleton corresponding to the current highest bid held by i.

[15]See the definition of $A_i(B, p, C)$ for the definition of feasible actions.

t, and whether or not bidder i is the current winning bidder.[16] In particular, a Markovian strategy is independent of the history of previous prices and winning bidders.

For our analysis, we are interested in strategies that are weakly conditionally undominated, as defined formally below.

Definition 2.1. Strategy s_i' weakly conditionally dominates s_i at $I_i \in \widehat{\mathcal{I}}_i$ if (i) $s_i'(I_i) \neq s_i(I_i)$; (ii) for all $s_{-i} \in S_{-i}$ and $I_{-i} \in \mathcal{I}_{-i}(I_i)$, $u_i(s_i', s_{-i}|I_i, I_{-i}) \geq u_i(s_i, s_{-i}|I_i, I_{-i})$; and (iii) for some $s_{-i} \in S_{-i}$ and $I_{-i} \in \mathcal{I}_{-i}(I_i)$, $u_i(s_i', s_{-i}|I_i, I_{-i}) > u_i(s_i, s_{-i}|I_i, I_{-i})$. If there exists some s_i' such that s_i' weakly conditionally dominates s_i at I_i, then we say that s_i is *weakly conditionally dominated* at I_i. We say that s_i is *weakly conditionally undominated* if s_i is not weakly conditionally dominated at any $I_i \in \widehat{\mathcal{I}}_i$.

To provide both sufficient and necessary conditions for weakly conditionally undominated strategies, let t^* be defined as the unique $t < 0$ satisfying $1 - e^{t\lambda} = e^{\lambda(n-1)t}$.[17] Intuitively, t^* is the time at which a player is indifferent between becoming the winning bidder at some price $p < v$, assuming that this event triggers every other bidder to place a bid of v upon first arrival, versus passing on the bidding opportunity and waiting for the next opportunity, assuming that if doing so, no other bidder will ever increase her bid. Given this interpretation for t^*, it is clear that at times after t^*, becoming the winning bidder, given $B < v$, is strictly preferred by i even with the most pessimistic belief regarding the continuation strategies of others conditional on this event, and the most optimistic belief regarding the continuation strategies conditional on abstaining at the current time.

Below we establish both sufficient and necessary conditions for a strategy to be weakly conditionally undominated. All proofs are relegated to the Appendix Section A. First we establish sufficient conditions. The following claim shows that if a strategy never calls for placing a bid higher than v, calls for losing bidders to place a bid at times near the end of the auction whenever it is possible to win the auction at a price strictly below v, and never calls on increasing the winning bid when the player is winning, then the strategy is weakly conditionally undominated. The last condition is not necessary, but all equilibria that we construct in the paper satisfy this requirement and imposing it simplifies the proof.

Claim 2.2. *Assume s_i satisfies the following properties: (i) at all $I_i \in \widehat{\mathcal{I}}_i$, s_i does not specify placing a bid $b > v$; (ii) if $I_i \in \widehat{\mathcal{I}}_i^t$, $t \geq t^*$, and $B(I_i)$ contains an element strictly lower than*

[16] Note that calendar time is indeed payoff relevant as it determines the distribution of future arrival sequences by the bidders (in particular, the probability that the given bidder will not get another chance to place a bid).

[17] Note that the left side is strictly decreasing and continuous in t, while the right side is strictly increasing and continuous in t. Moreover, for small enough t the left hand side is clearly larger than the right hand side, while for t close to 0 the right hand side is larger. Hence t^* exists and is unique.

v, then s_i specifies placing some bid; and (iii) if at $I_i \in \widehat{\mathcal{I}}_i$, bidder i is the winning bidder then s_i specifies abstaining. Then s_i is weakly conditionally undominated.

Next we show in the following two claims that essentially conditions (i) and (ii) (with a slight modification) of the above claim are indeed also necessary conditions for a strategy to be weakly conditionally undominated. These necessary conditions will be useful for bounding bidder payoffs and revenue to the seller across all equilibria in weakly conditionally undominated strategies.

Claim 2.3. *Suppose that s_i specifies placing a bid $b > v$ at some $I_i^* \in \widehat{\mathcal{I}}_i$ for which $B(I_i^*)$ contains an element weakly lower than v. Then s_i is weakly conditionally dominated at I_i^*. As a result, if all players play weakly conditionally undominated strategies, along the path of play, no player places a bid strictly higher than v.*

Claim 2.4. *Suppose that s_i prescribes abstaining at some $I_i^* \in \widehat{\mathcal{I}}_i^t$ for which (i) i is a losing bidder, (ii) $t \geq t^*$, and (iii) $B(I_i^*)$ contains an element strictly lower than v. Then s_i is weakly conditionally dominated at I_i^*.*

In the subsequent analysis we focus on perfect Bayesian equilibria in which all players use public strategies that are weakly conditionally undominated, which henceforth we refer to simply as equilibria. If additionally all strategies are Markovian, we refer to the equilibrium as a Markovian equilibrium.

2.3 Benchmark Truthful Equilibrium

One equilibrium that exists for any parametrization of the game is a strategy profile in which each bidder i bids v at any bidding opportunity if and only if i is a losing bidder and the current price is strictly below v and otherwise abstains. We refer to this equilibrium as the *truthful equilibrium* and summarize its properties in the following claim.

Claim 2.5. *There exists an equilibrium in which all losing bidders bid v at any bidding opportunity if and only if $p < v$ and otherwise abstains. In particular this equilibrium has the following features: i) it maximizes the seller's profits among all equilibria; ii) it minimizes the ex-ante payoffs of the bidders among all equilibria.*

The proof of this claim is straightforward and is relegated to Section A.4 in the Appendix. Due to the properties of the truthful equilibrium highlighted above and its close analogy to the unique equilibrium in weakly undominated strategies of a static second-price sealed bid auction, the truthful equilibrium provides a natural benchmark against which to compare other equilibria presented in the subsequent analysis.

3 Short Auctions

In this section we consider auctions that are short enough such that there exist equilibria in which any losing bidder wants to place a bid whenever he obtains a bidding opportunity. In particular, we analytically construct a Markovian equilibrium in which losing bidders upon arrival bid the minimal increment necessary to become the winning bidder. As we will see, these equilibria will serve as a natural starting point for our analysis in long auctions as well.

In Subsection 3.1 we provide an example of an equilibrium involving completely gradual bidding, and explain the main features of the dynamic strategic interaction in such an equilibrium. In Subsection 3.2, we formally construct a Markovian equilibrium with completely gradual bidding for general v, n, λ, Δ when $|T|$ is sufficiently small.

3.1 Example of a short auction with two bidders

Consider an auction with 2 symmetric bidders with values and arrival rates given by $v = 4$, $\lambda = 1$, $\Delta = 1$, and let $T = -1$. We now construct an equilibrium in which bidders make only the minimal bid necessary to hold the current high bid, whenever they arrive.

Formally, we consider a strategy profile in which a losing bidder bids $\max\{p+1, 1\}$ when $p \in \{\emptyset, 0, 1, 2, 3\}$ and abstains from bidding when $p \geq 4$.[18] At the same time, a winning bidder abstains from bidding even if she is given a bidding opportunity. Note that under this strategy profile, within a single bidding opportunity, the losing bidder places potentially multiple bids, until either she becomes the winning bidder or the price reaches 4.

To see that this simple strategy profile is indeed an equilibrium of the auction, first let us denote by $W(p,t)$ and $L(p,t)$ respectively the expected continuation payoffs of a winning bidder and a losing bidder conditional on a public history h^t on the *path of play of the above strategy profile* with $P(h^t) = p$ when all players play according to the strategy profile above. Note that we suppress the dependence of the continuation payoff on the current high bid as this is uniquely determined by the price along the path of play.[19]

Trivially, $W(4,t) = 0$ and $L(p,t) = 0$ for $p = 3, 4$. At prices $p = 2, 3$, the winning bidder obtains a payoff of $v - p$ if the other bidder obtains no further bidding opportunities before the end of the auction and 0 otherwise. Therefore $W(3,t) = e^t$ and $W(2,t) = 2e^t$.

The expected value of a losing bidder at $p = 2$, $L(2,t)$ is similarly derived by using the properties of a Poisson process. Note that if there are at least two arrivals by the losing

[18] Recall the convention that $\emptyset + \Delta = 0$.

[19] Note that at h^t on the path of play, the current highest bid is either (i) 0 if $p = \emptyset$, (ii) $p+1$ if $p = 0, 1, 2, 3$, or (iii) 4 if $p = 4$.

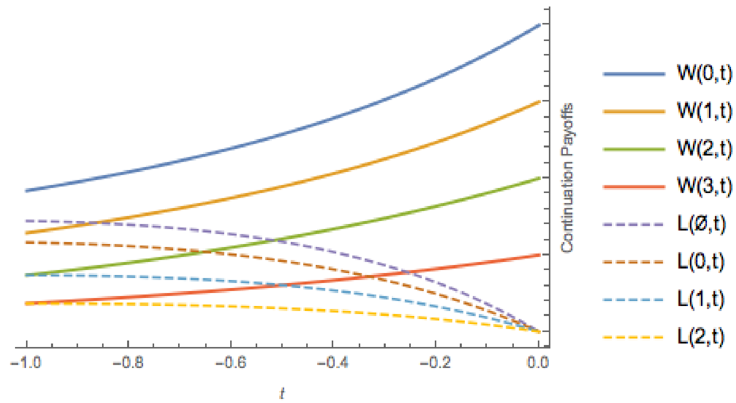

Figure 1: Continuation value functions on short time horizon

bidders, then this bidder obtains a payoff of zero.[20] Thus the only event under which this bidder obtains a positive payoff is if there is exactly one arrival by losing bidders. Given such an event, he obtains a payoff of $4 - 3 = 1$ since he must pay a price of 3. Under the Poisson arrival process specified in the model, this event occurs with probability equal to $|t|e^t$. Thus $L(2,t)$ is $L(2,t) = |t|e^t$. Similar arguments yield $L(1,t) = 2|t|e^t$.

Following along the same lines, the only events under which a winning bidder at a price of 1 obtains positive payoffs is if there are exactly zero or two arrivals by losing bidders. Under the first event, he obtains a payoff of 3 and while under the latter event, he obtains a payoff of 1. Using the basic properties of a Poisson process again, we obtain $W(1,t) = 3e^t + \frac{|t|^2}{2}e^t$. Using similar reasoning, being the winning bidder at $p = 0$ gives a continuation payoff of $W(0,t) = 4e^t + 2\frac{t^2}{2}e^t = 4e^t + t^2e^t$, whereas being the losing bidder at $p = 0$ yields $L(0,t) = 3|t|e^t + \frac{|t|^3}{3!}e^t$.

Before any bids have been placed, neither bidder holds the high bid and so the expected payoff is the expectation over becoming either the winning bidder at $p = 0$ or the losing bidder at $p = 0$ at the first bidding opportunity for the bidders with equal probability:

$$L(\emptyset,t) = \int_t^0 e^{-2(\tau - t)}(L(0,\tau) + W(0,\tau))d\tau.$$

Figure 1 depicts the expected continuation payoffs of winning and losing bidders at different prices, for the time horizon of the game. Note that indeed for all $t \geq -1$, $L(p,t) \leq W(p+1,t)$ for all $p = \emptyset, 1, 2, 3$, so that whenever $p = \emptyset, 1, 2, 3$, a losing bidder always prefers to become the winning bidder rather than abstaining and remaining the

[20]This occurs for example if the initial losing bidder obtains a bidding opportunity and then the subsequent losing bidder at a price of 3 again obtains a bidding opportunity before the end of the auction.

losing bidder.[21] Moreover, for all $t \geq -1$, being a winning bidder at a lower price is strictly better than at a higher price. As we will see in the proof of Theorem 3.2, this is sufficient to show that both (i) losing bidders have no incentive to place larger bids and (ii) winning bidders have no incentive to increase his current highest bid. As a result, the *completely gradual bidding strategy profile* described above is an equilibrium.

Intuitively, in the completely gradual bidding equilibrium illustrated above, a *losing bidder* faces a clear trade-off in her bidding decision. On the one hand, because the auction is relatively short, placing a bid and becoming the current winner increases her chance of winning the auction.[22] The downside is that placing a bid activates the other bidder and raises the expected price at which the good will sell. Placing a bid greater than the increment prescribed in equilibrium increases the downside without affecting the upside and hence if she chooses to place a bid it will also be incremental. Furthermore, the upside is decreasing in $|t|$ while the downside is increasing in $|t|$. If an auction is short enough, it will support a gradual bidding equilibrium in which bids are placed at every arrival by a losing bidder. This argument also hints that in longer auctions, equilibrium with gradual bidding also requires periods of waiting (losing bidders passing on opportunities to place a bid and take over the lead), as the incentive to slow down the increase of the current price may become stronger than the incentive to take over the lead. We discuss gradual bidding equilibria with delay in long auctions in Section 4.

We conclude this subsection by noting that in the above equilibrium, a bidder's expected payoff is $L(\emptyset, -1) \approx 1.45$, and the expected revenue for the seller is $(1 - e^{-2})4 - 2L(\emptyset, -1) \approx 0.57$. Because of the existence of gradual bidding behavior in equilibrium, these expected payoffs are considerably more favorable to the bidders than those in the truthful equilibrium, in which the expected payoffs are roughly 0.93 and 1.60 for the bidders and seller respectively.

3.2 Symmetric Markovian equilibria in short auctions

We now generalize the above construction for any v, λ, Δ, n. In particular, we show the existence of a completely gradual bidding equilibrium for any v, λ, Δ, n as long as the time horizon, T, is sufficiently short.

Definition 3.1. The *completely gradual bidding* strategy for bidder i, denoted s_i^g, is a public Markovian strategy in which at every public history h^t, each bidder i places a bid of

[21] Recall the convention that $\emptyset + \Delta = \emptyset + 1 = 0$.

[22] When the auction is sufficiently short, her chance of winning the auction at a final price $p < v$ increases upon becoming the winning bidder, because the probability that a losing bidder obtains a further bidding opportunity before the end of the auction is small. Indeed if the auction is longer, this probability may decrease upon becoming the winning bidder.

$\max\{\Delta, P(h^t) + \Delta\}$ if and only if $i \neq J(h^t)$ and $P(h^t) \leq v - \Delta$ or $P(h^t) = \emptyset$.[23] At all other histories, bidder i plays a. The *completely gradual bidding strategy profile*, s^g, is the public Markovian strategy profile in which each bidder i plays s_i^g.

Note that under the completely gradual bidding strategy profile, winning bidders never place a bid. Additionally, a losing bidder may potentially place multiple bids upon arrival. To see this, suppose that the current price is $P = p$ and the current highest bid is $B = p+\Delta < v$. Then a losing bidder upon arrival will first place a bid of $p+\Delta$. However, this only raises the current price to $P = p+\Delta$ while keeping unchanged the identity of the highest bidder. Thus in the same arrival opportunity, the losing bidder will submit an additional bid of $P = p+2\Delta$ after which the bidder becomes the new leading bidder and the price remains unchanged at $P = p + \Delta$. Because the bidder is now the winning bidder, he will cease further bidding at the current arrival.

We now extend the observations demonstrated in the example in Subsection 3.1 to show generally that the completely gradual bidding strategy profile defined above is indeed an equilibrium whenever the auction is sufficiently short.

Theorem 3.2. *Fix v, n, λ, Δ. There exists $T^* < 0$ such that the completely gradual bidding strategy profile, s^g, is a Markovian equilibrium in weakly conditionally undominated strategies in the auction $\Gamma(v, n, \lambda, \Delta, T)$ if and only if $T \geq T^*$. In particular, when $n = 2$, $T^* = -1/\lambda$.*

To prove the theorem, it is useful to first define the following value functions $W_\Delta(p, t)$ for all $p = 0, \Delta, \dots, v$, and $L_\Delta(p, t)$ for $p = \emptyset, 0, \dots, v$.[24] $W_\Delta(p, t)$ is the expected payoff to a winning bidder at time t conditional on a current price p and current highest bid of $\min\{p + \Delta, v\}$ when all players play according to s^g from time t until the end of the auction. We analogously define $L_\Delta(p, t)$ as the expected payoff to a losing bidder at time t conditional on a history at which the current price is p and current highest bid is believed to be $\min\{p + \Delta, v\}$ when all players again play according to s^g from time t until the end of the auction. These continuation values, $W_\Delta(p, t)$ and $L_\Delta(p, t)$, will be those associated respectively with a winning and losing bidder at a history with current price p at time t on the path of play of s^g. In Section B of the Appendix, we explicitly compute the expressions for $W_\Delta(p, t)$ and $L_\Delta(p, t)$ and illustrate some of their important properties, which we evoke in the proof below.

Proof of Theorem 3.2: Define $T^* := \inf\{t : W_\Delta(p + \Delta, t) \geq L_\Delta(p, t) \ \ \forall p = \emptyset, 0, \dots, v - \Delta\}$. Because $W_\Delta(v, t) = L_\Delta(v - \Delta, t) = 0$ for all t, for all $p = \emptyset, 0, \dots, v - 2\Delta$, $W_\Delta(p + \Delta, 0) =$

[23] Again recall the convention that $\emptyset + \Delta = 0$.
[24] Even though $\Delta > 0$ is fixed in the statement of the theorem, we subscript the value functions by $\Delta > 0$ to highlight its dependence on Δ as this will be useful later in Section 4.

$v - p - \Delta > 0 = L_\Delta(p, 0)$, and $W_\Delta(p + \Delta, t)$ and $L_\Delta(p, t)$ are continuous in t (see Appendix Corollary B.1), we have $T^* < 0$. Now note that if $T < T^*$, there exists some $t \in [T, T^*)$ and some $p = \emptyset, 0, \ldots, v - \Delta$ for which $W_\Delta(p + \Delta, t) < L_\Delta(p, t)$. Thus, clearly s^g is not an equilibrium for such T and proves the "only if" direction of the claim.

Conversely, suppose that $T \geq T^*$. We want to show that s^g is an equilibrium. First note that conditions (i), (ii), and (iii) of Claim 2.2 are trivially satisfied, which then guarantees that these strategy profiles are indeed weakly conditionally undominated.

We now prove that indeed s_i^g satisfies incentive compatibility at all public histories h^t. We first analyze the incentives of a winning bidder. Suppose that $i = I(h^t)$, $P(h^t) = p$, and the current highest bid, which i knows, is $b \geq p$. Suppose first that $b \geq v$. Then bidder i obtains a payoff of $(v - p)e^{\lambda(n-1)t}$ regardless of what strategy i plays. Thus it is incentive compatible for bidder i to play a. Suppose next that $p \leq b < v$. Then by following s_i^g, bidder i obtains $(v - p)e^{\lambda(n-1)t} + \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau - t)} L_\Delta(b, \tau) d\tau$. On the other hand, by placing any bid $b' > b$ followed by a, bidder i obtains $(v - p)e^{\lambda(n-1)t} + \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau - t)} L_\Delta(b', \tau) d\tau$. Note that the latter is weakly less than the former due to Lemma B.2 in the Appendix. Thus, it is optimal to play a.

Now we consider the incentives of the losing bidder and suppose that $i \neq I(h^t)$, $p = P(h^t)$. At such histories, bidder i does not observe the highest bid and so we define μ to be the belief that player i holds about the current high bid at history h^t. Note that at public histories h^t *on the path of play*, any such bidder believes that the highest bid is either v if $P(h^t) = v$ or $P(h^t) + \Delta$ if $P(h^t) = \emptyset, 0, \ldots, v - \Delta$. However, we must show incentive compatibility at *all* public histories h^t including those *off the path of play*. At public histories h^t off the path of play, μ may be different to those beliefs held by losing bidders on the path of play. In order to show incentive compatibility for *all* public histories h^t, we show that incentive compatibility holds for any possible belief μ that players can hold about the current highest bid at such a history.

To simplify notation let us define the following set for any $p = 0, \Delta, \ldots, v - \Delta$: $\mathcal{B}(p) = \{k\Delta : k \in \mathbb{Z}_{++}, p \leq k\Delta \leq v - \Delta\}$. First, if $p \geq v$, then incentives are trivial. Secondly, suppose that $0 \leq p < v$. Given any μ at h^t, by playing s_i^g, bidder i obtains an expected payoff of $\sum_{b \in \mathcal{B}(p)} \mu(b) W_\Delta(b, t)$.[25] On the other hand, a one-stage devation to playing a today

[25]This expected payoff is derived from the fact under s_i^g, the losing bidder i upon a bidding opportunity will place potentially multiple bids until either he becomes the winning bidder or the price reaches v.

yields bidder i a payoff of:

$$\sum_{b \in \mathcal{B}(p)} \mu(b) \int_{t}^{0} \lambda e^{-\lambda(n-1)(\tau-t)} \left(W_\Delta(b,\tau) + (n-2)L_\Delta(b,\tau)\right) d\tau$$
$$= \sum_{b \in \mathcal{B}(p)} \mu(b) L_\Delta(b-\Delta, t) \leq \sum_{b \in \mathcal{B}(p)} \mu(b) W_\Delta(b,t),$$

where the inequality follows from the assumption that $t \geq T^*$. Moreover consider a one-stage deviation of placing some bid $b' > p + \Delta$ followed by s_i^g afterwards. If $b' \geq v$, then this strategy yields a payoff of $\sum_{b \in \mathcal{B}(p)} \mu(b)(v-b)e^{\lambda(n-1)t} \leq \sum_{b \in \mathcal{B}(p)} \mu(b) W_\Delta(b,t)$, where the inequality follows since the winning bidder at a price of b always is guaranteed a payoff of at least $(v-b)e^{\lambda(n-1)t}$. If on the other hand, $b' \leq v - \Delta$, then this one-stage deviation yields a payoff of:

$$\sum_{\{b \in \mathcal{B}(p): b \leq b' - \Delta\}} \mu(b) \left((v-b)e^{\lambda(n-1)t} + \int_{t}^{0} \lambda(n-1) e^{-\lambda(n-1)(\tau-t)} L_\Delta(b', \tau) d\tau \right)$$
$$+ \sum_{\{b \in \mathcal{B}(p): b \geq b'\}} \mu(b) W_\Delta(b,t)$$
$$\leq \sum_{\{b \in \mathcal{B}(p): b \leq b' - \Delta\}} \mu(b) \left((v-b)e^{\lambda(n-1)t} + \int_{t}^{0} \lambda(n-1) e^{-\lambda(n-1)(\tau-t)} L_\Delta(b+\Delta, \tau) d\tau \right)$$
$$+ \sum_{\{b \in \mathcal{B}(p): b \geq b'\}} \mu(b) W_\Delta(b,t) = \sum_{b \in \mathcal{B}(p)} \mu(b) W_\Delta(b,t),$$

where the inequality follows from Lemma B.2 in the Appendix. Thus, following the strategy s_i^g is indeed optimal against all one-stage deviations.

Finally suppose that $p = \emptyset$. Here bidder i knows that indeed $B = 0$ since no bids have yet been placed. Therefore, playing s_i^g yields a payoff of $W_\Delta(0,t)$. A one-stage deviation to playing a today yields to bidder i a payoff of $L_\Delta(\emptyset, t) \leq W_\Delta(0,t)$ since $t \geq T \geq T^*$. Finally consider a one-stage deviation to bidding $b > \Delta$ and then following s_i^g. Such a strategy yields a payoff to bidder i of

$$ve^{\lambda(n-1)t} + \int_{t}^{0} \lambda(n-1) e^{-\lambda(n-1)(\tau-t)} L_\Delta(\min\{b, v\}, \tau) d\tau$$
$$\leq ve^{\lambda(n-1)t} + \int_{t}^{0} \lambda(n-1) e^{-\lambda(n-1)(\tau-t)} L_\Delta(\Delta, \tau) d\tau = W_\Delta(0,t),$$

where the inequality again follows from Lemma B.2. Thus following s_i^g is optimal. This then concludes the proof that s^g is an equilibrium if and only if $T \geq T^*$. The proof of the claim that $T^* = -1/\lambda$ for $n = 2$ is relegated to Section C in the Appendix. ∎

4 Long auctions

In this section we examine completely gradual bidding equilibria in auctions with longer time horizons. Supporting such equilibria requires periods for which losing bidders abstain from bidding. In particular, we consider equilibria in which players wait until times relatively close to the end of the auction before placing any bid. Further periods of inactivity, after a bid has already been placed, are also possible. Intuitively such delays in bidding are incentive compatible for bidders in these equilibria due to the (rational) fear that an early bid will trigger a dramatic increase in competition via a reversion to the truthful equilibrium. Notably because of the existence of such periods of inactivity at the beginning of the auction, regardless of the length of the auction, the expected revenue of the seller can potentially be very small relative to v.

In Subsection 4.1, we provide an example of a gradual bidding equilibrium with delays which generate small expected revenue to the seller even for very long auctions. In Subsection 4.2, we generalize this observation to show the existence of equilibria in which the expected revenue is a negligible fraction of the valuation v when the bid increment Δ is sufficiently small (relative to the value of the object) even when the time horizon is arbitrarily long. To show the existence of such equilibria, we construct a class of equilibria which we call *completely gradual bidding equilibria over delay sequences* that have this particular feature.

4.1 Example of a long auction with two bidders

The failure of the gradual bidding equilibria with no delays in long auctions can be seen by extending the length of the 2-bidder auction example from the previous section such that $T = -2$. Figure 2 plots the non-trivial bidder value functions in the completely gradual bidding strategy profile over the interval $[-2, 0]$. As we demonstrated in Subsection 3.1, for all $p = \emptyset, 1, 2, 3$ and $t > -1$, placing a bid is optimal as $W(p+1, t) \geq L(p, t)$. However, at any time $t < -1$, a winning bidder's expected value at $p = 3$ is lower than a losing bidder's expected value at $p = 2$ and hence a losing bidder facing a price of 2 would find it profitable in expectation to wait until $t > -1$ to place a bid.

Sustaining gradual bidding in equilibrium requires intervals during which losing bidders abstain from bidding even though the price is below their value. Losing bidders choose to

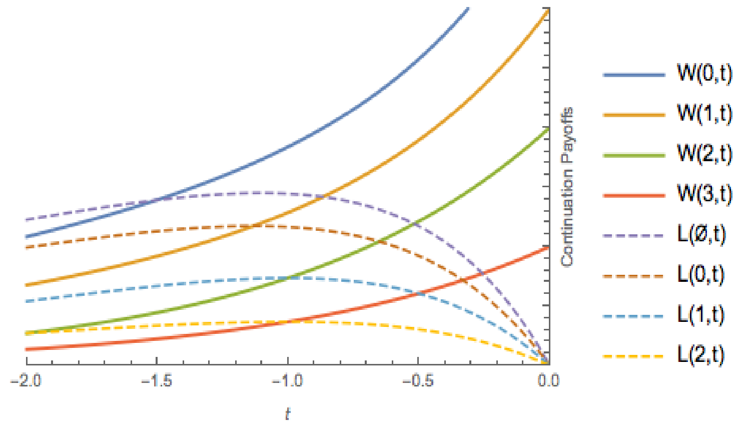

Figure 2: Continuation value functions with no delay

wait when the cost of increasing the price outweighs the likelihood of winning the object with the current bid. In our example, the trade-off is straightforward. Bidding at $p = 2$, since it induces the other player trying to bid again, yields a positive payoff only in the event that the other bidder does not return to the auction. This occurs with decreasing likelihood as we extend the time remaining in the auction. On the other hand, the likelihood of returning to the auction at the same price but closer to the end of the auction, and thereby face a more favorable trade-off, is increasing in the time remaining in the auction. For these reasons, at times far from the deadline, a losing bidder at $p = 2$ prefers to wait (if all others play the completely gradual bidding strategy with no delays).

We now construct an equilibrium that involves both gradual bidding and delays on the path of play for any time horizon T. Indeed we will construct a Markovian strategy in which for each price $p = \emptyset, 0, 1, 2, 3$, there will be a corresponding threshold time τ_p such that

1. winning bidders never place any bids;

2. at each price of p, the losing bidder places either a bid of v if $p = 3$ or the minimal bid necessary to become the new winning bidder if $p = \emptyset, 0, \ldots, 2$ if and only if $t \geq t_p$; otherwise a losing bidder abstains.

We will denote this strategy profile with cutoff times to be determined by s^*.

To construct these thresholds, recall the expressions $W(p, t)$ and $L(p, t)$ which respectively represent the values at the corresponding prices of a winning and losing bidder when all players play the completely gradual bidding strategy profile *without delays* from Subsection 3.1. First we let $\tau_3 = T$ so that at a price of $p = 3$ when the current high bid is 4, the losing bidder places a bid of 4 upon his next bidding opportunity to raise the price to

4.[26] Next consider the scenario in which the current price is 2 and the current high bid is 3. Under this scenario, let $\tau_2 = -1$, which is the unique time at which a losing bidder is indifferent between overtaking the current high bid to become the new winning bidder and waiting for the next bidding opportunity to place a bid, conditional on all players playing according to s^* in the future.

Next consider the incentives of a losing bidder at a current price of 1 and current high bid of 2. Note that if $\tau_1 > T$, then at time $t = \tau_1$, a losing bidder must be indifferent between placing a bid of 3 to become the new winning bidder and waiting for the next opportunity to place a bid. The former yields a payoff $W(2, \max\{t, \tau_2\})$ while the latter yields a payoff of

$$\int_t^0 e^{-(\tau-t)} W(2, \max\{\tau, \tau_2\}) d\tau = \begin{cases} L(2,t) & \text{if } t \geq \tau_2, \\ \left(1 - e^{-(t_0 - t)}\right) W(2, \tau_2) + e^{-(t_0-t)} L(1, \tau_2) & \text{if } t < \tau_2. \end{cases}$$

But clearly this is always weakly less than $W(2, \max\{t, \tau_2\})$ at all times t since $W(2, \tau_2) \geq L(2, \tau_2)$. Thus we proceed to set $\tau_1 = T$.

Proceeding in a similar manner, we can compute all the thresholds and find that $\tau_\emptyset = T, \tau_0 = -17/15, \tau_1 = T, \tau_2 = -1, \tau_3 = T$. It is straightforward to check that the strategy profile s^* with these threshold times is indeed an equilibrium for any arbitrary time horizon T. Figure 3 plots the relevant continuation value functions at each price of the equilibrium strategy profile s^*. Note the equilibrium path of play is divided into three phases of play. In the first phase, players initiate the bidding but further bidding does not take place keeping the price at 0 until τ_0 is reached. In the second phase, an arriving losing bidder bids the minimal increment necessary to take over the lead if $p \in \{0, 1\}$ but passes on opportunities to bid if $p = 2$. Finally, in the third phase, a losing bidder upon arrival bids the minimal incremental necessary to take over the lead until either the auction ends or the price reaches $p = 4$.

When $T = -2$, the equilibrium expected bidder payoff and seller revenue in this example are 1.44 and 1.05 respectively, while in the benchmark truthful equilibrium the bidder and seller respectively earn payoffs of 0.47 and 2.99. The seller revenue compares favorably to that of the short auction due to the prolonged length of the auction but is still significantly less than in the benchmark truthful equilibrium. Moreover, as $T \to -\infty$, the expected bidder payoff and seller revenue respectively converge to 1.44 and 1.13 in contrast to the corresponding truthful bidding equilibrium limit payoffs, 0 and 4. The reason is that an

[26] We could have also specified in equilibrium that the losing bidder abstains from bidding at such histories, which leads to different equilibrium payoffs. This construction however simplifies computation.

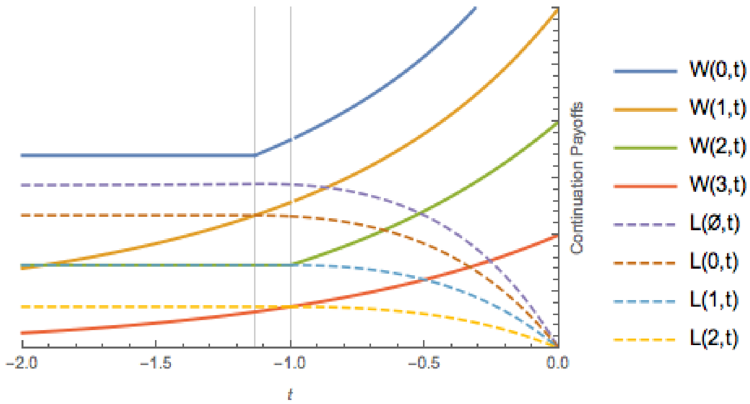

Figure 3: Continuation value functions with delay in equilibrium

increase in the length of the auction only serves to increase the length of time in the constructed equilibrium in which no bidding occurs after the initial bid (due to players waiting at all times $t < \tau_0$ at a price of 0).

4.2 Completely Gradual Bidding with Delays

As we have just seen in Subsection 4.1, because of the possibility of an initial period of inactivity in equilibrium, equilibria can potentially yield low revenue to the seller even when the time horizon is arbitrarily long. In fact such equilibria exist generally in any auction environment and is especially problematic for the revenue of the seller when Δ is small (relative to v). To state our main theorem of this section, let us introduce some additional notation. Given a strategy profile s of the auction $\Gamma(v, n, \lambda, \Delta, T)$, we denote by $R(s, \Gamma(v, n, \lambda, \Delta, T))$ the expected revenue that accrues to the seller when bidders play the strategy profile s in the auction $\Gamma(v, n, \lambda, \Delta, T)$.

Theorem 4.1. *Fix v, λ, n. For every $\varepsilon > 0$, there exists some Δ^* such that for all $\Delta < \Delta^*$ (such that $v/\Delta \in \mathbb{Z}_{++}$), and any T, there exists some equilibrium s of the auction $\Gamma(v, n, \lambda, \Delta, T)$ such that $R(s, \Gamma(v, n, \lambda, \Delta, T)) < \varepsilon$.*

To prove the above theorem, we construct a simple class of equilibria involving both gradual bidding and delays along the equilibrium path of play. In these equilibria, there will be a threshold time τ_p associated with each price $p = \emptyset, 0, \Delta, \ldots, v - \Delta$ such that all bidders delay bidding at a price of p until times after τ_p at which point losing bidders place the minimal bid necessary to become the winning bidder. As a preview to our proof, delay in bidding before τ_p in these equilibria is sustained by a threat of a reversion to the truthful equilibrium.

To define the strategy profiles formally, we partition the set of public histories into two categories. We start first with a delay sequence $\mathcal{T} = (\tau_\emptyset, \tau_0, \ldots, \tau_{v-\Delta})$, which, as we will see, describes the threshold times after which players begin to place bids at each corresponding price. Given $\mathcal{T}$, we define the set of public histories $\mathcal{H}^{t,\mathcal{T}}$ which will form the set of on-path public histories in the strategy profile to be constructed. Conversely, any history $h^t \notin \mathcal{H}^{t,\mathcal{T}}$ will be a public history off the path of play in the strategy profiles to be constructed.

Formally, given $\mathcal{T}$, a public history h^t is an element of $\mathcal{H}^{t,\mathcal{T}}$ if and only if either $h^t = (\emptyset)$ or $h^t = (p_\ell, i_\ell, t_\ell)_{\ell=0}^k$ satisfies the following conditions:

1. $t_0 < t_1 < \cdots < t_k$ and $\tau_\emptyset \leq t_0, \tau_0 \leq t_1, \ldots, \tau_{(k-1)\Delta} \leq t_k$;
2. $p_0 = 0, p_1 = \Delta, \ldots, p_k = k\Delta \leq v$;
3. $i_0 \neq i_1, i_1 \neq i_2, \ldots, i_{k-1} \neq i_k$.

To simplify notation, we partition $\mathcal{H}^{t,\mathcal{T}}$ into sets even further. For each bidder i, let

$$\mathcal{H}_i^{t,\mathcal{T},b} = \left\{h^t \in \mathcal{H}^{t,\mathcal{T}} : i \neq J(h^t), v > P(h^t), t \geq \tau_{P(h^t)}\right\}, \mathcal{H}_i^{t,\mathcal{T},a} = \mathcal{H}^{t,\mathcal{T}} \setminus \mathcal{H}_i^{t,\mathcal{T},b}.$$

Note that $\mathcal{H}_i^{t,\mathcal{T},b}$ consists of histories in $\mathcal{H}^{t,\mathcal{T}}$ where i is a losing bidder, the current price is below v, and time has surpassed the threshold associated with the current price. As a preview of our equilibrium constructions to follow, these histories will be precisely the histories in $\mathcal{H}^{t,\mathcal{T}}$ at which losing bidders place a bid. At all other histories in $\mathcal{H}^{t,\mathcal{T}}$, losing bidders will abstain.

Definition 4.2. Let $\mathcal{T} = (\tau_\emptyset, \tau_0, \ldots, \tau_{v-\Delta})$ and consider any auction environment $\Gamma(v, n, \lambda, \Delta, T)$. We define the player i strategy denoted $s_i^{\Delta,\mathcal{T}}$ in this auction as follows:

$$s_i^{\Delta,\mathcal{T}} = \begin{cases} v & \text{if } h^t \notin \mathcal{H}^{t,\mathcal{T}}, i \neq J(h^t), \text{ and } P(h^t) < v, \\ a & \text{if } h^t \in \mathcal{H}_i^{t,\mathcal{T},a}, i = J(h^t), \text{ or } P(h^t) \geq v, \\ \min\{v, P(h^t) + 2\Delta\} & \text{if } h^t \in \mathcal{H}_i^{t,\mathcal{T},b} \text{ and } P(h^t) \geq 0, \\ \Delta & \text{if } h^t \in \mathcal{H}_i^{t,\mathcal{T},b} \text{ and } P(h^t) = \emptyset. \end{cases}$$

The strategy profile, $s^{\Delta,\mathcal{T}} := \left(s_1^{\Delta,\mathcal{T}}, \ldots, s_n^{\Delta,\mathcal{T}}\right)$, in which each bidder i plays the strategy $s_i^{\Delta,\mathcal{T}}$ is called the *completely gradual bidding strategy profile with delay sequence* $\mathcal{T}$.

Remark. Note that if $\mathcal{T} = (T, T, \ldots, T)$, then $s^{\Delta,\mathcal{T}}$ is exactly the strategy profile s^g from Theorem 3.2.

The strategy profile $s^{\Delta,\mathcal{T}}$ is quite intuitive. On the path of play, bidders begin with a period of delay until time $\tau_{\emptyset}$ is reached after which the first arriving bidder places a bid of Δ to become the winning bidder. If $\tau_0 > \tau_{\emptyset}$, another period of inactivity ensues until time τ_0 arrives after which the next osing bidder to arrive places the minimal bid (2Δ) necessary to become the new winning bidder. If on the other hand, $\tau_0 < \tau_{\emptyset}$, the period of inactivity at a price of 0 is non-existent and the next arriving losing bidder becomes the new winning bidder with a high bid of 2Δ. Play then proceeds in a similar manner until either the price reaches v or the auction ends. We can now state the following proposition that shows the existence of a delay sequence $\mathcal{T}^{\Delta}$ for which $s^{\Delta,\mathcal{T}^{\Delta}}$ is an equilibrium.

Proposition 4.3. *Fix v, λ, n. Then for each $\Delta > 0$, there exists some $\mathcal{T}^{\Delta} = (\tau_{\emptyset}^{\Delta}, \tau_0^{\Delta}, \dots, \tau_{v-\Delta}^{\Delta})$ such that $s^{\Delta,\mathcal{T}^{\Delta}}$ is an equilibrium in the auction $\Gamma(v, n, \lambda, \Delta, T)$ for all $T < 0$.*

The complete proof of Proposition 4.3 is relegated to the Appendix. However, let us provide the construction of $\mathcal{T}^{\Delta}$ as well as an intuition for why the proposed equilibrium is indeed an equilibrium. To construct $\mathcal{T}^{\Delta}$, given any price p and time t, recall from Section 3 that $W_{\Delta}(p,t)$ and $L_{\Delta}(p,t)$ denote the respective continuation values to the winning and losing bidders when all bidders follow the strategy profile s^g (from Theorem 3.2) at all times after t in the dynamic auction $\Gamma(v, n, \lambda, \Delta, T)$. We first define τ_0^{Δ} as the unique time at which $(v-\Delta)e^{\lambda(n-1)\tau_0^{\Delta}} = L_{\Delta}(0, \tau_0^{\Delta})$.[27] We then define $\tau_{\Delta}^{\Delta} = \dots = \tau_{v-\Delta}^{\Delta} = \tau_0^{\Delta}$. To construct $\tau_{\emptyset}^{\Delta}$, we define the following function:

$$\tilde{L}_{\Delta}(\emptyset, t) := \int_t^0 \lambda e^{-\lambda n(\tau - t)} \left(W_{\Delta}(0, \tau \vee \tau_0^{\Delta}) + (n-1)L_{\Delta}(0, \tau \vee \tau_0^{\Delta})\right) d\tau$$

where $x \vee y = \max\{x, y\}$. Intuitively the above function represents the expected payoff that bidders would receive at time t and $P(h^t) = \emptyset$, when all players play according to the strategy $s^{\Delta,\hat{\mathcal{T}}}$ where $\hat{\mathcal{T}} = (t, \tau_0^{\Delta}, \tau_0^{\Delta}, \dots, \tau_0^{\Delta})$. Then we define $\tau_{\Delta}^{\emptyset}$ to be the unique time at which $ve^{\lambda(n-1)\tau_{\emptyset}^{\Delta}} = \tilde{L}_{\Delta}(0, \tau_{\emptyset}^{\Delta})$.[28] Thus the appropriate delay sequence is given by $\mathcal{T}^{\Delta} := (\tau_{\emptyset}^{\Delta}, \tau_0^{\Delta}, \dots, \tau_0^{\Delta})$.

Intuitively in equilibrium, bidding is deterred before the threshold $\tau_{\emptyset}^{\Delta}$ due to the fear that an early bid will induce a substantial increase in competition which in our equilibrium construction manifests through a reversion to the truthful equilibrium. More formally, to see why $s^{\Delta,\mathcal{T}^{\Delta}}$ is an equilibrium, consider the incentives of a bidder at a history h^t when $P(h^t) = \emptyset$ and $t < \tau_{\emptyset}^{\Delta}$.[29] At such a history, note that if a bidder chooses to place a bid, he

[27] Proof of existence and uniqueness of t_0^{Δ} is contained in the Appendix.

[28] Again the proof of existence and uniqueness of $\tau_{\emptyset}^{\Delta}$ is in the Appendix.

[29] Incentives at other histories follow similar arguments and the complete proof is provided in detail in

will face a continuation strategy of truthful bidding by the opponents giving him a payoff of $ve^{\lambda(n-1)t}$. On the other hand, following the strategy profile $s^{\Delta,\mathcal{T}^\Delta}$ gives him a payoff of $\tilde{L}_\Delta(\emptyset,t)$, which by definition of $\tau_\emptyset^\Delta$ is bigger than $ve^{\lambda(n-1)t}$ at times $t<\tau_\emptyset^\Delta$. Conversely, when $P(h^t)=\emptyset$ and $t\geq\tau_\emptyset^\Delta$, not even the threat of reversion to truthful bidding is sufficient to deter bidding.

Given the above proposition, it is straightforward to prove Theorem 4.1.

Proof of Theorem 4.1. To prove Theorem 4.1, consider the delay sequence $\mathcal{T}^\Delta$ constructed in the discussion succeeding Proposition 4.3. Recall the following equations that define $\tau_\emptyset^\Delta$ and τ_0^Δ: $L_\Delta(0,\tau_0^\Delta)=(v-\Delta)e^{\lambda(n-1)\tau_0^\Delta}$ and $\tilde{L}_\Delta(\emptyset,\tau_\emptyset^\Delta)=ve^{\lambda(n-1)\tau_\emptyset^\Delta}$.

For any t, by Corollary B.1 in the Appendix,

$$W_\Delta(0,t)=e^{\lambda(n-1)t}\sum_{k=0}^{v/\Delta}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n}+\frac{(-1)^k(n-1)}{n}\right)(v-k\Delta)$$

$$L_\Delta(0,t)=e^{\lambda(n-1)t}\sum_{k=0}^{v/\Delta}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k+(-1)^{k+1}}{n}\right)(v-k\Delta)$$

Note that both $W_\Delta(0,t)$ and $L_\Delta(0,t)$ are strictly decreasing in Δ and therefore, both τ_0^Δ and $\tau_\emptyset^\Delta$ are strictly decreasing in Δ. As a result, there exists $\bar{t}$, namely $\bar{t}=\min\{\tau_0^v,\tau_\emptyset^v\}$, such that for all $\Delta\in(0,v]$, $\bar{t}\leq\tau_0^\Delta,\tau_\emptyset^\Delta$.

Now we bound $R(s^{\Delta,\mathcal{T}^\Delta},\Gamma(v,n,\lambda,\Delta,T))$ as follows. Consider instead the strategy profile $s^{\Delta,\hat{\mathcal{T}}}$ where $\hat{\mathcal{T}}=(\bar{t},\bar{t},\ldots,\bar{t})$. This is the strategy profile (not necessarily an equilibrium) in which all bidders abstain from bidding until $\bar{t}$ after which bidders bid completely gradually. Since bidding takes place earlier under $s^{\Delta,\hat{\mathcal{T}}}$ relative to $s^{\Delta,\mathcal{T}^\Delta}$ (with respect to first order stochastic dominance), it is clear that $R(s^{\Delta,\mathcal{T}^\Delta},\Gamma(v,n,\lambda,\Delta,T))\leq R(s^{\Delta,\hat{\mathcal{T}}},\Gamma(v,n,\lambda,\Delta,T))$. But note that

$$\sup_{T<0}R(s^{\Delta,\hat{\mathcal{T}}},\Gamma(v,n,\lambda,\Delta,T))\leq R(s^{\Delta,\hat{\mathcal{T}}},\Gamma(v,n,\lambda,\Delta,\bar{t}))\leq e^{\lambda n\bar{t}}\sum_{k=0}^{v/\Delta}\frac{(\lambda n|\bar{t}|)^k}{k!}k\Delta\to 0,$$

as $\Delta\to 0$, where convergence follows from the Lebesgue dominated convergence theorem. This concludes the proof. ∎

The intuition for Theorem 4.1 is simple once we understand Proposition 4.3. Because of the possibility of arbitrarily long periods of delay in equilibrium (which becomes even more prolonged when Δ is small), bidding activity is concentrated on a short period of time near the end of the auction in these equilibria. But in such a short period of time, the expected

Section D of the Appendix.

number of bidding opportunities is bounded which leads to very low revenue for the seller when Δ is small.

5 A Comparative Static in Competition

We noted in Section 4 that as $\Delta \to 0$, the worst case revenue (among all equilibria) accruing to the seller converges to zero regardless of the time horizon. This is because of the existence of equilibria in which bidding activity is concentrated near the end of the auction at which point all bidders place the minimal bid necessary to become the new winning bidder.

In this section, we study whether the seller of such auctions can benefit from an increase in competition. The answer is yes as we will see in the following proposition. In particular, holding fixed $\Delta > 0$ (as well as v, λ, T), when the number of bidders is sufficiently large, the seller is able to extract almost all surplus (at least $v - \Delta$) from the bidders in *all equilibria in weakly conditionally undominated strategies.*

Proposition 5.1. *Fix v, Δ, λ, T. Then for every $\varepsilon > 0$ there exists some N^* such that for all $n \geq N^*$ and any equilibrium in weakly conditionally undominated strategies, s, of the dynamic auction $\Gamma(v, n, \Delta, \lambda, T)$, $R(s, \Gamma(v, n, \Delta, \lambda, T)) \geq v - \Delta - \varepsilon$.*

Proof. Let s be an equilibrium in weakly conditionally undominated strategies. Let t_n^* be the threshold such that $1 - e^{\lambda t_n^*} = e^{\lambda(n-1)t_n^*}$. First note that as $n \to \infty$, $t_n^* \to 0$. Thus, $\lim_{n\to\infty} e^{\lambda(n-1)t_n^*} = 0$ which means that $\lim_{n\to\infty}(n-1)t_n^* = -\infty$.

By Claim 2.4, s must prescribe bidding for losing bidders at all times $t \geq t_n^*$ when $P(h^t) = \emptyset, \Delta, \ldots, v - 2\Delta$. Since losing bidders arrive at least at Poisson rate $\lambda(n-1)$ at all times, the following is a lower bound on the revenue to the seller:[30]

$$R(s, \Gamma(v, n, \lambda, \Delta, T)) \geq \left(1 - e^{\lambda(n-1)t_n^*} \sum_{k=1}^{v/\Delta - 1} \frac{(\lambda(n-1)|t_n^*|)^k}{k!}\right)(v - \Delta).$$

The above is a lower bound on the revenue to the seller in any equilibrium because

$$\left(1 - e^{\lambda(n-1)t_n^*} \sum_{k=1}^{v/\Delta - 1} \frac{(\lambda(n-1)|t_n^*|)^k}{k!}\right)$$

is the probability that there are at least v/Δ arrivals of a Poisson process with arrival rate $\lambda(n-1)$ in the time interval $[t_n^*, 0]$. Under this event, the seller obtains a revenue of at least

[30] Note that losing bidders actually arrive at a faster rate of λn when $P(h^t) = \emptyset$.

$v-\Delta$. Note then that the lower bound above converges to $v-\Delta$ since $\lim_{n\to\infty} \lambda(n-1)|t_n^*| = -\infty$.

Then given any $\varepsilon > 0$, this implies that there exists some N^* such that for all $n \geq N^*$ and any s, an equilibrium in weakly conditionally undominated strategies, $R(s, \Gamma(v, n, \lambda, \Delta, T)) \geq v - \Delta - \varepsilon$. This concludes the proof. ∎

Remark. Note that the lower bound is $v - \Delta - \varepsilon$ rather than $v - \varepsilon$. The reason is that there may be equilibria in weakly conditionally undominated strategies in which no bidder places any further bids at a price of $v - \Delta$ on the equilibrium path. This is because at a public history $(p_\ell, i_\ell, t_\ell)_{\ell=1}^k$ where $p_k = v - \Delta$ and $i_k \neq i_{k-1}$, it is common knowledge that the highest bid is at least v, and as a result, strategies that prescribe bidders to play a at such histories need not be weakly conditionally dominated.[31]

The proposition above in conjunction with Theorem 4.1 highlight the importance of the relative sizes of $\Delta > 0$ and n for the revenue of a seller. On the one hand, when Δ is very small relative to n, Theorem 4.1 shows the existence of equilibria in which the revenue to the seller is driven down to almost zero. In contrast, the above proposition shows that when n is very large relative to Δ the seller is able to capture almost all surplus in *all equilibria* in weakly conditionally undominated strategies.

6 Extension: Time-dependent arrival rates

For expositional simplicity, we have assumed thus far that arrival rates are constant throughout the whole duration of the game. Here we discuss how our results extend to environments with time-dependent arrival rates. Other extensions of our model, such as asymmetric valuations and asymmetric information regarding valuations, are included in the Supplementary Appendix.

First, note that multiplying all arrival rates by a constant $\alpha > 0$ is equivalent to rescaling time by $\frac{1}{\alpha}$. In particular, if $\Gamma(v, n, \lambda, \Delta, T)$ has a completely gradual bidding equilibrium with delay sequence $\mathcal{T} := \{\tau_\emptyset, \tau_0, \ldots, \tau_{v-\Delta}\}$ then the dynamic auction $\Gamma(v, n, \alpha\lambda, \Delta, T)$ has a completely gradual bidding equilibrium equilibrium with delay sequence $\mathcal{T}_\alpha := \{\frac{1}{\alpha}\tau_\emptyset, ..., \frac{1}{\alpha}\tau_{v-\Delta}\}$. Furthermore, expected payoffs to the bidders and expected revenue in the completely gradual bidding equilibrium with delay sequence $\mathcal{T}$ in $\Gamma(v, n, \lambda, \Delta, T)$ are the same as those in the completely gradual bidding equilibrium with delay sequence $\mathcal{T}_\alpha$ in the dynamic auction with time horizon T/α, $\Gamma(v, n, \alpha\lambda, \Delta, T/\alpha)$. Moreover, if $T \leq \tau_\emptyset$ and $\alpha > 1$, then the expected

[31] Claim 2.4 does not state that such strategies at such histories are weakly conditionally dominated because at such histories $B(I_i)$ does not contain an element strictly smaller than v.

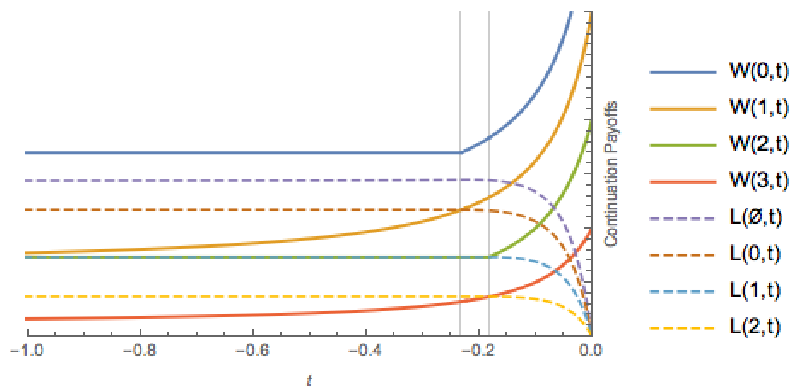

Figure 4: Equilibrium continuation values with time-dependent $\lambda(t)$

payoffs to the bidders and the expected revenue in $s^{\Delta,\mathcal{T}}$ in $\Gamma(v,n,\lambda,\Delta,T)$ is exactly the same as the corresponding equilibrium $s^{\Delta,\mathcal{T}_\alpha}$ in the auction, $\Gamma(v,n,\alpha\lambda,\Delta,T)$, with arrival rate $\alpha\lambda$ and time horizon T. Essentially an increase in the arrival rate leaves both the equilibrium expected payoffs and revenues unchanged across these two equilibria as it only shifts all thresholds to the right in an inversely proportional manner. Intuitively, if bidders obtain bidding opportunities more frequently, they become more willing to postpone bidding at each price level until times closer to the deadline to exactly offset the increase in the arrival rates of bidding opportunities.

We obtain qualitatively similar conclusions when allowing for time-varying arrival rates, as long as they stay bounded. For example, it is natural to assume that bidders are much more likely to check the status of the auction towards the end of the auction. We model this feature by assuming that the arrival rate of bidding opportunities is increasing over time.[32] As in the previous discussion, such time-dependence of bidding opportunities only serves as a rescaling of time, where incentives for bidders at time t in this model is exactly identical to the incentives for bidders at time $\hat{t} = -\int_t^0 \lambda(s)ds$ in a dynamic auction model with constant arrival rate 1 of bidding opportunities.

To illustrate this point more concretely, let us return to the symmetric 2-bidder auction example from Section 4.1 where $v = 4$ and $T = -2$ but now let the arrival rate be given by the strictly increasing function $\lambda(t) = \frac{a}{(1-bt)^2}$ where $a, b > 0$. The arrival rate at the end of the auction, $\lambda(0)$, is then determined by the parameter a and the parameter b determines how steeply arrival rates increase at the end of the auction. The average arrival rate over the auction is given by $\bar{\lambda} = \int_{-2}^0 \frac{a}{(1-bt)^2}dt$. We choose $a = 10$ and $b = 9/2$ which gives an average arrival rate of 1 as in our original example. Figure 4 shows the effect of an

[32] Presumably endogenizing arrival rates would lead to such time patterns of arrival rates, although we do not pursue this direction formally here, due to its analytical complexity.

increasing arrival rate on bidder value functions in a completely gradual equilibrium with delays. The qualitative structure of the equilibrium is identical to the equilibrium constructed in Subsection 4.1 as bidding begins initially after which bidders wait to bid at $p = 0$ until after a threshold time. However, because bidder's arrival rates are increasing over the course of the auction, the thresholds are much closer to the end of the auction, which reinforces the high frequency of late-bidding in equilibrium. As long as arrival rates are bounded, incremental equilibria are robust to increasing arrival rates. Note also that in line with our previous discussion expected payoffs and revenue in the equilibrium with time-dependent arrival rates are exactly the same as those in the equilibrium with a constant arrival rate.

7 Conclusion

This paper shows that in online auctions like eBay where bidders can leave proxy bids, if bidders get limited and random opportunities to place bids then many different equilibria arise in weakly undominated strategies. Bidders can implicitly collude by bidding gradually or by waiting to place bids, in a self-enforcing manner, slowing down the increase of leading price. Such equilibria can leave the seller with very low revenue especially when the minimal bid increment is small. These features of our model are consistent with the empirical observations that both gradual bidding and sniping are common bidder behaviors on ebay.

Our investigation suggests that given a fixed set of bidders, running an ascending auction with a long time horizon (long enough that bidders cannot continuously participate) has the potential to affect the revenue of the seller very adversely, even when proxy bidding is possible, relative to running a prompt auction. Hence, introducing a time element can only be beneficial if it takes time for potential bidders to discover the auction.[33] It is an open question what mechanism guarantees the highest possible revenue for the seller in such environments. In order to prevent implicit collusive equilibria, sellers may want to set high reservation prices and/or minimum bid increments. They may also want to allow each bidder to submit at most one bid over the course of an auction, although in practice this might be difficult to enforce, given that the same person can have multiple online identities. We leave the formal investigation of these issues to future research.

[33] Fuchs and Skrzypacz (2010) consider the arrival of new buyers over time, but in a dynamic bargaining context in which the seller cannot commit to a mechanism. Another difference compared to our setting is that in their model once a buyer arrives, she is continuously present until the end of negotiations.

A Proofs for Section 2

A.1 Proof of Claim 2.2

Proof. We will show that for any t and $I_i \in \widehat{\mathcal{I}}_i^t$ there is no $s_i' \in S_i$ that weakly conditionally dominates s_i at I_i.

Losing Bidder Information Sets

Case 1. Consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is a losing bidder and $B(I_i)$ does not contain any element strictly smaller than v. Then the maximum payoff player i can obtain is 0. Since s_i satisfies condition (i), s_i achieves 0 payoff and hence cannot be weakly conditionally dominated at I_i.

Case 2. Consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is a losing bidder, $B(I_i)$ contains an element strictly lower than v, $t < t^*$ and $s_i(I_i) = a$. Suppose by contradiction that there exists some $s_i' \in S_i$ that weakly conditionally dominates s_i at I_i. Then $s_i'(I_i) \neq a$. Because $B(I_i)$ contains an element strictly lower than v, there exists $I_{-i} \in \mathcal{I}_{-i}(I_i)$ be such that $B(I_i, I_{-i}) < v$. Define s_{-i} with continuation strategies following I_{-i} such that if either the price or winning bidder identity changes at exactly t then all players in $N/\{i\}$ bid v at the first bidding opportunity, while if both the price nor the winning bidder identity remains unchanged at t, then all players in $N/\{i\}$ abstain at all times past t. Then $u_i(s_i', s_{-i} \mid I_i, I_{-i}) \leq (v - B(I_i, I_{-i}))e^{\lambda(n-1)t}$, while $u_i(s_i, s_{-i} \mid I_i, I_{-i}) \geq (v - B(I_i, I_{-i}))\left(1 - e^{\lambda t^* \lambda}\right)$ because s_i satisfies condition (ii). Since $t < t^*$, the latter lower bound is strictly higher than the former upper bound, contradicting that s_i' weakly conditionally dominates s_i at I_i.

Case 3. Consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is a losing bidder, $B(I_i)$ contains an element strictly lower than v, $t < t^*$, and $s_i(I_i) = b \neq a$. Let p be the current price at information set I_i. Suppose by way of contradiction that there exists some $s_i' \in S_i$ that weakly conditionally dominates s_i at I_i.

First consider the case in which $b < v$. Then there exists some $I_{-i} \in \mathcal{I}_{-i}(I_i)$ be such that $b = B(I_i, I_{-i})$. Define $s_{-i} \in S_{-i}$ that specify continuation strategies after I_{-i} such that if at t, there is a *direct jump* of the winning price from p to b without a change in the identity of the winning bidder, all bidders in $N \setminus \{i\}$ always abstain after t, and otherwise bid v whenever possible. Then $u_i(s_i', s_{-i} \mid I_i, I_{-i}) \leq (v-b)e^{\lambda(n-1)t}$ while $u_i(s_i, s_{-i} \mid I_i, I_{-i}) \geq (v-b)\left(1 - e^{\lambda t^*}\right)$ because s_i satisfies condition (ii). Since $t < t^*$, the latter lower bound is strictly higher than the former upper bound, contradicting that s_i' weakly conditionally dominates s_i at I_i.

Next consider $b = v$. First consider $I_{-i} \in \mathcal{I}_{-i}(I_i)$ be such that $B(I_i, I_{-i}) = v - \Delta$. Now consider strategies $s_{-i} \in S_{-i}$ that specify continuation strategies after I_{-i} such that if at time t, the identity of the winning bidder changes and there is either a *direct jump* of the winning price from p to $v - \Delta$ if the current price $p < v - \Delta$ or no price change if

$p = v - \Delta$, then all bidders in $N \setminus \{i\}$ abstain after t, while otherwise they bid v whenever possible. Note then that against the strategy profile s_{-i}, $u_i(s_i, s_{-i} \mid I_i, I_{-i}) = \Delta$ while any strategy specifying either a or some bid strictly lower than v yields a continuation payoff of at most $\Delta e^{\lambda(n-1)t} < \Delta$. Hence, $s_i'(I_i) = b' > v$. But then consider $I_{-i} \in \mathcal{I}_{-i}(I_i)$ such that $B(I_i, I_{-i}) = v$ and consider strategies $s_{-i} \in S_{-i}$ that specify continuation strategies after I_{-i} such that every other player bids b' whenever possible (and otherwise abstains). In this case s_i' against s_{-i} yields a strictly negative continuation payoff. Meanwhile, because of condition (i), s_i yields a continuation payoff of zero, contradicting that s_i' weakly conditionally dominates s_i at I_i.

Case 4. Consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is a losing bidder, $B(I_i)$ contains an element strictly lower than v, and $t \geq t^*$. Let p be the current price at I_i. Suppose by contradiction that there exists some $s_i' \in S_i$ that weakly conditionally dominates s_i at I_i. By condition (ii), $s_i(I_i) = b \neq a$. First suppose that $b < v$. Let $I_{-i} \in \mathcal{I}_{-i}(I_i)$ be such that $B(I_i, I_{-i}) = b$, and consider strategies $s_{-i} \in S_{-i}$ that specify continuation strategies after I_{-i} such that if at time t, there is a *direct jump* of the winning price from p to b without the identity of the winning bidder changing, then all bidders in $N \setminus \{i\}$ abstain after t, while otherwise they bid v at the first bidding opportunity. Given condition (ii) in the statement of the claim, by playing s_i, bidder i after placing a bid of b immediately places another bid higher than b (since b is not enough to become the winning bidder). As a result, $u_i(s_i, s_{-i} \mid I_i, I_{-i}) = v - b$. Since $s_i'(I_i) \neq b$, by playing s_i', the first price change leads to either a change in the winning bidder or to a change in price not equal to b. Thus $u_i(s_i', s_{-i} \mid I_i, I_{-i}) \leq (v-b)e^{\lambda(n-1)t} < v-b$. This contradicts the assumption that s_i' weakly conditionally dominates s_i at I_i. Deriving a contradiction for the case $b = v$ is exactly analogous to Case 3.

Winning Bidder Information Sets

Case 5. Consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is the winning bidder with current highest bid $B < v - \Delta$. Suppose there exists some $s_i' \in S_i$ that weakly conditionally dominates s_i at I_i. By condition (iii), note that $s_i(I_i) = a$ and therefore, $s_i'(I_i) \neq a$. Let p be the current price at information sets I_i. Consider any $I_{-i} \in \mathcal{I}_{-i}(I_i)$, and strategies $s_{-i} \in S_{-i}$ that specify continuation strategies after I_{-i} such that whenever possible, all bidders in $N \setminus \{i\}$ bid $B + \Delta$ after which, if such a bid results in becoming the winning bidder then abstains for the rest of the auction, and otherwise immediately bids v. Then $u_i(s_i', s_{-i} \mid I_i, I_{-i}) \leq (v-p)e^{\lambda(n-1)t}$. Meanwhile, because s_i satisfies condition (ii),

$$u_i(s_i, s_{-i} \mid I_i, I_{-i}) \geq (v-p)e^{\lambda(n-1)t} + \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau - t)}(v - B - \Delta)\left(1 - e^{\lambda(t^* \vee \tau)}\right) d\tau.$$

Note that the last term above is strictly positive, which then contradicts the assumption that s_i' weakly conditionally dominates s_i at I_i.

Case 6. Consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is the winning bidder with $B = v - \Delta$. Suppose there exists some $s_i' \in S_i$ weakly conditionally dominating s_i at I_i. By condition (iii), $s_i(I_i) = a$ and therefore, $s_i'(I_i) \neq a$. Then for any $I_{-i} \in \mathcal{I}_{-i}(I_i)$ and any $s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i} \mid I_i, I_{-i}) \geq u_i(s_i', s_{-i} \mid I_i, I_{-i})$. This is because in the continuation game following (I_i, I_{-i}), against any $s_{-i} \in S_{-i}$, s_i and s_i' do not influence the path of information sets reached (and hence actions played) by opponents until one of the opponents bids v or higher. But in such an event, s_i yields 0 continuation payoff, while s_i' yields a continuation payoff bounded from above by 0. This contradicts the assumption that s_i' weakly conditionally dominates s_i at I_i.

Case 7. Finally consider any $I_i \in \widehat{\mathcal{I}}_i^t$ for which i is the winning bidder and $B \geq v$. Let p be the current price at I_i. Suppose that there exists some $s_i' \in S_i$ weakly conditionally dominating s_i at I_i. Again by condition (iii), $s_i(I_i) = a$ and therefore, $s_i'(I_i) = b \neq a$. Note that $b > B \geq v$. Now consider any $I_{-i} \in \mathcal{I}_{-i}(I_i)$ and a strategy profile s_{-i} in which bidders in $N \setminus \{i\}$ all bid b whenever possible (and otherwise abstain). Note that by condition (iii), $u_i(s_i, s_{-i} \mid I_i, I_{-i}) = (v - p)e^{\lambda(n-1)t}$. On the other hand, playing s_i' yields a payoff of at most $(v-p)e^{\lambda(n-1)t} + \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau - t)}(v - b) < (v-p)e^{\lambda(n-1)t}$. This then contradicts the assumption that s_i' weakly conditionally dominates s_i at I_i. ∎

A.2 Proof of Claim 2.3

Proof. Suppose that $I_i^* \in \widehat{\mathcal{I}}_i^t$ and $s_i(I_i^*) = b > v$. Define s_i' such that $s_i'(I_i) = s_i(I_i)$ for all $I_i \in \widehat{\mathcal{I}}_i$ except at I_i^* and all successors of I_i^* in $\widehat{\mathcal{I}}_i$ at which s_i specifies placing a bid strictly higher than v. For the latter information sets, let s_i' specify bidding v if feasible and otherwise abstaining. Note that conditional on I_i^*, playing s_i versus s_i' does not affect which opponents' information sets are reached nor the opponents' actions until the first time some player in $N \setminus \{i\}$ places a bid strictly higher than v. Conditional on such an event after I_i^*, the highest payoff i can obtain is 0, and s_i' as defined above guarantees this payoff. This concludes that conditional on I_i^* s_i' yields a weakly higher continuation payoff than s_i against any $s_{-i} \in S_{-i}$.

Suppose first that at I_i^*, i is a losing bidder. Now consider $I_{-i}^* \in \mathcal{I}_{-i}(I_i^*)$ such that $B(I_i^*, I_{-i}^*) \leq v$ and consider strategies s_{-i} that after I_{-i}, specify placing a bid $v + \Delta$ whenever possible and otherwise abstain. Then $u_i(s_i', s_{-i} \mid I_i^*, I_{-i}^*) = (v - B(I_i^*, I_{-i}^*))e^{\lambda(n-1)t}$. On the other hand, $u_i(s_i, s_{-i} \mid I_i^*, I_{-i}^*) = (v - B(I_i^*, I_{-i}^*))e^{\lambda(n-1)t} - \Delta\left(1 - e^{\lambda(n-1)t}\right)$ which is strictly smaller. Therefore s_i' weakly conditionally dominates s_i at I_i^*.

Now suppose that at I_i^*, i is the winning bidder with current high bid B and current price p. Consider any $I_{-i}^* \in \mathcal{I}_{-i}(I_i^*)$. and consider strategies s_{-i} that after I_{-i}, specify placing a bid $v+\Delta$ whenever possible and otherwise abstain. Then $u_i(s_i', s_{-i} \mid I_i^*, I_{-i}^*) = (v-p)e^{\lambda(n-1)t}$ while $u_i(s_i, s_{-i} \mid I_i^*, I_{-i}^*) = (v-p)e^{\lambda(n-1)t} - \Delta\left(1-e^{\lambda(n-1)t}\right)$ which is again strictly smaller. Therefore, s_i' again weakly conditionally dominates s_i at I_i^*. ∎

A.3 Proof of Claim 2.4

Proof. Suppose that $I_i^* \in \widehat{\mathcal{I}}_i^t$ such that $t \geq t^*$ and $B(I_I^*)$ contains an element strictly smaller than v. Suppose also that $s_i(I_i^*) = a$. We now construct $s_i' \in S_i$ that weakly dominates s_i at I_i^*.

Define s_i' such that $s_i'(I_i) = s_i(I_i)$ for all $I_i \in \widehat{\mathcal{I}}_i$ except at I_i^* and all successors of I_i^* in $\widehat{\mathcal{I}}_i$. At I_i^* let s_i' specify bidding v, and at all successors of I_i^* let s_i' specify abstaining. First note that for all $I_{-i} \in \mathcal{I}_{-i}(I_i^*)$ for which $B(I_i^*, I_{-i}) \geq v$, the highest possible payoff any strategy can yield for i conditional on (I_i^*, I_{-i}) is 0. Note that s_i' guarantees zero payoff, and hence conditional on (I_i^*, I_{-i}), s_i' yields weakly higher payoff against any $s_{-i} \in S_{-i}$ than s_i. Consider now any $I_{-i} \in \mathcal{I}_{-i}(I_i^*)$ for which $B(I_i^*, I_{-i}) < v$. Then $u_i(s_i', s_{-i} \mid I_i^*, I_{-i}) \geq (v - B(I_i^*, I_{-i}))e^{\lambda(n-1)t}$ for all s_{-i} and all I_{-i}. On the other hand, because $s_i(I_i^*) = a$, $u_i(s_i, s_{-i} \mid I_i^*, I_{-i}) \leq (v - B(I_i^*, I_{-i}))\left(1-e^{\lambda(n-1)t}\right)$, which corresponds to payoff in the best case scenario for player i in which all opponents always abstain after time t. By the definition of t^* and that $t \geq t^*$, the above implies that $u_i(s_i', s_{-i} \mid I_i^*, I_{-i}) \geq u_i(s_i, s_{-i} \mid I_i^*, I_{-i})$ for all s_{-i} and all $I_{-i} \in \mathcal{I}_{-i}(I_i^*)$.

Consider now $s_{-i} \in S_{-i}$ such that all other players abstain from bidding after t and let $I_{-i}^* \in \mathcal{I}_{-i}(I_i^*)$ be such that $B(I_i^*, I_{-i}^*) < v$. Then as argued above, $u_i(s_i, s_{-i} \mid I_i^*, I_{-i}^*) \leq (1-e^{\lambda t})(v - B(I_i^*, I_{-i}^*))$. On the other hand, s_i' yields a strictly higher payoff of $v - B(I_i^*, I_{-i}^*)$. This concludes that s_i' weakly conditionally dominates s_i at I_i^*. ∎

A.4 Proof of Claim 2.5

Proof. We first prove that the truthful strategy profile is indeed an equilibrium. First note that conditions (i), (ii), and (iii) of Claim 2.2 are all satisfied so that strategies are weakly conditionally undominated. Now consider the incentives of a bidder i. Suppose first that at time t and price p, i is the winning bidder. If $p > v$, then regardless of what i plays, i obtains a payoff of $p - v$ and so it is optimal to abstain. If $p \leq v$, then regardless of what i plays, i obtains a payoff of $(v-p)e^{\lambda(n-1)t}$. Thus if $B \geq v$, then it is optimal to abstain while if $B < v$, it is optimal to place a bid of v.

Now consider the incentives of a losing bidder i. Suppose that $p \geq v$. Then regardless of the strategy that i plays, bidder i obtains a payoff of 0 and so it is in particular optimal to abstain. Suppose instead that $p = \emptyset, 0, \Delta, \ldots, v - \Delta$. Note that since i does not observe the current highest bid, he forms some belief μ about about the current highest bid. We will show that bidding v is always optimal for any belief μ. To see this, note that by bidding v, player i obtains a payoff of $\sum_{k=0}^{(v-p)/\Delta} \mu(p+k\Delta)(v-p-k\Delta)e^{\lambda(n-1)t}$. A one-stage deviation to playing a instead yields a payoff of $\left(1-e^{\lambda t}\right) \sum_{k=0}^{(v-p)/\Delta} \mu(p+k\Delta)(v-p-k\Delta)e^{\lambda(n-1)t}$ which is clearly suboptimal. Placing a bid of $b < v$ yields a payoff of:

$$\sum_{k=0}^{(b-p)/\Delta} \mu(p+k\Delta)(v-p-k\Delta)e^{\lambda(n-1)t} + \left(1-e^{\lambda t}\right) \sum_{k=(b-p)/\Delta+1}^{(v-p)/\Delta} (v-p-k\Delta)e^{\lambda(n-1)t}$$

which is again suboptimal. Finally placing a bid $b > v$ yields a payoff of $\sum_{k=0}^{(b-p)/\Delta} \mu(p+k\Delta)(v-p-k\Delta)e^{\lambda(n-1)t}$, which is again suboptimal. Thus we have shown that the truthful strategy profile is indeed an equilibrium.

Note that by Claim 2.3, the price never rises above v in any equilibrium. Therefore, an upper bound for the revenue to the seller in the auction $\Gamma(v, n, \lambda, \Delta, T)$ is given by $v \int_T^0 \lambda n e^{-\lambda n(\tau - T)} \left(1 - e^{\lambda(n-1)\tau}\right) d\tau$. But this is exactly the revenue generated by the truthful equilibrium.

Finally note that in any equilibrium in weakly conditionally undominated strategies, bidder i obtains at least a payoff of $v\left(1-e^{\lambda T}\right) e^{\lambda(n-1)T}$. This is because bidder i can guarantee this payoff in any equilibrium by placing a bid at his first bidding opportunity. But note that this is the payoff that bidder i obtains in the truthful equilibrium. This concludes the proof of Claim 2.5. ∎

B Preliminaries for Equilibrium Analysis

Recall the definitions of $W_\Delta(p,t)$ and $L_\Delta(p,t)$ from the main text. We now proceed to explicitly compute these value functions.

To calculate these value functions, we look at the probability of the following events conditional on every player playing s^g from time t until the end of the auction at time 0. We define $E_k^W(t)$ to be the probability of the event in which conditional on starting time t, (and highest bid $p + \Delta$), k bidding opportunities for the losing bidders arrive in $[t, 0]$, and the winning bidder at time t is the last bidder to obtain a bidding opportunity.[34] Similarly we define $E_k^L(t)$ to be the probability of the event in which conditional on starting time t, k bidding opportunities for the losing bidders arrive in $[t, T]$ and a particular losing bidder at

[34]Note that we can ignore the bidding opportunities of the winning bidder since she never places a bid.

time t is also last bidder to obtain a bidding opportunity.[35]

Given these definitions, it is clear that

$$W_\Delta(p,t) = \sum_{k=0}^{(v-p)/\Delta} E_k^W(t)(v-p-k\Delta), \quad L_\Delta(p,t) = \sum_{k=1}^{(v-p)/\Delta} E_k^L(t)(v-p-k\Delta).$$

In the Supplementary Appendix, we compute $E_k^W(t)$ and $E_k^L(t)$ explicitly to show that

$$\begin{aligned} E_k^W(t) &= e^{\lambda(n-1)t}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} + (-1)^k\frac{n-1}{n}\right), \\ E_k^L(t) &= e^{\lambda(n-1)t}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} - (-1)^k\frac{1}{n}\right). \end{aligned}$$

As a result, we have the following expressions for $W_\Delta(p,t)$ and $L_\Delta(p,t)$.

Corollary B.1. *For every $p = 0, 1, 2, \dots, v$,*

$$\begin{aligned} W_\Delta(p,t) &= e^{\lambda(n-1)t}\sum_{k=0}^{(v-p)/\Delta}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} + (-1)^k\frac{n-1}{n}\right)(v-p-k\Delta), \\ L_\Delta(p,t) &= e^{\lambda(n-1)t}\sum_{k=0}^{(v-p)/\Delta}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} - (-1)^k\frac{1}{n}\right)(v-p-k\Delta). \end{aligned}$$

Given the value functions derived above, the following properties of these value functions will be useful for our analysis.

Lemma B.2. *For every $p = 0, 1, 2, \dots, v-\Delta$ and every t,*

$$L_\Delta(p,t) \geq L_\Delta(p+\Delta,t), W_\Delta(p,t) \geq W_\Delta(p+\Delta,t).$$

Proof. This is clear from Corollary B.1:

$$\begin{aligned} L_\Delta(p,t) &= e^{\lambda(n-1)t}\sum_{k=0}^{(v-p)/\Delta}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} - (-1)^k\frac{1}{n}\right)(v-p-k\Delta) \\ &\geq e^{\lambda(n-1)t}\sum_{k=0}^{(v-p)/\Delta-1}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} - (-1)^k\frac{1}{n}\right)(v-p-(k+1)\Delta) = L_\Delta(p+\Delta,t) \end{aligned}$$

where the inequality follows from the observation that for all k, $\left(\frac{(n-1)^k}{n} - (-1)^k\frac{1}{n}\right) \geq 0$. Furthermore, the inequality for the winning bidder trivially holds when $p = v-\Delta$. So

[35] Note that this event does not depend on the particular identity of the losing bidder at time t since all arrival rates of losing bidders are symmetric.

assume that $p < v - \Delta$. Then by the fact that $L_\Delta(p+\Delta,\tau) \geq L_\Delta(p+2\Delta, 2\tau)$ for all τ,

$$\begin{aligned} W_\Delta(p,t) &= \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau-t)} L_\Delta(p+\Delta,\tau)d\tau \\ &\geq \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau-t)} L_\Delta(p+2\Delta,\tau)d\tau = W_\Delta(p+\Delta,t). \end{aligned}$$

This concludes the proof. ∎

C Omitted Part of Proof of Theorem 3.2

Proof. It remains to show that when $n = 2$, $T^* = -1/\lambda$. To this end, fix any arbitrary $\Delta > 0$. First note that for all $t < -1/\lambda$,

$$W_\Delta(v-\Delta,t) = \Delta e^{\lambda t} < \Delta\lambda|t|e^{\lambda(n-1)t} = L_\Delta(v-2\Delta,t).$$

Thus $T^* \geq -1/\lambda$. To see that $T^* \leq -1/\lambda$, we show that $W_\Delta(p+\Delta,t) \geq L_\Delta(p,t)$ for all $p = \emptyset, 0, \Delta, v-\Delta$ and all $t \geq -1/\lambda$.

Consider any $t \geq -1/\lambda$. First suppose that $p = v-\Delta$. Then $W_\Delta(v,t) = L_\Delta(v-\Delta,t) = 0$ for all t and so the inequality holds trivially. Secondly, consider any price $p = 0, \Delta, \ldots, v-2\Delta$. Then by Corollary B.1,

$$\begin{aligned} W_\Delta(p+\Delta,t) &= e^{\lambda(n-1)t} \sum_{k=0}^{(v-p)/\Delta-2} \frac{(\lambda|t|)^k}{k!}\left(\frac{1}{2} + (-1)^k\frac{1}{2}\right)(v-p-(k+1)\Delta), \\ L_\Delta(p,t) &= e^{\lambda(n-1)t} \sum_{k=1}^{(v-p)/\Delta-1} \frac{(\lambda|t|)^k}{k!}\left(\frac{1}{2} - (-1)^k\frac{1}{2}\right)(v-p-k\Delta) \\ &= e^{\lambda(n-1)t} \sum_{k=0}^{(v-p)/\Delta-2} \frac{(\lambda|t|)^{(k+1)}}{(k+1)!}\left(\frac{1}{2} + (-1)^k\frac{1}{2}\right)(v-p-(k+1)\Delta). \end{aligned}$$

Thus,

$$\begin{aligned} &W_\Delta(p+\Delta,t) - L_\Delta(p,t) \\ &= e^{\lambda(n-1)t} \sum_{k=0}^{(v-p)/\Delta-2} \frac{(\lambda|t|)^k}{k!}\left(1 - \frac{\lambda|t|}{k+1}\right)\left(\frac{1}{2} + (-1)^k\frac{1}{2}\right)(v-p-(k+1)\Delta). \end{aligned}$$

Note that the latter is weakly positive since $\lambda|t| \leq 1$.

We now prove that $L_\Delta(\emptyset, t) \leq W_\Delta(0, t)$ for all $t \geq -1/\lambda$. Consider the following function $\hat{L}_\Delta(\emptyset, t)$ defined as follows: $\hat{L}_\Delta(\emptyset, t) = \int_t^0 \lambda e^{-\lambda(\tau - t)} W_\Delta(0, \tau) d\tau$. We will use this value function as a bound for $L_\Delta(\emptyset, t)$. Note that an easy computation together with Corollary B.1 yield the following:

$$\begin{aligned}
\hat{L}_\Delta(\emptyset, t) &= e^{\lambda t} \sum_{k=1}^{v/\Delta} \frac{(\lambda|t|)^k}{k!} \left(\frac{1}{2} - (-1)^k \frac{1}{2} \right) (v - (k-1)\Delta) \\
&\geq e^{\lambda t} \sum_{k=1}^{v/\Delta} \frac{(\lambda|t|)^k}{k!} \left(\frac{1}{2} - (-1)^k \frac{1}{2} \right) (v - k\Delta) = L_\Delta(0, t).
\end{aligned}$$

Furthermore, by the recursive representation of $\hat{L}_\Delta(\emptyset, t)$, we obtain:

$$\begin{aligned}
\hat{L}_\Delta(\emptyset, t) &= \int_t^0 \lambda e^{-2\lambda(\tau - t)} \left(W_\Delta(0, \tau) + \hat{L}_\Delta(\emptyset, \tau) \right) d\tau \\
&\geq \int_t^0 \lambda e^{-2\lambda(\tau - t)} \left(W_\Delta(0, \tau) + L_\Delta(0, \tau) \right) d\tau = L_\Delta(\emptyset, t).
\end{aligned}$$

Finally, note that an explicit computation of $\hat{L}_\Delta(\emptyset, t)$ using the same arguments as in Corollary B.1 yields:

$$\begin{aligned}
W_\Delta(0, t) &= e^{\lambda t} \sum_{k=0}^{v/\Delta - 1} \frac{(\lambda|t|)^k}{k!} \left(\frac{1}{2} + (-1)^k \frac{1}{2} \right) (v - k\Delta), \\
\hat{L}_\Delta(\emptyset, t) &= e^{\lambda t} \sum_{k=1}^{v/\Delta} \frac{(\lambda|t|)^k}{k!} \left(\frac{1}{2} - (-1)^k \frac{1}{2} \right) (v - (k-1)\Delta) \\
&= e^{\lambda t} \sum_{k=0}^{v/\Delta - 1} \frac{(\lambda|t|)^{k+1}}{(k+1)!} \left(\frac{1}{2} + (-1)^k \frac{1}{2} \right) (v - k\Delta)
\end{aligned}$$

Therefore,

$$W_\Delta(0, t) - \hat{L}_\Delta(\emptyset, t) = e^{\lambda t} \sum_{k=0}^{v/\Delta - 1} \frac{(\lambda|t|)^k}{k!} \left(1 - \frac{\lambda|t|}{k+1} \right) \left(\frac{1}{2} + (-1)^k \frac{1}{2} \right) (v - k\Delta)$$

which is again weakly positive for all $|t| \leq 1/\lambda$. This implies that $W_\Delta(0, t) \geq \hat{L}_\Delta(\emptyset, t) \geq L_\Delta(\emptyset, t)$ for all $t \geq -1/\lambda$. This concludes the proof. ∎

D Proof of Proposition 4.3

Recall the definition of $\tilde{L}_\Delta(\emptyset, t)$ from the main text:

$$\tilde{L}_\Delta(\emptyset, t) = \int_t^0 \lambda e^{-\lambda n(\tau - t)} \left(W_\Delta(0, \tau \vee \tau_0^\Delta) + (n-1) L_\Delta(0, \tau \vee \tau_0^\Delta) \right) d\tau.$$

We begin with a lemma that shows the uniqueness of τ_0^Δ and $\tau_\emptyset^\Delta$ defined in the main text.

Lemma D.1. *There exist unique $\tau_0^\Delta < 0$ and $\tau_\emptyset^\Delta < 0$ that satisfy the following:*

$$(v - \Delta) e^{\lambda(n-1)\tau_0^\Delta} = L_\Delta(0, \tau_0^\Delta), \, v e^{\lambda(n-1)\tau_\emptyset^\Delta} = \tilde{L}_\Delta(\emptyset, \tau_\emptyset^\Delta).$$

Furthermore $(v-\Delta)e^{\lambda(n-1)t} > L_\Delta(0,t)$ if and only if $t \in (\tau_0^\Delta, 0]$ and $v e^{\lambda(n-1)t} > \tilde{L}_\Delta(\emptyset, t)$ if and only if $t \in (\tau_\emptyset^\Delta, 0]$.

Proof. We first show the existence and uniqueness of τ_0^Δ. By Corollary B.1,

$$e^{-\lambda(n-1)t} L_\Delta(0, t) = \sum_{k=0}^{(v-p)/\Delta} \frac{(\lambda |t|)^k}{k!} \left(\frac{(n-1)^k}{n} - (-1)^k \frac{1}{n} \right) (v - p - k\Delta)$$

Therefore existence is immediate since $\lim_{t \to 0} e^{-\lambda(n-1)t} L_\Delta(0,t) = 0$ and $\lim_{t\to\infty} e^{-\lambda(n-1)t} L_\Delta(0,) = +\infty$. Moreover, note that $e^{-\lambda(n-1)t} L_\Delta(0,t)$ is strictly decreasing in t (strictly increasing in $|t|$) and therefore, τ_0^Δ must be unique.

Let us now prove the existence and uniqueness of $\tau_\emptyset^\Delta$. By Corollary B.1 and the definition of $\tilde{L}_\Delta(\emptyset, t)$,

$$e^{-\lambda(n-t)t} \tilde{L}_\Delta(\emptyset, t) = \sum_{k=0}^{v/\Delta} \int_t^0 \lambda e^{-\lambda(\tau - t)} \frac{(\lambda(n-1)|\tau \vee t_0^\Delta|)^k}{k!} (v - k\Delta) d\tau.$$

Define for each $k = 0, 1, \ldots, v/\Delta$, $f_k : \mathbb{R} \to \mathbb{R}$ as follows:

$$f_k(\tau) = \begin{cases} \frac{(\lambda(n-1)|\tau \vee \tau_0^\Delta|)^k}{k!}(v - k\Delta) & \text{if } \tau \le 0, \\ 0 & \text{if } \tau > 0. \end{cases}$$

Then note that $e^{-\lambda(n-1)t} \tilde{L}_\Delta(\emptyset, t) = \sum_{k=0}^{v/\Delta} \int_t^\infty \lambda e^{-\lambda(\tau - t)} f_k(\tau) d\tau$. As a result, $\lim_{t\to 0} e^{-\lambda(n-1)t} \tilde{L}_\Delta(\emptyset, t) = 0$ and $\lim_{t \to +\infty} e^{-\lambda(n-1)t} \tilde{L}_\Delta(\emptyset, t) = +\infty$ and so existence follows. To see uniqueness, note that f_k is a weakly decreasing function which is strictly decreasing in the region $[\tau_0^\Delta, 0]$. Therefore, $e^{-\lambda(n-1)t} \tilde{L}_\Delta(\emptyset, t)$ is strictly decreasing in t and hence, $\tau_\emptyset^\Delta$ is unique. ∎

Before proceeding to the proof of Proposition 4.3, the following lemma is useful.

Lemma D.2. *For every $p = 0, \Delta, \ldots, v - 2\Delta$ and every t, $\frac{L_\Delta(p,t)}{v-p-\Delta} \leq \frac{L_\Delta(0,t)}{v-\Delta}$.*

Proof. We know that $\frac{L_\Delta(p,t)}{v-p-\Delta} = \sum_{k=1}^{(v-p)/\Delta} E_k^L(t)\frac{v-p-k\Delta}{v-p-\Delta}$ and $\frac{L_\Delta(0,t)}{v-\Delta} = \sum_{k=1}^{v/\Delta} E_k^L(t)\frac{v-k\Delta}{v-\Delta}$. Note that for each k, $\frac{v-p-k\Delta}{v-p-\Delta} \leq \frac{v-k\Delta}{v-\Delta}$, which implies that $\frac{L_\Delta(p,t)}{v-p-\Delta} \leq \frac{L_\Delta(0,t)}{v-\Delta}$. ∎

We are now ready to prove Proposition 4.3.

Proof of Proposition 4.3. We need to prove that $s^{\Delta,\mathcal{T}^\Delta}$ is indeed an equilibrium where $\mathcal{T}^\Delta = (\tau_\emptyset^\Delta, \tau_0^\Delta, \ldots, \tau_0^\Delta)$. First we show that all sufficient conditions for conditional weak undomination are satisfied. Note that conditions (i) and (iii) from Claim 2.2 are satisfied trivially. To prove that condition (ii) from Claim 2.2 is also satisfied, it is sufficient to show that $\tau_0^\Delta, \tau_\emptyset^\Delta \leq t^*$. Recall that $(v-\Delta)e^{\lambda(n-1)\tau_0^\Delta} = L_\Delta(0,\tau_0^\Delta)$ and $ve^{\lambda(n-1)\tau_\emptyset^\Delta} = \tilde{L}_\Delta(0,\tau_\emptyset^\Delta)$.

It is clear that $L_\Delta(0,\tau_0^\Delta) \leq (v-\Delta)\left(1 - e^{\lambda\tau_0^\Delta}\right)$ since the latter is the continuation value to the best case scenario in which all opposing bidders cease to place bids in the future. Similarly, $\tilde{L}_\Delta(0,\tau_\emptyset^\Delta) \leq v\left(1 - e^{\lambda\tau_\emptyset^\Delta}\right)$. Thus we have $e^{\lambda(n-1)\tau_0^\Delta} \leq \left(1 - e^{\lambda\tau_0^\Delta}\right)$ and $e^{\lambda(n-1)\tau_\emptyset^\Delta} \leq \left(1 - e^{\lambda\tau_\emptyset^\Delta}\right)$. As a result, $\tau_0^\Delta, \tau_\emptyset^\Delta \leq t^*$.

It remains to show that $s_i^{\Delta,\mathcal{T}^\Delta}$ for bidder i satisfies incentive compatibility at all histories h^t. There are four cases to consider: 1) $i = J(h^t)$; 2) $i \neq J(h^t)$ and $h^t \notin \mathcal{H}^{t,C}$; 3) $i \neq J(h^t)$ and $h^t \in \mathcal{H}^{t,C,a}$; and 4) $i \neq J(h^t)$ and $h^t \in \mathcal{H}^{t,C,b}$.

Case 1: Suppose that $i = J(h^t)$ and suppose first that $h^t \notin \mathcal{H}^{t,C}$. Then regardless of what bidder i does at such a history, i obtains a payoff of $(v - P(h^t))e^{\lambda(n-1)t}$. Thus it is optimal to choose a. Suppose on the other hand that $h^t \in \mathcal{H}^{t,C}$. In this case, if $p = P(h^t)$, then the highest current bid (which is observed by i) will be $p + \Delta$.

If $p = v-\Delta$, incentives are trivial. If $p = \Delta, \ldots, v-2\Delta$, then since $h^t \in \mathcal{H}^{t,C}$, $t \geq \tau_0^\Delta$. At such a history, by playing a, bidder i can guarantee himself at least a payoff of $(v-p)e^{\lambda(n-1)t}$. On the other hand, by deviating to place some bid $b \geq p + 2\Delta$, bidder i' deviation will be detected upon the next bidding opportunity by a losing bidder at which time bidders revert to truthful bidding. Thus, the most that bidder i can obtain to deviating is $(v-p)e^{\lambda(n-1)t}$, showing that it is optimal to play a.

Finally if $p = 0$, then bidder i by playing a can guarantee himself at a least payoff of $ve^{\lambda(n-1)(t\vee\tau_0^\Delta)}$. By deviating to placing some bid $b \geq 2\Delta$, bidder i's deviation will be detected upon the next bidding opportunity arriving to a losing bidder after time τ_0^Δ after which all bidders revert to truthful bidding. Thus, the most that bidder i can obtain to placing a bid is $ve^{\lambda(n-1)(t\vee\tau_0^\Delta)}$, again showing that is optimal to play a.

Case 2: Suppose now that $i \neq J(h^t)$ and $h^t \notin \mathcal{H}^{t,C}$. Suppose that bidder i has belief $\mu \in \mathcal{M}(\{p, p+\Delta, \ldots, v\})$ about the highest bid at history h^t. Then bidder i's expected payoff to bidding v is $e^{\lambda(n-1)t}\sum_{k=0}^{(v-p)/\Delta}\mu(p+k\Delta)(v-p-k\Delta)$. On the other hand, bidding until $b<v$ and then playing a followed by s_i^{Δ,C^Δ} in the future yields a payoff of:

$$e^{\lambda(n-1)t}\left(\sum_{k=0}^{(b-p)/\Delta-1}\mu(p+k\Delta)(v-p-k\Delta)+\left(1-e^{\lambda t}\right)\sum_{k=(b-p)/\Delta}^{(v-p)/\Delta}\mu(p+k\Delta)(v-p-k\Delta)\right).$$

Finally playing a and following s^{Δ,C^Δ} in the future yields a payoff of $\sum_{k=0}^{(v-p)/\Delta}\mu(p+k\Delta)(v-p-k\Delta)(1-e^{\lambda t})e^{\lambda(n-1)t}$. Thus, bidding v is optimal.

Case 3: Suppose next that $i\neq J(h^t)$ and $h^t\in\mathcal{H}^{t,C,a}$. Note that such histories occur either when $P(h^t)=0$ and $t<\tau_0^\Delta$, $P(h^t)=\emptyset$ and $t\leq\tau_\emptyset^\Delta$, or $P(h^t)=v$. First consider the case in which $P(h^t)=0$ and $t<\tau_0^\Delta$. First note that since $h^t\in\mathcal{H}^{t,C}$, losing bidders believe that the highest bid is Δ. Then playing a yields a payoff of $L_\Delta(0,\tau_0^\Delta)=(v-\Delta)e^{\lambda(n-1)\tau_0^\Delta}$, where the equality follows from the definition of τ_0^Δ. On the other hand, if bidder i places any bid, then this deviation is immediately detected, triggering reversion to truthful bidding which yields a payoff of at most $(v-\Delta)e^{\lambda(n-1)t}<(v-\Delta)e^{\lambda(n-1)\tau_0^\Delta}$. Thus playing a is optimal.

Suppose next that $P(h^t)=\emptyset$ and $t<\tau_\emptyset^\Delta$. Then playing a yields a payoff of $\tilde{L}_\Delta(\emptyset,\tau_\emptyset^\Delta)=ve^{\lambda(n-1)\tau_\emptyset^\Delta}$, where the equality follows again from the definition of $\tau_\emptyset^\Delta$. On the other hand, by placing any bid, bidder i triggers a reversion to truthful bidding which yields a payoff of $ve^{\lambda(n-1)t}<ve^{\lambda(n-1)\tau_\emptyset^\Delta}$. Again playing a is optimal. Finally if $P(h^t)=v$, then bidders' incentives are trivial and again it is optimal to play a.

Case 4: Finally suppose that $i\neq J(h^t)$ and $h^t\in\mathcal{H}^{t,C,b}$. Suppose first that $P(h^t)=\emptyset$ in which case, because $h^t\in\mathcal{H}^{t,C,b}$, $t\geq\tau_\emptyset^\Delta$. By playing s_i^{Δ,C^Δ}, bidder i obtains a payoff of $W_\Delta(0,t)$, which by Corollary B.1, is weakly greater than $ve^{\lambda(n-1)t}$. On the other hand, any deviation to some bid $b>\Delta$ is detected upon arrival of the next losing bidder at which time all bidders switch to truthful bidding. Thus, any such deviation yields a payoff of $ve^{\lambda(n-1)t}$. Finally by playing a, bidder i obtains a payoff of $\tilde{L}_\Delta(\emptyset,t)$ which by Lemma D.1 is weakly less than $ve^{\lambda(n-1)t}$. Thus bidding Δ is optimal.

Next suppose that $0\leq p=P(h^t)\leq v-\Delta$ in which case, because $h^t\in\mathcal{H}^{t,C,b}$, $t\geq\tau_0^\Delta$. Again since $h^t\in\mathcal{H}^{t,C,b}\subseteq\mathcal{H}^{t,C}$, losing bidders believe the highest bid to be $p+\Delta$. If $p=v-\Delta$, then incentives are trivial, so let us assume that $p<v-\Delta$. If in this case, players follow the equilibrium strategy of $s_i^{\Delta,\mathcal{T}^\Delta}$, then bidder i bids $p+2\Delta$ and obtains $W_\Delta(p+\Delta,t)\geq(v-p-\Delta)e^{\lambda(n-1)t}$ where the inequality follows from Corollary B.1. On the other hand, placing a bid $b\neq p+2\Delta$ leads to a history that is off of the equilibrium path. This then leads to a reversion to truthful bidding in the future which yields a payoff of at most $(v-p-\Delta)e^{\lambda(n-1)t}\leq W_\Delta(p+\Delta,t)$. Furthermore, by playing a, bidder i obtains a payoff of $L_\Delta(p,t)$. By Lemma D.2 and the fact that $t\geq\tau_0^\Delta$, $\frac{L_\Delta(p,t)}{v-p-\Delta}\leq\frac{L_\Delta(0,t)}{v-\Delta}\leq e^{\lambda(n-1)t}$. As

a result, we have $L_\Delta(p,t) \leq (v-p-\Delta)e^{\lambda(n-1)t} \leq W_\Delta(p+\Delta,t)$. Therefore, it is optimal for bidder i to bid $p+2\Delta$ at h^t. This concludes the proof. ∎

References

Ambrus, Attila and Shih En Lu, "A continuous-time model of multilateral bargaining," *American Economic Journal: Microeconomics*, 2015, *7* (1), 208–49.

Arcidiacono, Peter, Patrick Bayer, Jason R Blevins, and Paul B Ellickson, "Estimation of dynamic discrete choice models in continuous time with an application to retail competition," *The Review of Economic Studies*, 2016, *83* (3), 889–931.

Ariely, Dan, Axel Ockenfels, and Alvin E. Roth, "An Experimental Analysis of Ending Rules in Internet Auctions," *The RAND Journal of Economics*, 2005, *36* (4), 890–907.

Bajari, Patrick and Ali Hortaçsu, "Economic Insights from Internet Auctions," *Journal of Economic Literature*, June 2004, *42* (2), 457–486.

Calcagno, Riccardo, Yuichiro Kamada, Stefano Lovo, and Takuo Sugaya, "Asynchronicity and coordination in common and opposing interest games," *Theoretical Economics*, June 2014, *9* (2), 409–434.

Compte, Olivier and Philippe Jehiel, "The Wait-and-See Option in Ascending Price Auctions," *Journal of the European Economic Association*, 2004, *2* (2-3), 494–503.

Doraszelski, Ulrich and Kenneth L. Judd, "Avoiding the curse of dimensionality in dynamic stochastic games," *Quantitative Economics*, March 2012, *3* (1), 53–93.

Ely, Jeffrey C. and Tanjim Hossain, "Sniping and Squatting in Auction Markets," *American Economic Journal: Microeconomics*, August 2009, *1* (2), 68–94.

Fu, Rao, "Are Incremental Bidders Really Naive? Theory and Evidence from Competing Online Auctions," June 2009. Mimeo, University of Texas Austin.

Fuchs, William and Andrzej Skrzypacz, "Bargaining with Arrival of New Traders," *American Economic Review*, June 2010, *100* (3), 802–36.

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

Hendricks, Kenneth, Ilke Onur, and Thomas Wiseman, "Last-Minute Bidding in Sequential Auctions with Unobserved, Stochastic Entry," *Review of Industrial Organization*, Feb 2012, *40* (1), 1–19.

Hossain, Tanjim, "Learning by bidding," *The RAND Journal of Economics*, September 2008, *39* (2), 509–529.

_ **and John Morgan**, "...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay," *Advances in Economic Analysis & Policy*, January 2006, *5* (2).

Kamada, Yuichiro and Michihiro Kandori, "Revision Games," 2009. Mimeo, Harvard University.

Ku, Gillian, Deepak Malhotra, and J. Keith Murnighan, "Towards a competitive arousal model of decision-making: A study of auction fever in live and Internet auctions," *Organizational Behavior and Human Decision Processes*, 2005, *96* (2), 89 – 103.

Ockenfels, Axel and Alvin E. Roth, "Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction," *Games and Economic Behavior*, 2006, *55* (2), 297 – 320. Mini Special Issue: Electronic Market Design.

Peters, Michael and Sergei Severinov, "Internet auctions with many traders," *Journal of Economic Theory*, 2006, *130* (1), 220 – 245.

Rasmusen, Eric B., "Strategic Implications of Uncertainty over One's Own Private Value in Auctions," *Advances in Theoretical Economics*, 2006, *6* (1), 1–22.

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

Zeithammer, Robert and Christopher Adams, "The Sealed-Bid Abstraction in Online Auctions," *Marketing Science*, 2010, *29* (6), 964–987.

E Supplementary Appendix: Not for Publication

E.1 Computation of $E_k^W(t)$ and $E_k^L(t)$

Lemma E.1.

$$E_k^W(t) = e^{\lambda(n-1)t}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} + (-1)^k\frac{n-1}{n}\right),$$
$$E_k^L(t) = e^{\lambda(n-1)t}\frac{(\lambda|t|)^k}{k!}\left(\frac{(n-1)^k}{n} - (-1)^k\frac{1}{n}\right).$$

Proof. To see this, we proceed by induction. Note that clearly $E_0^W(t) = e^{\lambda(n-1)t}$ and $E_0^L(t) = 0$. Suppose that the above is satisfied for k. Then note that

$$\begin{aligned} E_{k+1}^W(t) &= \int_t^0 \lambda(n-1)e^{-\lambda(n-1)(\tau-t)}E_k^L(\tau)d\tau \\ &= \lambda e^{\lambda(n-1)t}\left(\frac{(n-1)^{k+1}}{n} - (-1)^k\frac{n-1}{n}\right)\int_t^0 \frac{(\lambda|\tau|)^k}{k!}d\tau \\ &= e^{\lambda(n-1)t}\frac{(\lambda|t|)^{k+1}}{(k+1)!}\left(\frac{(n-1)^{k+1}}{n} + (-1)^{k+1}\frac{n-1}{n}\right). \end{aligned}$$

Furthermore note that $E_{k+1}^W(t)+(n-1)E_{k+1}^L(t)$ is the probability of $k+1$ bidding opportunities arriving in $[t,0]$ for the losing bidders. Therefore,

$$\begin{aligned} E_{k+1}^L(t) &= \frac{1}{n-1}\left(e^{\lambda(n-1)t}\frac{(\lambda|t|)^{k+1}}{(k+1)!} - E_{k+1}^W(t)\right) \\ &= e^{\lambda(n-1)t}\frac{(\lambda|t|)^{k+1}}{(k+1)!}\left(\frac{(n-1)^{k+1}}{n} - (-1)^{k+1}\frac{1}{n}\right). \end{aligned}$$

This concludes the proof. ∎

E.2 Additional Extensions

In this section we show that the existence of equilibria with gradual bidding and waiting generalizes to the case of bidders with asymmetric valuations, as well as to situations in which bidders are uncertain about the valuations of other bidders. As these environments are analytically more difficult, we restrict our attention to the case of two possible valuations.

E.2.1 Asymmetric Values

Here we consider the case with two bidders, who have commonly known but different valuations for the object. We show that for any pair of valuations in which the lower valuation exceeds the minimum bid Δ, there exists an equilibrium in which the initial bid by a low valuation bidder is gradual. In this equilibrium, the low valuation bidder wins the object with non-trivial frequency even for auctions with an arbitrarily long time horizon.

Proposition E.2. *For any 2-bidder auction with bidder values $v_H > v_L > \Delta$, symmetric arrival rates and $|T|$ sufficiently large, there exists an equilibrium with gradual bidding.*

In our equilibrium construction, the low valuation bidder plays a public strategy in which he bids

1. $\Delta > 0$ at all times whenever $p = \emptyset$,
2. v_L whenever the other player has placed at least one bid and v_L is feasible,
3. and otherwise abstains.

Note that the low valuation bidder places a bid of v_L as long as the other player has placed at least one bid even when he is a winning bidder as long as v_L is a feasible bid. The high valuation bidder plays a public strategy characterized by two threshold times $\tau^H_\emptyset, \tau^H_0 > -\infty$ and $\tau^H_p = -\infty$ for all $p \geq \Delta$. Given these thresholds, she bids

1. v_H if $p = \emptyset$ and $t \geq \tau^H_\emptyset$;
2. v_H if only the opponent placed a bid in the past (so that $p = 0$ and she is the losing bidder) and $t \geq \tau^H_0$;
3. v_H if she previously placed at least one bid and v_H is feasible;
4. and otherwise abstains.

Hence, in this strategy profile the low valuation bidder bids (partially) gradually while the high valuation bidder waits until near the end of the auction to place any bids. The high valuation bidder waits as bidding too early increases the probability that the ultimate winning price is v_L instead of Δ.

Before proceeding to the proof, there are a couple important features of the equilibrium to highlight. The first is that $\tau^H_p < \tau^H_0$ for all $p \geq \Delta$ and therefore, on the path of play, there is only delay at prices of $\emptyset$ and 0. Secondly, if $|T|$ is sufficiently large, because all initial bidding opportunities when no bids have yet been placed by the high valuation bidder are relinquished, with probability close to one, the low valuation bidder initiates bidding at a time $t < \tau_{\emptyset}^H$ by placing a bid of Δ. As a result, the ex-ante expected payoffs in the limit as $T \to -\infty$ of the two bidders are respectively:

$$L_H(\emptyset, -\infty) := \lim_{T \to -\infty} L_H(\emptyset, T) = (v_H - v_L)(1 - e^{\lambda \tau_0^H}) + \lambda |\tau_0^H| e^{\lambda \tau_0^H}(v_L - \Delta),$$
$$L_L(\emptyset, -\infty) := \lim_{T \to -\infty} L_L(\emptyset, T) = v_L e^{\lambda \tau_0^H}.$$

As a result, because of the existence of delay by the high valuation bidder in equilibrium, the low valuation bidder wins the auction with non-trivial probability even as $T \to -\infty$.

To see this more concretely, consider an auction with $v_H = 6$, $v_L = 4$, $\Delta = 1$, $T = -\infty$, and $\lambda = 1$. In the benchmark truthful equilibrium, the high and low valuation bidders respectively bid 6 and 4 at their first opportunity. This implies that as $T \to -\infty$, with probability arbitrarily close to one, the high valuation bidder wins and obtains a payoff of 2, giving the seller a revenue of 4. In contrast, in the equilibrium constructed in Proposition E.2, the low valuation bidder's expected payoff is $v_L e^{\lambda \tau_0^H} = 4e^{-\frac{5}{3}} \approx 0.76$. The low type's chance of winning is $e^{-\frac{5}{3}} \approx 19\%$ whereas the high valuation bidder has a $(1 - e^{-\frac{5}{3}}) \approx 81\%$ chance of winning. The total expected payoff among both bidders is equal to 3.32 (with a payoff of 2.57 to the high type) versus 2 in the benchmark equilibrium, and the seller's expected revenue falls to roughly 2.3. Since the losing bidder places at least one bid with probability one, there is no inefficiency in this equilibrium due to no bidding. There is however inefficiency due to the fact that the low valuation bidder wins the object with some probability.

Proof of Proposition E.2. Let us first construct the threshold times. To this end, define the following value functions:

$$W_H(\Delta, t) = (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t}(v_H - \Delta),$$
$$W_H(0, t) = (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t} v_H,$$
$$\hat{L}_H(0, t) = (v_H - v_L)(1 - e^{\lambda t}) + \lambda |t| e^{\lambda t}(v_L - \Delta).$$

To construct the threshold times, we first define τ_0^H to be the unique time at which $W_H(\Delta, \tau_0^H) = \hat{L}_H(0, \tau_0^H)$:

$$W_H(\Delta, \tau_0^H) = \hat{L}_H(0, \tau_0^H) \iff \lambda |\tau_0^H| = \frac{v_H - \Delta}{v_L - \Delta}.$$

Given this, define the value function at a price of 0 as $L_H(0, t) = \hat{L}_H(0, t \vee \tau_0^H)$. To define

$\tau_\emptyset^H$, first define the following value function:

$$\hat{L}_H(\emptyset, t) = \int_t^0 \lambda e^{-2\lambda(s-t)} \left(W_H(0,s) + L_H(0,s)\right) ds.$$

Let $\tau_\emptyset^H$ be the following:

$$\tau_\emptyset^H := \max\{t : \hat{L}_H(\emptyset, t) = W_H(0,t)\}.$$

Note that $\{t : \hat{L}_H(\emptyset, t) = W_H(0,t)\}$ is non-empty because $\hat{L}_H(0,t)$ and $W_H(0,t)$ are continuous in t, $\hat{L}_H(\emptyset, 0) = 0 < W_H(0,0) = v_H$, and

$$\lim_{t\to-\infty} \hat{L}_H(\emptyset, t) = \frac{1}{2}(v_H - v_L) + \frac{1}{2}L_H(0, \tau_0^H) > v_H - v_L = \lim_{t\to-\infty} W_H(0,t).$$

Moreover, because of the continuity of $\hat{L}_H(\emptyset, \cdot)$ and $W_H(0,\cdot)$ and the fact that $\hat{L}_H(\emptyset, 0) < W_H(0,0)$, $\tau_\emptyset^H < 0$ and $W_H(0,t) > \hat{L}_H(\emptyset, t)$ for all $t \in (\tau_\emptyset^H, 0]$. Then we can define the value function for the high valuation bidder at $p = \emptyset$ as $L_H(\emptyset, t) = \hat{L}_H(\emptyset, t \vee \tau_\emptyset^H)$.

Finally we show that $\tau_\emptyset^H < \tau_0^H$. To see this, note that for all $t \geq \tau_0^H$,

$$\begin{aligned}
\hat{L}_H(\emptyset, t) &= \int_t^0 \lambda e^{-2\lambda(s-t)} \left(W_H(0,s) + L_H(0,s)\right) ds \\
&= \int_t^0 \lambda e^{-2\lambda(s-t)} \big(2(v_H - v_L)(1 - e^{\lambda s}) + e^{\lambda s} v_H - \lambda s e^{\lambda s}(v_L - \Delta)\big) ds \\
&= \int_t^0 \lambda e^{-2\lambda(s-t)} \big(2(v_H - v_L) + e^{\lambda s}(2v_L - v_H) - \lambda s e^{\lambda s}(v_L - \Delta)\big) ds \\
&= (v_H - v_L)\left(1 - e^{2\lambda t}\right) + e^{\lambda t}\left(1 - e^{\lambda t}\right)(2v_L - v_H) - \int_t^0 \lambda e^{-2\lambda(s-t)} (\lambda s) e^{\lambda s}(v_L - \Delta) ds \\
&= (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t}(1 - e^{\lambda t}) v_L - \lambda e^{2\lambda t}(v_L - \Delta) \int_t^0 \lambda s e^{-\lambda s} ds \\
&= (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t}(1 - e^{\lambda t}) v_L + \lambda |t| e^{\lambda t}(v_L - \Delta) - e^{\lambda t}\left(1 - e^{\lambda t}\right)(v_L - \Delta) \\
&= (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t}(1 - e^{\lambda t})\Delta + \lambda |t| e^{\lambda t}(v_L - \Delta) \\
&\leq (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t}\Delta + e^{\lambda t}(v_H - \Delta) = (v_H - v_L)(1 - e^{\lambda t}) + e^{\lambda t} v_H.
\end{aligned}$$

Thus, $\tau_\emptyset^H \leq \tau_0^H$.

Having defined the cutoffs τ_0^H and $\tau_\emptyset^H$, we now check that indeed the specified strategy profile with these cutoffs indeed constitute an equilibrium. First consider the incentives of the high valuation bidder.

Case 1: Suppose that $p = \emptyset$ and that $t < \tau_\emptyset^H$. By construction of the value functions above, following the prescribed strategy yields a payoff of $L_H(\emptyset, t) = L_H(\emptyset, \tau_\emptyset^H)$. On the other hand, placing a bid today induces the opponent to place a bid of v_L upon first arrival yielding a payoff of at most $v_H e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right)$. Note that by construction of $\tau_\emptyset^H$,

$$v_H e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right) < v_H e^{\lambda \tau_\emptyset^H} + (v_H - v_L)\left(1 - e^{\lambda \tau_\emptyset^H}\right) = L_H(\emptyset, \tau_\emptyset^H) = L_H(\emptyset, t).$$

Therefore abstaining is optimal.

Case 2: Suppose that $p = \emptyset$ and that $t \geq \tau_\emptyset^H$. By following the prescribed strategy, the high valuation bidder obtains a payoff of $W_H(0, t) = v_H e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right)$. Any other bid will result in a payoff of at most $v_H e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right)$ since it induces the opponent to bid v_L at his first bidding opportunity. Finally a one-stage deviation to abstaining today results in a payoff of $L_H(\emptyset, t)$. This is less than $W_H(0, t)$ since $t \geq \tau_\emptyset^H$. As a result, placing a bid of v_H is optimal.

Case 3: Suppose that the opponent is the only player to have placed a bid and $t < \tau_0^H$. At such a history, the high valuation bidder is the losing bidder at a price of $p = 0$. Placing any bid induces the opponent to bid v_L at his first bidding opportunity. Thus placing a bid yields at most a payoff of $(v_H - \Delta)e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right)$. On the other hand following the prescribed strategy yields a payoff of $L_H(0, t) = L_H(0, \tau_0^H)$. Note that

$$\begin{aligned} L_H(0, t) = L_H(0, \tau_0^H) &= (v_H - \Delta)e^{\lambda \tau_0^H} + (v_H - v_L)\left(1 - e^{\lambda \tau_0^H}\right) \\ &> (v_H - \Delta)e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right). \end{aligned}$$

Thus, the prescribed strategy is optimal.

Case 4: Suppose that the opponent is the only player to have placed a bid and $t \geq \tau_0^H$. At such a history, the high valuation bidder is the losing bidder at a price of $p = 0$. Note that at such a history, the high valuation bidder has a belief that the current highest bid is Δ. Thus playing the prescribed strategy yields a payoff of $W(\Delta, t) = (v_H - \Delta)e^{\lambda t} + (v_H - v_L)\left(1 - e^{\lambda t}\right)$. Placing any other bid yields at most a payoff of $W(0, t)$ and abstaining yields a payoff of $L(0, t)$ which by construction of τ_0^H is weakly less than $W(\Delta, t)$. Again the prescribed strategy is optimal.

Case 5: Finally suppose that the high valuation bidder had placed at least one bid. If v_H

is not a feasible bid, then incentives are trivial. Thus suppose that v_H is feasible. Note that at such a history, the opponent bids v_L upon the first opportunity if such a bid is indeed feasible. Against such a strategy, it is immediate that bidding v_H today is optimal.

Now consider the incentives of the low valuation bidder.

Case 1a: Suppose that $p = \emptyset$ and that $t \geq \tau_{\emptyset}^H \vee \tau_0^H$. In this case, clearly placing any bid strictly above v_L is suboptimal, placing any bid weakly less than v_L yields a payoff of $v_L e^{\lambda t}$, while abstaining yields a payoff of $v_L \left(1 - e^{\lambda t}\right) e^{\lambda t} < v_L e^{\lambda t}$. Thus bidding Δ is optimal.

Case 1b: Suppose that $p = \emptyset$ and that $t \in [\tau_{\emptyset}^H, \tau_0^H)$. Clearly any bid above v_L is suboptimal while any bid weakly less than v_L yields a payoff of $v_L e^{\lambda \tau_0^H}$. On the other hand, abstaining today yields

$$L_L(\emptyset, t) := \left(1 - e^{-2\lambda(\tau_0^H - t)}\right) \frac{1}{2} v_L e^{\lambda \tau_0^H} + e^{-2\lambda(\tau_0^H - t)} v_L \left(1 - e^{\lambda \tau_0^H}\right) e^{\lambda \tau_0^H} \leq v_L e^{\lambda \tau_0^H}.$$

So bidding Δ is optimal.

Case 1c: Suppose that $p = \emptyset$ and that $t < \tau_{\emptyset}^H \leq \tau_0^H$. Again placing any bid strictly above v_L is suboptimal while placing any bid weakly less than v_L yields a payoff of $v_L e^{\lambda \tau_0^H}$. Abstaining today yields a payoff of

$$\left(1 - e^{-\lambda(\tau_{\emptyset}^H - t)}\right) v_L e^{\lambda \tau_0^H} + e^{-\lambda(\tau_{\emptyset}^H - t)} L_L(\emptyset, \tau_{\emptyset}^H) \leq v_L e^{\lambda \tau_0^H}.$$

Again bidding Δ is optimal.

Case 2: Suppose that the other player has placed at least one bid. If v_L is not feasible, then incentives are trivial. Suppose otherwise that v_L is feasible. Then in the continuation game, the opponent places a bid of v_H upon her first bidding opportunity whenever feasible. Against such a strategy, it is immediate that it is optimal to place a bid of v_L.

Case 3: The only remaining case is when the other player has not placed any bids and the low valuation bidder is the current winning bidder. In such a case, regardless of whether he abstains or places a bid today, he obtains a payoff of $v_L e^{\lambda(t \vee \tau_0^H)}$. Thus, in particular, it is optimal to abstain. ∎

E.2.2 Asymmetric Information

Thus far we have considered auctions in which each bidder's value is common knowledge. We now consider the case when valuations are privately known. For simplicity, we restrict attention to two bidders with identically and independently drawn valuations with binary support. In this environment it is possible to construct equilibria in which a bidder can only make inferences on the other bidder's type once it is no longer relevant to her bidding decisions. This greatly simplifies the calculation of cutoff points and the verification of incentive constraints. However, we conjecture that the games at hand have many more complicated incremental equilibria in which bidders draw nontrivial inferences on each others' types along the equilibrium path.

Proposition E.3. *Assume that there are 2 bidders, whose valuations are drawn iid, taking value $v_L > \Delta$ with probability $q \in (0,1)$ and $v_H > v_L$ with probability $1-q$. Then there exists an equilibrium in which bidders with valuation v_H bid gradually, and both types of bidders abstain from bidding at certain histories along the equilibrium path.*

The equilibrium we construct in the proof of Proposition E.3 takes the following form. Low types play the following strategy:

1. Bid v_L whenever $p = \emptyset$;
2. If the opponent is the only one to place a bid, then place a bid of v_L if $t \geq \bar{t} = -\frac{1}{\lambda(1-q)}$;
3. If the bidder himself has placed a bid and v_L is feasible, place a bid of v_L.
4. otherwise abstain.

High types play a similar strategy:

1. If $p = \emptyset$, bid v_L;
2. If the opponent is the only one to place a bid, then place a bid of v_H if $t \geq \bar{t} = -\frac{1}{\lambda(1-q)}$;
3. If the opponent hasn't placed a bid, but she herself has, then place a bid of v_L if feasible;
4. If both players have placed a bid, bid v_H whenever feasible;
5. otherwise abstain.

Note that under this strategy profile, the high type plays a (partially) gradual bidding strategy in which he initially bids v_L which is strictly less than his valuation for the object. Moreover the cutoff time satisfies $\bar{t} = -\frac{1}{\lambda(1-q)}$. To understand the role of the threshold $\bar{t}$, consider the high valuation bidder's incentives. Note that the benefit to placing a bid of v_H at some time t is that the high valuation bidder is able to guarantee himself a payoff of $v_H - v_L$ in case the opponent is a low valuation type. However, the cost is that in case the opponent is a high valuation type, he increases the chance that the other opponent arrives again before the end of the auction to increase the final price to v_H instead of v_L. Note that the cost overwhelms the benefit at times t far away from the end of the auction, leading the high valuation bidder to abstain at such times. Furthermore, when q is large the cost is less severe leading to a threshold time $\bar{t}$ further away from the end of the auction.

Note that because of the threshold time $\bar{t}$, the expected payoff as $T \to -\infty$ for a low valuation type is given by: $\frac{1}{2} v_L e^{\lambda \bar{t}}$, while the high valuation type's expected payoff as $T \to -\infty$ is:

$$\frac{1}{2}\left(v_H e^{\lambda \bar{t}} + \left(1 - e^{\lambda \bar{t}}\right) q(v_H - v_L)\right) + \frac{1}{2}\left(q\left(1 - e^{\lambda \bar{t}}\right) + (1-q)\lambda |t^*| e^{\lambda \bar{t}}\right)(v_H - v_L).$$

Proof of Proposition E.3. Consider the strategy profile described above with $\bar{t} = -1/\lambda$. We now show that the strategy profile is indeed an equilibrium. Consider first the incentives of a low valuation type.

Case 1: Suppose that $p = \emptyset$. Clearly placing any bid $b > v_L$ is suboptimal. Placing any bid $b = \Delta, \ldots, v_L$ yields a payoff of $v_L e^{\lambda(t \vee \bar{t})}$. On the other hand, abstaining today yields a payoff of

$$\begin{cases} \left(1 - e^{-2\lambda(\bar{t}-t)}\right) \frac{1}{2} v_L e^{\lambda \bar{t}} + e^{-2\lambda(\bar{t}-t)} v_L \left(1 - e^{\lambda \bar{t}}\right) e^{\lambda \bar{t}} & \text{if } t < \bar{t} \\ \left(1 - e^{\lambda t}\right) v_L e^{\lambda t} & \text{if } t \geq \bar{t}. \end{cases}$$

Clearly the latter is less than $v_L e^{\lambda(t \vee \bar{t})}$ and so bidding v_L is optimal.

Case 2: Suppose that the opponent is the only one to place a bid. At such a history, this bidder believes that the opponent has placed a bid of v_L. Thus the best payoff he can hope to obtain is 0. As a result, he is indifferent between abstaining and placing any bid weakly less than v_L. As a result, it is optimal for the bidder to abstain for all $t < \bar{t}$ and to place a bid of v_L at $t \geq \bar{t}$.

Case 3: Suppose that the opponent has placed no bids and he is the current winner. If the current high bid $B \geq v_L$, then clearly it is suboptimal to place any bid $b' \geq v_L + \Delta$. As a result, it is optimal to abstain. If instead the current high bid $B < v_L$, then again it is suboptimal to place any bid $b' \geq v_L + \Delta$. By abstaining or placing any bid $b' \leq v_L$, the bidder obtains the same payoff of $v_L e^{\lambda(t \vee \bar{t})}$. Thus it is optimal to place a bid of v_L.

Case 4: Suppose that both players have already placed a bid. Then in the continuation game, the opponent bids truthfully upon the first opportunity. Therefore, it is optimal to bid v_L if it is feasible and otherwise abstain.

We now study the incentives of the high valuation bidder.

Case 1: If $p = \emptyset$, note that the posterior belief that the opponent is a low valuation type is q. By placing any bid weakly above v_L, the bidder obtains the same payoff of:

$$W(0,t) := e^{\lambda(t \vee \bar{t})} v_H + \left(1 - e^{\lambda(t \vee \bar{t})}\right) q(v_H - v_L).$$

Similarly, let $L(0,t)$ be the continuation value to a low valuation bidder at a history at time t when the opponent is the only one that has placed a bid. Then, abstaining yields a payoff of:

$$L(\emptyset,t) := \int_t^0 \lambda e^{-2\lambda(s-t)} \left(W(0,s) + L(0,s)\right) ds.$$

For $t \geq \bar{t}$, note that

$$\begin{aligned}
L(\emptyset,t) &= \int_t^0 \lambda e^{-2\lambda(s-t)} \left(2q(v_H - v_L)\left(1 - e^{\lambda s}\right) + v_H e^{\lambda s} + (1-q)(v_H - v_L)\lambda |s| e^{\lambda s}\right) ds \\
&= \left(1 - e^{2\lambda t}\right) q(v_H - v_L) + \int_t^0 \lambda e^{-2\lambda(s-t)} \left((v_H - 2q(v_H - v_L))e^{\lambda s} + (1-q)(v_H - v_L)\lambda |s| e^{\lambda s}\right) ds \\
&= \left(1 - e^{\lambda t}\right) q(v_H - v_L) + e^{\lambda t}\left(1 - e^{\lambda t}\right) v_L + \lambda |t| e^{\lambda t}(1-q)(v_H - v_L) \\
&\leq \left(1 - e^{\lambda t}\right) q(v_H - v_L) + e^{\lambda t} v_L + e^{\lambda t}(v_H - v_L) = W(0,t).
\end{aligned}$$

Furthermore for all $t < \bar{t}$,

$$L(\emptyset,t) = \left(1 - e^{-2\lambda(\bar{t}-t)}\right)\left(\frac{1}{2}W(0,\bar{t}) + \frac{1}{2}L(0,\bar{t})\right) + e^{-2\lambda(\bar{t}-t)} L(\emptyset,\bar{t}).$$

Note that $\bar{t}$ is constructed so that $L(0,\bar{t}) = q(v_H - v_L) + (1-q)(v_H - v_L)e^{\lambda \bar{t}}$ and therefore, $L(0,\bar{t}) \leq W(0,\bar{t})$. As a result,

$$L(\emptyset,t) \leq \left(1 - e^{-2\lambda(\bar{t}-t)}\right) W(0,\bar{t}) + e^{-2\lambda(\bar{t}-t)} L(\emptyset,\bar{t}) \leq W(0,\bar{t}).$$

Thus, abstaining is suboptimal at such a history at all times t.

Now consider any the strategy of placing of any bid $b' < v_L$ at time t. Note that if the opponent does arrive in the time interval $[t \vee \bar{t}, 0)$, the most that this bidder can obtain is $(v_H - v_L)$ if the opponent is a low type and 0 if the opponent is a high type. Therefore, such a bid gives this bidder a payoff of at most $e^{\lambda(t \vee \bar{t})} v_H + \left(1 - e^{\lambda(t\vee\bar{t})}\right) q(v_H - v_L) = W(0,t)$. Hence, we have shown that bidding v_L is optimal.

Case 2a: Suppose that the opponent is the only one to place a bid and $t \geq \bar{t}$. At such a history, the bidder believes that the current high bid is v_L. Furthermore, this bidder's posterior belief about the opponent is still the prior q that he is a low type. This is because at both histories where either $p = \emptyset$ or $p = 0$ and the opponent is the current winner, the high and low types behave the same way. Thus by placing any bid strictly above v_L, this bidder obtains a payoff of

$$q(v_H - v_L) + (1-q)(v_H - v_L)e^{\lambda t}.$$

Placing a bid weakly below v_L yields a payoff of

$$q(v_H - v_L)\left(1 - e^{\lambda t}\right) + (1-q)(v_H - v_L)\left(1 - e^{\lambda t}\right)e^{\lambda t}$$

which is clearly less than the former. Abstaining on the other hand yields a payoff of

$$q(v_H - v_L)\left(1 - e^{\lambda t}\right) + (1-q)(v_H - v_L)\lambda|t|e^{\lambda t}$$

which is less than the payoff to bidding v_H if $t \geq \bar{t}$. Thus placing a bid of v_H is optimal.
Case 2b: Suppose that the opponent is the only one to place a bid and $t < \bar{t}$. The arguments here are the same as in Case 2a except the payoff to abstaining is clearly greater than the payoff to placing a bid strictly above v_L since $t < \bar{t}$. Thus abstaining is optimal.
Case 3: Suppose the opponent hasn't yet placed a bid, but she is the current winning bidder. The arguments here are exactly the same as the argument used in Case 1 to show suboptimality of placing a bid $b' < v_L$. Thus, placing a bid of v_L is optimal at such a history.
Case 4: If both players have already placed a bid, then in the continuation game, the opponent bids truthfully. Therefore, it is optimal to bid truthfully as well. ∎