Internet Auctions with Artificial Adaptive Agents: A Study on Market Design*

John Duffy[†]
University of Pittsburgh

M. Utku Ünver[‡]
University of Pittsburgh

February 10, 2007

Abstract

Many internet auction sites implement ascending-bid, second-price auctions. Empirically, last-minute or "late" bidding is frequently observed in "hard-close" but not in "soft-close" versions of these auctions. In this paper, we introduce an independent private-value repeated internet auction model to explain this observed difference in bidding behavior. We use finite automata to model the repeated auction strategies. We report results from simulations involving populations of artificial bidders who update their strategies via a genetic algorithm. We show that our model can deliver late or early bidding behavior, depending on the auction closing rule in accordance with the empirical evidence. Among other findings, we observe that hard-close auctions raise less revenue than soft-close auctions. We also investigate interesting properties of the evolving strategies and arrive at some conclusions regarding both auction designs from a market design point of view.

Keywords: auctions, artificial agent simulations, genetic algorithm, finite automata
J.E.L. Classification Numbers: D44, D83, C63, C99

*We would like to thank two anonymous referees, Alexander Elbittar and Jeff MacKie-Mason for useful comments on earlier versions of this paper. Ünver gratefully acknowledges the support of the Turkish Academy of Sciences Distinguished Young Scholars Support Program (TÜBA-GEBİP), and both Duffy and Ünver thank the National Science Foundation for their support.

[†]University of Pittsburgh, Department of Economics, Pittsburgh, PA 15260, USA. jduffy@pitt.edu
[‡]University of Pittsburgh, Department of Economics, Pittsburgh, PA 15260, USA. uunver@pitt.edu

1 Introduction

Since the advent of the world wide web, an increasing number of goods are being traded using on–line, "internet auctions".[1] The most popular internet auction sites are those run by eBay, Amazon, and Yahoo! While these sites implement several different auction formats, the most common and widely used format is an ascending-bid, second-price format that is a hybrid of the ascending-bid English auction and the second-price sealed bid auction. Indeed, this internet auction format has led to a new theoretical and empirical literature.[2]

There are two popular ascending-bid, second-price formats used by internet auction websites. The first type, used e.g., by eBay, is a "hard-close" auction where bidding closes at the end of a known preset time period, typically seven days. The high bidder wins the object by paying the second highest bid plus some small increment. The second type, used e.g., by Amazon, is a "soft-close" auction where bidding closes at the end of a known, preset time period if and only if no bidder submits a "late" bid within a certain interval of time near the scheduled closing time (e.g., within the last 10 minutes). Otherwise, the auction is extended for a fixed and known additional period of time (e.g., 10 more minutes), starting from the time of submission of the last bid.

An interesting phenomenon, known to participants in *hard*-close internet auction sites and empirically documented by Roth and Ockenfels (2002) and Bajari and Hortaçsu (2003) is that of "last–minute" or "late" bidding, which practitioners call "sniping." Specifically, many more bids are submitted very close to, or just at the end of a hard-close auction than are submitted near the scheduled end of soft-close auctions. Further, the number of bids per bidder is higher in hard-close than in soft-close auctions. The evidence for late bidding in hard-close auctions is puzzling since successful internet auction bidders pay only the second-highest price and we know from Vickrey (1961) that in (static) sealed-bid second price auctions it is a weakly dominant strategy for bidders to simply bid their reservation values. Indeed, eBay even advises bidders in its hard-close auctions to "enter the maximum amount you are willing to pay for the item."

Ockenfels and Roth (2006) present a model that can rationalize late bidding as an equilibrium strategy in hard-close auctions under the assumption of either private- or common-values.[3] Their model, based on tacit collusion among bidders, relies heavily on the assumption that there is a small probability that bids submitted in the final moments of a hard-close auction may fail to be properly transmitted to the auction software. This end-of-auction congestion creates a potentially large ex-ante surplus for bidders. The potential to collude so as to capture this surplus is what rationalizes the use of late-bidding strategies.[4]

[1] For instance, eBay Inc. reports steady growth in its annual gross merchandize volume – the total value of all successfully closed listings on eBay's trading platforms (primarily auction listings) – from $95 million in 1997 to $44.3 billion in 2005.

[2] See Bajari and Hortaçsu (2004) for a survey.

[3] The type of item sold determines whether the auction is a common-value or a private-value auction. Antique coin auctions are examples of common-value auctions and computer part auctions are examples of private-value auctions. In this paper, we focus on private-value auctions only.

[4] Ockenfels and Roth (2006) also note that late bidding can be caused by the presence of naive incremental bidders. As

However, late bidding is just one equilibrium possibility in their model; all bidders bidding early remains another. Ariely, Ockenfels and Roth (2005) report results from a laboratory human subject experiment which confirms that hard-close auctions are prone to late bidding while soft-close auctions are not.

While this theoretical and experimental work rationalizes or confirms the late (early) bidding phenomenon in hard (soft) close internet auctions, it does not tell us about the *types* of bidding strategies that are being used in the different internet auction formats. In both the field and the laboratory we do not observe the strategies that individuals play in internet auctions – only the players' actions (bids). Further, it is difficult to assess the sensitivity of bidding strategies to changes in auction rules, environmental parameters or other design features in the field (where such features are fixed) or in the laboratory where such changes are costly in the sense that additional sessions with paid human subjects have to be conducted for each change in auction rules.

Our research aim in this paper is to develop a model of dynamic, internet auctions that delivers both the macro-phenomena of late or early bidding consistent with the auction having a hard or soft close, but which also enables us to assess at the micro-level, the bidding strategies that are being employed by individual agents. Because of the complexity of the strategy space, we adopt the agent-based methodology. Agent-based modeling represents a different approach to empirical understanding wherein the researcher shows –typically using computer simulations – how a population of autonomous, heterogeneous and boundedly rational adaptive agents interacting with one another in some well-defined environment can generate certain macroscopic empirical phenomenon.[5] In our case, the macroscopic phenomenon will be early or late bidding. An important by-product of generating macroscopic phenomena is that the modeler can investigate the individual behavior that supports that phenomena, in our case, the individual bidding strategies.

Our agent-based model involves a single seller offering an item for which bidders have independent, private-values. The bidders, who are the only active players in our model, play strategies against one another repeatedly in either hard- or soft-close multi-period auctions. Bidding strategies evolve via a genetic algorithm (Holland, 1975), which is a versatile search and optimization tool ideally suited to the large strategy spaces that are possible in dynamic, internet auctions. While the bidding strategies we consider are quite simple, they are flexible enough to distinguish between early or late bidding as well as between history contingent or unconditional-bidding behavior. We use simulations of our agent-based model to address several research questions regarding behavior in internet auctions.

First, holding the format of agent strategies and the method by which these strategies evolve constant,

evidence of this, they show instances of the multiple submission of bids by the same bidder. On the other hand, Bajari and Hortaçsu (2003) explain multiplicity of bids by on-going updating related to the common-value aspect of many items being auctioned. Empirically, Gonzalez, Hasker and Sickles (2004) show that the type of late bidding equilibrium introduced by Ockenfels and Roth (2006) and a type of early bidding equilibrium (a variant of the equilibria introduced by Avery (1998) for English auctions) are not actually played by bidders.

[5] See, e.g., Tesfatsion and Judd (2006) or Epstein (2006) for introductions to agent-based modeling in economic systems.

is a change in the auction closing rule by itself sufficient to generate a difference in the frequency of late bidding and if so, is this in the direction consistent with the empirical evidence? Indeed, we consistently find that hard-close auctions, together with a small failure rate for last minute bids, lead to much more frequent use of late-bidding strategies than do soft-close auctions, even as we vary certain parameters for both auction formats such as the number of bidders. We perform an extensive sensitivity analysis of the robustness of this finding and our other findings to changes in model parameters governing admissible bids, failure rates for last minute bids, the number of auctions played, bidding behavior in extension periods of soft-close auctions, and the distribution and discreteness of private valuations. We note that the evidence our model provides on bidding behavior between auction format is more direct than can be ascertained using *field data*, where comparisons between auctions (eBay, Amazon) may be complicated by factors other than the closing rule (e.g., the greater trading volume on eBay) or inferences from *laboratory experiments* where control over factors such as subjects' risk attitudes or prior experience with internet auctions is imperfect. Further, theoretical models rationalizing late bidding (e.g., Ockenfels and Roth, 2006) as Bayesian Nash equilibria admit many other equilibria as well, e.g., all early bidding, with no refinement for selecting from among these equilibria. The adaptive search algorithm we use can be viewed as an equilibrium selection device.

Second, we consider the implications of the different closing rules for seller and buyer surpluses. Again, this is a difficult question to address using field data as the same items are not typically auctioned using different internet auction formats.[6] We find that sellers are relatively worse off in hard-close auctions than in soft-close auctions in that their average revenue in hard-close auctions is lower. Not surprisingly, the reverse finding holds for buyers who do better (receive, on average, a larger surplus) in hard-close auctions than in soft-close auctions.

Third, we ask what type of bidder strategies are being used in the two different internet auction formats and whether differences in these strategies help account for the different bidding behavior between hard- and soft-close auctions. We find that in hard-close auctions with small numbers of bidders, the evolved bidding strategies condition on the history of rival bids and bidders typically wait until the last period to bid their full private valuation for the item. In soft-close auctions, the vast majority of strategies involve unconditional early bidding of the bidder's full valuation.

Finally, we consider the evolution of adaptive bidding strategies in the presence of naive incremental bidding agents who repeatedly increment the amount of their bid until they are the auction's current high bidder. The existence of such incremental bidders in internet auctions is well documented both in the field and in an experiment by Ariely, Ockenfels and Roth (2005). The addition of incremental bidders serves only to widen the difference we observe in the frequency of late-bidding behavior between hard-close (where it remains high) and soft-close auctions (where it is infrequent) relative to the difference we find in simulations without these incremental bidders. Moreover, with incremental bidders, a large difference in

[6] See, however, Lucking-Reiley (1999) for a field-experiment approach.

the frequency of late bidding between hard- and soft- close auctions is obtained without the need for any congestion in last minute bidding (failure rates of zero).

We conclude that our agent-based model appears to capture some of the empirical regularities regarding bidding behavior in internet auctions, giving us some confidence in the use of this or related models as tools for market designers interested in predicting outcomes under a variety of different internet auction formats.

2 A Model of Internet Auctions

We present a simplified model of an internet auction that captures essential features of these auctions, though the correspondence between our model and internet auctions formats (eBay, Amazon) is not precise. In particular, we suppose there is a single seller, offering an indivisible object without any reserve price and there are n risk neutral bidders. As the bidders are the only active players in our environment, we refer to them as the "agents." In our simulation analysis we will vary n from 2 to 5. Each agent has a private valuation for the auction item that is an independently and identically distributed draw from a discrete probability density function g, that is the same for all agents. Agents' valuations are all drawn at the beginning of each internet auction and remain fixed for the duration of that auction. The density g is the discrete uniform density, which has n_V equally distant mass points in the interval $[m - \frac{\epsilon}{2}, m + \frac{\epsilon}{2}]$, where $m, \epsilon \in \mathbb{R}_{++}$ are such that $m \geq \frac{\epsilon}{2}$. Specifically:

$$g(v) = \begin{cases} \frac{1}{n_V} & \text{if } v = m - \frac{\epsilon}{2} + \frac{k\epsilon}{n_V - 1} \text{ for some } k \in \{0, 1, ..., n_V - 1\}, \\ 0 & \text{otherwise.} \end{cases}$$

An internet auction is a dynamic auction repeated over T consecutive bidding periods, indexed by $t = 1, 2, ..., T$. In our benchmark simulations we set $T = 8$. In each period t, each of the n bidders can submit at most a single bid, which may or may not comprise a *valid* bid.[7] A valid internet auction bid must be greater than what eBay and Amazon call the "current bid," which is the second highest amount bid in the auction to date - call this amount b_2; if the current high bidder were to win the auction, b_2 is the price that bidder would pay. In our design a new valid bid must exceed the current bid b_2 by the bid increment Δ. Bids below $b_2 + \Delta$ are declared invalid and do not affect any auction values. When there are no bids submitted, the smallest valid bid is set equal to the bid increment Δ. When the first bid greater than Δ arrives, the current bid $b_2 = \Delta$. The owner of the first bid becomes the auction's "high bidder." The high bidder's bid in all internet auctions (eBay, Amazon) is not a static value, but adjusts according to a known, *proxy bidding rule* that automatically increments the high bidders's bid as new valid bids come in to challenge that bid.

[7]Bidders can choose whether or not to bid in any given period as discussed in detail below.

More precisely, suppose b^* is the highest bid submitted to date, bid b_2 is the second highest "current bid" and bidder i^* is the identity of the current high bidder. All bidders observe only the second highest bid, b_2, and the identity of the current high bidder, i^*. All bidders who choose to submit a bid in each period do so simultaneously. However, in our set-up, the arrival order of these bids to the computer program (internet website) is not simultaneous, and is determined by a random permutation of the n bidders, with each possible permutation being equally likely.

The proxy bidding rule of internet auctions works as follows. Suppose a new valid bid b is submitted by bidder i. If $b > b^*$ then b becomes the new high bid, bidder i is identified as the current high bidder and b_2 is set equal to the minimum of $\{b^* + \Delta, \quad b\}$. On the other hand, if $b \leq b^*$ then the high bid b^* remains unchanged as does the identity of the high bidder i^*, and b_2 is set equal to the minimum of $\{b + \Delta, \quad b^*\}$. After the first new valid bid b is processed, subsequent valid bids are processed in a similar fashion using the updated information concerning the high bid and the second, "current bid". Processing of bids involves checking first whether a bid is valid or not and then processing only the valid bids in the manner described above. This processing continues until all n bidders' bids have been considered at which point the period t ends and we move on to period $t + 1$, subject to the ending rule in operation.

We suppose that any valid bid submitted in one of the first $T - 1$ bidding periods is always correctly registered by the software program. However, a valid bid that is submitted in the final period T is correctly registered by the software program only with some fixed probability $\rho \leq 1$. The latter assumption captures congestion effects or bidder timing mistakes near the scheduled end period of an auction. For our benchmark simulations we set $\rho = .9$. In our sensitivity analysis we will also consider $\rho = .8$ and $\rho = 1$ (no congestion).

As mentioned in the introduction, our main focus is on the two ending procedures for the internet auction. In a "hard-close" auction, the auction closes after the last valid bid of period T (that does not succumb to the congestion effect) has been processed. The high bidder at the end of period T wins the object, paying the second highest current bid, b_2.

In a soft-close auction, the auction closes after period T if and only if a valid bid (that does not succumb to the congestion effect) is not registered in period T. If a valid bid does register in period T, the auction is extended for one more period. All n agents may again simultaneously submit bids and the probability of successful bid submission, ρ, continues to apply in every extension period. If a valid bid does not register in the first extension period the auction closes, but if one does, the auction continues on for a second extension period and so on until no new valid bid registers.[8] The high bidder at the end of a soft-close auction wins the object, paying the second highest current bid, b_2. Note that, as there are a finite number of bidders, and bidding above one's private valuation is prohibited in our model (see bidding strategies

[8]Notice that the probability of successful registration of a bid equals ρ in period T of both the hard and soft-close auctions. Further the probability of successful registration of a bid also equals ρ in every extension period of a soft-close auction, i.e. $T + 1$, $T + 2$,.. etc. This design allows us to maintain comparability between the hard- and soft-close auction formats, in that the *only* difference between the two internet auction formats is in the closing procedure for the auction. Later in the paper we will consider the case of $\rho = 1$ (no congestion).

below), a finite end is guaranteed to the soft-close auction; one agent will have the highest valuation, and bidding cannot exceed that value.

In our simulations, each agent plays a block of R dynamic (T-period) internet auctions. Each dynamic auction in a block is also referred to as a "stage game" or "round", indexed by $r \in \{1, 2, ..., R\}$. The valuation or "type" of each agent is redetermined via another draw from g at the beginning of each new round. Agent i's utility is simply the summation of payoffs earned from all rounds played.

After each auction, all bidders can observe a "bid history" h_r, consisting of a list of the current (second) bids made in all periods $t = 1, 2, ..., T$ of all rounds $r = 1, ..., r-1$ played along with the identities of the high bidder at all times. History h_1 is the empty set. With such a rich history, stage game strategies can be quite complicated objects, not to mention repeated game strategies. For modeling purposes, it will be useful to place some restrictions on the set of admissible internet auction strategies.

2.1 Internet Auction Strategies

We make several assumptions regarding admissible internet auction strategies.

1. Each bidder i's history, h_r^i, is restricted to include only the timing of bids by rival bidders in the last round $r-1$. Further, only two possible histories are permitted. History $h_r^i =$*late* denotes the state where, in round $r-1$, a valid rival bid arrived in period $t \geq T$. History $h_r^i =$*early* denotes the state were, in round $r-1$, no valid rival bid arrived in any period $t \geq T$.

2. A stage game strategy j for agent i conditions only on the previous bids of agent i and his own type (valuation). Further, in each period of a stage game (T-period auction), bidders are permitted to make just one of four possible bids. The bid of bidder i with type v_i in any period $t = 1, 2, ..., T$ can be either: v_i, $\frac{2}{3}v_i$, $\frac{1}{3}v_i$ or 0. Bidding 0 means that bidder i does not submit a bid in that period.[9] Bidder i's stage game strategy j, σ_j^i, in a hard-close internet auction can be compactly written as a three-element vector:

   Stage game strategy for a hard-close auction:

$$\sigma_j^i = \left(\overbrace{\underbrace{\sigma_{j1}^i}_{\text{when to bid } v_i}, \underbrace{\sigma_{j2}^i}_{\text{when to bid } \frac{2}{3}v_i}, \underbrace{\sigma_{j3}^i}_{\text{when to bid } \frac{1}{3}v_i}}^{\text{period numbers}} \right)$$

   where the σ_{jk}^i are integers representing period numbers in $\{0, 1, 2, ..., T\}$. The first number in the string, σ_{j1}^i, is the period when bidder i will bid his entire value. The second number, σ_{j2}^i, is the period

[9] We do not allow bids greater than the bidder's private value simply because bidding above one's private value is always weakly dominated by a bid equal to the private value. The argument for this result is the same as Vickrey's (1961) argument for second price auctions with independent private values.

when bidder i will bid two thirds of his value. The third number, σ_{j3}^i, is the period when bidder i will bid one third of his value. Period 0 means that bidder i will not bid that particular fraction in any period of the auction. Period 1 means the bidder will bid that particular fraction in period 1, etc.[10]

3. In extension periods of a soft-close auction $t > T$, we permit just two possible bids: v_i or 0. (We later consider a variation on this assumption in our sensitivity analysis). Further, bidder i's bid is restricted to be the same in every extension period. Bidder i's j^{th} stage game strategy in a soft-close internet auction is thus written as a four element vector:

Stage game strategy for a soft-close auction:

$$\sigma_j^i = \left(\overbrace{\underbrace{\sigma_{j1}^i}_{\text{when to bid } v_i}, \underbrace{\sigma_{j2}^i}_{\text{when to bid } \frac{2}{3}v_i}, \underbrace{\sigma_{j3}^i}_{\text{when to bid } \frac{1}{3}v_i}, \underbrace{\sigma_{j4}^i}_{\text{whether to bid } v_i \text{ in an extension period}}}^{\text{period numbers}} \right)$$

where σ_{j4}^i alone takes on a binary value $\{0, 1\}$. When $\sigma_{j4}^i = 0$, bidder i will not bid anything in any extension period, otherwise he bids his full valuation v_i.

This simplified model captures many essential features of internet auctions, but some of our simplifications have important consequences. Assumption (1) eliminates the price and identity dimensions from the history of previous auctions and focuses on the timing issues that are the focus of our analysis. Assumption (2) eliminates incremental bidding within an auction (stage game) in response to rival bids by other bidders. While seemingly restrictive, this assumption nevertheless leads to different bidding behavior in hard- and soft-close auctions consistent with the empirical evidence. However, as incremental bidding is thought to play a role in bidding behavior, later in the paper we will exogenously introduce (pre-programmed) naive incremental bidders who submit only incrementally higher bids and we will investigate the evolution of adaptive bidding strategies in the presence of these naive bidders. Assumption (3) is just a simplification that eliminates cumbersome bidding strategies in extension periods.[11] An analysis of more complicated internet bidding strategies that relax some or all of these assumptions may lead to greater insights than can be provided using our model. Nevertheless, as we show below, our model suffices to generate differences in bidding behavior between the two auction formats that is consistent with the empirical evidence.

[10] If a bidder's strategy calls for bidding a larger value, e.g., the full value v_i, in some period t and a smaller value (e.g., $1/3(v_i)$) in some later period, $t + k$, the smaller, later bid is not submitted. Even if we did not restrict the possibility of declining bids by the same bidder, the bid improvement rule of the internet auction insures that any bid smaller than the current bid plus some increment never registers as a valid bid. Hence our restriction simply amounts to reducing the number of invalid bids submitted.

[11] This assumption is not unrealistic as our simulations reveal that in both the hard- and soft-close formats, nearly all bidders learn to bid their full private values.

2.2 Finite Automata Representation of Repeated Auction Strategies

To allow for dynamic, history-contingent repeated game strategies, we endow each bidder i with *two* stage-game strategies, represented by the vectors σ_1^i and σ_2^i.[12] The dynamic, repeated auction strategy of each bidder i, s^i, is characterized using a "finite automata" representation (Moore, 1956).[13] An "automaton" is represented here as a string of integers that describe the two stage game strategies and the transitions between those strategies for each possible history. In our implementation, there are just two histories, "early" and "late," upon which bidders can condition their bidding strategy.

A finite state automaton representation of bidder i's dynamic repeated auction strategy s^i, is the string of integers:

$$s^i = \left(\left(\sigma_1^i\right), \left[L_1^i,\ E_1^i\right];\ \left(\sigma_2^i\right), \left[L_2^i,\ E_2^i\right]\right).$$

where σ_1^i, σ_2^i are the two possible stage game strategies in the format described above (and depending on whether the auction is hard- or soft-close) and L_k^i, $E_k^i \in \{1, 2\}$, are transition indexes associated with each stage game strategy, $k = 1, 2$, that characterize the transitions between the two stage game strategies. For instance, following play of stage game strategy j, if the bid history is "early", then in the next round (auction), bidder i plays stage game strategy $\sigma^i_{E^i_j}$, and if it is late, then in the next round he plays stage game strategy $\sigma^i_{L^i_j}$. In the very first round of the repeated internet auction, we assume that bidder i always selects the first stage game strategy in the finite automaton s^i. Following the first round, the transitions between the two automata are dictated by the bid history and the transition index values of the automaton.

To help clarify the finite automata representation of bidder strategies, we give an example below:
Example 1: Suppose $T = 8$ and that bidder i has the following soft-close strategy:

$$s^i = \left(\left(\mathbf{0}, \mathbf{2}, \mathbf{3}, \mathbf{1}\right), [1,\ 2];\ \left(\mathbf{8}, \mathbf{4}, \mathbf{5}, \mathbf{0}\right), [2,\ 1]\right).$$

Let v_i be bidder i's value. The first stage game strategy tells bidder i to bid $\frac{2}{3}v_i$ in period 2 (as indicated by 2 in the 2$^{\text{nd}}$ digit) and v_i in every extension period (as indicated by the 1 in the 4$^{\text{th}}$ digit). Notice that the strategy also tells the bidder to bid $\frac{1}{3}v_i$ in period 3 (the 3 in the 3$^{\text{rd}}$ digit). However, this part of the strategy is superseded by the part stipulating a bid of $\frac{2}{3}v_i$ in period 2, and so we can ignore the lower bid stipulated for period 3.

If a rival submits a bid in period 8 or in an extension period then the history is late; in that case, the transition index associated with the first stage game strategy calls for bidder i to repeat use of this first

[12] This is the minimal number necessary for implementation of repeated game strategies. Nevertheless, as we show below, our implementation is sufficiently general to characterize a wide range of possible bidding strategies in repeated internet auctions.

[13] Finite automata are frequently used in representing repeated game strategies theoretically, computationally and experimentally. See Abreu and Rubinstein (1988) for the theory of Nash equilibria with finite automata in repeated games. See Miller (1996) for an application with genetic algorithms in repeated prisoner's dilemma. See Engle-Warnick and Slonim (2006) for an application to inference of human strategies from experimental data.

stage game strategy in the next auction (as indicated by the 1 in the 5^{th} digit). Otherwise, the history is "early," and the second stage game strategy is used in the next auction (as indicated by the 2 in the 6^{th} digit).

The second stage game strategy, which starts at the 7^{th} digit, tells bidder i to bid $\frac{2}{3}v_i$ in period 4 (as indicated by the 4 in the 8^{th} digit) and to bid v_i in period 8 (as indicated by the 8 the 7^{th} digit). If the auction is extended, no bid will be submitted (as indicated by the 0 in the10^{th} digit). If a rival bid arrives in or later than period 8, then the history is "late" and the transition index associated with the second stage game strategy indicates that this same stage game strategy will be played again (as indicated by the 2 in the 11^{th} digit), otherwise the history is "early" and the first stage game strategy will be played next (as indicated by the 1 in the 12^{th} digit).

In the next section, we provide some theoretical analysis of the internet auction game.

3 Theoretical Analysis

Here we establish a theoretical result regarding the timing of bidding behavior in a single play of the internet auction using the restricted stage-game strategies of the environment we consider. In particular, we can show that, regardless of whether the internet auction has a hard- or soft- close, it is a weakly dominant strategy for bidders to bid their full valuation (fraction 1 in our implementation) in one of the first $T-1$ periods of the dynamic auction (single-stage game) when the bidding increment is zero. For simplicity, we consider the case with just 2 bidders, though the proof is readily extended to the more general case with more than 2 bidders.

Theorem: Under our modeling choice of strategies, regardless of the ending rule of the internet auction, any strategy that involves bidding fraction 1 (the bidder's full valuation) before period T weakly dominates any other strategy that does not involve bidding fraction 1 early (in one of the first $T-1$ periods) in an internet auction with 2 bidders and with increment $\Delta = 0$.[14]

The dominance of early bidding of valuations by two bidders in a single play of the internet auction game is analogous to the dominance of the "defect" strategy in a single play of a two-person prisoner's dilemma game. We know that infinitely repeated prisoner dilemma games admit many more equilibria than the "always defect" equilibrium of the non-repeated game, and the same is true of the repeated internet auction game. In particular, in the hard-close format, if there is only some probability $\rho < 1$ that a bid registers in the final period, there can be gains to a strategy of late bidding in repeated internet auction games; delaying bids until the end of the auction serves to dampen the final price and raise the expected surplus of bidders adhering to such a collusive strategy. Of course, it is not obvious which repeated game

[14] See the Appendix for the proof.

strategy will be selected in practice, and hence we turn to a simulation analysis. In the next section, we explain how we analyze strategic bidding in a repeated internet auction.

4 Simulations with Artificial Adaptive Agents

We use a model of adaptation known as a "genetic algorithm." The genetic algorithm is a population-based, stochastic directed search algorithm based on principles of natural selection and genetics. These algorithms have powerful search capabilities for large strategy spaces, such as the one considered in this paper. In particular, genetic algorithms have been shown to optimize on the trade-off between exploring new strategies and exploiting strategies that have performed well in the past (see Holland, 1975). Further, these algorithms are frequently capable of finding optimal solutions in highly complex environments. We therefore regard the genetic algorithm as a useful equilibrium selection device. Economic applications of genetic algorithms are discussed and surveyed in Dawid (1999) and Arifovic (2000). The economic application most closely related to this one is Andreoni and Miller's (1995) use of genetic algorithms to find bidding strategies in a variety of different auction formats. However, Andreoni and Miller did not consider dynamic auctions with the bidding rules of internet auctions.

4.1 Algorithmic Details

The programming language we use is Delphi (Object Pascal). We ran the simulations using a Windows Desktop Pentium 4 PC. The code is available on request.

Our program operates on a population of strategies (finite automata) of the type described above. The size of this population is fixed at $N = 30$. The automata in this population are initially generated randomly subject to constraints on integer values, e.g., the digits indicating the periods in which various amounts are bid must lie between 0 and T, the number of periods in an auction. Over time, this population of strategies evolves via the genetic operations of the genetic algorithm as described below. This evolution step occurs only after the N strategies of the population have gained experience playing blocks of $R = 20$ internet auctions. Specifically, the genetic operators of the genetic algorithm are called on after a fixed number of blocks (300) or "tournament" has been played.

Each block proceeds as follows. First, a set of n finite automata (bidders) are randomly chosen from the 30-member population of finite automata. Our simulations were conducted separately for groups of $n = 2, 3, 4$ or 5 bidders. Second, these n bidders play against one another for $R = 20$ consecutive dynamic, internet auctions, each lasting $T = 8$ periods or possibly longer in the case of soft-close auctions. Our simulations are conducted separately for hard- and soft-close auction formats. At the start of each dynamic auction, each strategy draws a random valuation from the pdf g, and plays its strategy against the other n bidders (strategies). The bidder's (strategy's) payoff from an auction is the difference between the bidder's valuation and the price paid for the item, if the bidder (strategy) won the auction; otherwise the payoff is zero. At the end of these R auctions, each strategy is assigned a fitness score. The fitness of each strategy is its average payoff from all R auctions played in the block. Further blocks of auctions are then played in the same manner, always by first drawing n strategies at random and then having these same strategies play one another in R internet auctions.

After a fixed number of blocks has been played (300 in our simulations), average fitness levels are calculated for each strategy adjusting for the number of blocks that strategy participated in and using the average payoff that strategy earned in each block. These fitness scores are used to select strategies for reproduction in the next population, or "generation" of N strategies. These reproduced strategies may also undergo some recombination and mutations before becoming the strategies that make up the next generation of strategies as described below. Generation $G+1$ is called the "offspring" of generation G for every $G \geq 1$.

The genetic algorithm has three basic operators that are used to update the strategies in the population of $N = 30$ strategies.

1. Selection: Some number, $M = 6$, of the very best strategies of the current generation, as determined by fitness levels, are selected for inclusion (copied intact), in the set of N "offspring" strategies that comprise the next generation.[15] The remaining $N - M = 24$ next generation strategies are obtained using the crossover operation described next.

2. Crossover: We adopt the "selection with linear-crossover" procedure described in the following steps.

   (a) Two parent strategies are selected randomly from the population of N strings (with replacement) in proportion to their relative fitness. These two parent strategies are strings of real integers of length L; $L = 12$ in a soft-close auction strategy, and $L = 10$ in a hard-close auction strategy.

   (b) An arbitrary crossover point $\ell \in \{1, ..., L-1\}$ is randomly determined for this pair of parent strategies.

   (c) The first $\ell < L$ integers of the first parent strategy and the last $L - \ell$ integers of the second parent strategy are combined to form the first new offspring strategy. Similarly, the first $\ell < L$ integers of the second parent strategy and the last $L - \ell$ integers of the first parent strategy are combined to form the second new offspring strategy.

   (d) This crossover operation is repeated $(N - M)/2 = 12$ times so as to complete the set of $N = 30$ new offspring strategies for the next generation – the 6 strategies obtained via selection and the 24 strategies obtained via crossover.

3. Mutation: The mutation operation applies to all of the 30 *new* offspring strategies created via selection and crossover. Specifically, each integer of each offspring strategy is randomly changed

[15] This procedure is sometimes termed "elitist selection."

to another admissible integer value with a small fixed probability, .01 and unchanged otherwise. For non-binary values, e.g., the periods in the stage game strategy, we use a non-uniform mutation operation wherein mutations to digits closest to the current digit are more likely than are mutations to digits further away from the current digit.

We apply these operations repeatedly to each generation following the end of each tournament, using average fitness levels over all auctions played so as to create the next generation of strategies. We run each simulation for 4000 generations. Further, for each treatment condition, (number of bidders n; hard- or soft- close auction format) we run $S = 20$ simulations with different random seed values to obtain Monte Carlo estimates of different statistics.

As noted above, we set the number of bidding periods in a stage auction at $T = 8$. We consider $R = 20$ repeated auctions in each of the 300 tournament blocks run among the strategies of a single generation. We conduct benchmark simulations with either $n = 2, 3, 4$ and 5 bidders, and under either a hard- or soft-close format.[16] In these simulations, we set the expected value of the probability density function g used to draw valuations, at $m = 10^6$, and we set the spread of the support interval, $\epsilon = 1000$. We choose the number of mass points in the interval $[10^6 - 500, 10^6 + 500]$ as $n_V = 101$.[17] We set $\rho = 0.9$ as the probability of successful bid registration in each period $t \geq 8$. Finally, we set the bid increment, $\Delta = 1$. We also conducted and report on an extensive sensitivity analysis where we change many of these parameter values one at a time, and then compare new Monte Carlo estimates (again from $S = 20$ simulations) with the original ones.

In the next subsection, we introduce a method for classifying strategies that aids in our presentation of the simulation findings.

4.2 Classification of Strategies

We introduce a simple method for classifying the evolving strategies in our simulation exercises. Each distinct classification is called a "phenotype."[18] A phenotype is determined according to the following two criteria: (i) the period in which the bidder makes his final bid (any fraction of his valuation) for each strategy and (ii) the strategy he plays following each history. For the first criterion, we only take into account bidding in the first T periods. A stage game strategy is of type "E" if the bidder completes his bidding in one of the first T - 1 periods, and is of type "L" otherwise, i.e., the strategy calls for a bid (of any fraction) to be placed in period T. (The classification of "Early" or "Late" ignores strategic behavior in extension periods of soft-close auctions). The phenotype of a repeated game strategy is a characterization of the bidding behavior of each strategy (automata) and its transition indexes between strategies.

[16] The genetic algorithm parameters were chosen in accordance with parameters suggested by computer scientists who use genetic algorithms for complex numerical search tasks (e.g., see Goldberg, 1989).

[17] The mean value $m = 10^6$ is a pedagogical value only; it could represent any number of units of money, e.g., cents.

[18] This term is inspired by a similar term for classifying genes in biological evolution.

The idea of a phenotype is best illustrated via some examples. Phenotype "L 1 2 E 2 1" characterizes a strategy where, in the initial, first stage game strategy the bidder's final bid (any fraction of his value) is made late in period T; hence the "L" in the first position. His second stage game strategy involves placing a final bid (of any fraction) in any period prior to period T and is therefore labeled as "E" in position 4. Positions 2—3 and 5—6 in the phenotype indicate strategy transition behavior conditional on whether the realized history by the bidders' rival bidders was late or early (as in the characterization of strategies). If while playing the late-bidding strategy 1, any rival bid also arrives late, so that the history is L, the integer 1 in position 2 indicates that this bidder will continue with the late-bidding strategy 1 for one more auction. However, if no rival bids arrive late, so that the history is E, this bidder will transition to the early-bidding strategy 2 as indicated by the integer 2 in position 3. If while playing this early-bidding strategy, the history of rival bids is L, the bidder will stick with the early-bidding strategy 2, otherwise he transitions back to the late-bidding strategy 1 as indicated by the integers 2 and 1 in positions 5–6.

We note that there are just 22 phenotypes that are possible. This number is less than 2^6 because certain unconditional strategies reduce the set of phenotypes necessary to characterize strategies. For instance, phenotype "E 1 1 L 2 1" is more compactly characterized simply as "E" denoting unconditional early bidding; the strategy starts off bidding early (strategy 1) and never moves away from this strategy (it ignores the history of rival bids). Similarly, phenotype "L" denotes unconditional late bidding. The classification of repeated game strategies into phenotypes is further illustrated in the following examples.

Example 2: The phenotype of the soft-close strategy in Example 1, $((\mathbf{0}, \mathbf{2}, \mathbf{3}, \mathbf{1}), [1, 2]; (\mathbf{8}, \mathbf{4}, \mathbf{5}, \mathbf{0}), [2, 1])$ is "E 1 2 L 2 1."

Example 3: The following strategies have phenotype "L" (unconditional late bidding):
hard-close auction strategy $((\mathbf{8}, \mathbf{6}, \mathbf{1}), [1, 1]; (\mathbf{2}, \mathbf{4}, \mathbf{5}), [1, 2])$
soft-close auction strategy $((\mathbf{0}, \mathbf{8}, \mathbf{2}, \mathbf{0}), [2, 1]; (\mathbf{8}, \mathbf{1}, \mathbf{2}, \mathbf{0}), [1, 1])$

Example 4: The following strategies have phenotype "E" (unconditional early bidding):
soft-close auction strategy $((\mathbf{0}, \mathbf{7}, \mathbf{1}, \mathbf{1}), [1, 1]; (\mathbf{8}, \mathbf{4}, \mathbf{5}, \mathbf{1}), [1, 2])$
hard-close auction strategy $((\mathbf{0}, \mathbf{4}, \mathbf{2}), [2, 1]; (\mathbf{5}, \mathbf{8}, \mathbf{2}), [1, 2])$

In the next section, we use this phenotype classification scheme to characterize the main findings from our benchmark simulations.

5 Results

In this section and the next we summarize the main results from our simulation exercises.

Result 1: When $n < 5$, the percentage of late bidders (submitting bids in period T) is greater in hard-close auctions than in soft-close auctions.

HARD-CLOSE	*Number of Bidders - Benchmark Simulations*							
Statistics	$n = 2$		$n = 3$		$n = 4$		$n = 5$	
Number of early bids per bidder	1.96	(0.044)***	1.60	(0.043)	1.40	(0.042)	1.41	(0.046)**
Fraction submitting late bid	0.58	(0.024)***	0.32	(0.038)***	0.071	(0.022)**	0.0031	(0.00086)
Fraction learning to bid full value	0.99	(0.0062)**	0.998	(0.0010)	0.998	(0.00078)**	0.998	(0.00071)
Revenue of the seller/mean value	0.95	(0.0037)**	0.998	(0.00046)***	0.99999	(0.000037)	1.000161	(0.0000013)
Payoff of bidders/mean value	0.024	(0.0018)**	0.0013	(0.00015)***	0.000076	(0.0000091)	0.000033	(0.00000024)
Freq. of early-bidding automata	0.15	(0.019)***	0.55	(0.067)***	0.89	(0.11)	0.97	(0.12)
Freq. of late-bidding automata	0.0082	(0.0013)	0.053	(0.0074)***	0.0037	(0.00065)	0.0012	(0.00022)
Freq. of cond.-bidding automata	0.84	(0.101)***	0.40	(0.049)***	0.11	(0.012)***	0.027	(0.0034)
SOFT-CLOSE	*Number of Bidders - Benchmark Simulation*							
Statistics	$n = 2$		$n = 3$		$n = 4$		$n = 5$	
Number of early bids per bidder	1.48	(0.046)	1.53	(0.048)	1.36	(0.040)	1.30	(0.040)
Fraction submitting a late bid	0.088	(0.021)	0.036	(0.012)	0.023	(0.012)	0.0026	(0.0010)
Fraction learning to bid full value	0.96	(0.012)	0.998	(0.00074)	0.997	(0.0015)	0.998	(0.00077)
Revenue of the seller / mean value	0.97	(0.0076)	0.9998	(0.000099)	0.999956	(0.0000131)	1.000162	(0.00000079)
Payoff of bidders / mean value	0.014	(0.0038)	0.000168	(0.000033)	0.000084	(0.000054)	0.000033	(0.00000015)
Freq. of early-bidding automata	0.90	(0.11)	0.91	(0.11)	0.93	(0.11)	0.97	(0.12)
Freq. of late-bidding automata	0.0086	(0.0016)	0.0051	(0.0011)	0.0026	(0.00050)	0.0020	(0.00044)
Freq. of cond.-bidding automata	0.09	(0.011)	0.082	(0.0095)	0.063	(0.0068)	0.030	(0.0046)

Table 1: **Benchmark Simulation Results** (Averages and Standard Deviations Over the Last 100 Generations.)

Support for Result 1 is found in Table 1 (see the row labeled "Fraction submitting late bid") and Figure 1 where we observe that the fraction of bidders submitting a late bid is greater in hard-close (Figure1a) than in soft-close auctions (Figure 1b) with the same number of bidders, n =2, 3, or 4. When $n = 5$ there is no difference in the frequency of late bidding across auction formats.[19] These percentages and the ones reported for other statistics in the Tables and Figures that follow were obtained by taking averages over the last 100 generations of the $S = 20$ simulations run for each treatment.[20] The statistical significance of differences in the statistics reported in Table 1 for hard- and soft-close auctions with the same number of bidders was tested using two-sample, one-sided t-tests with $2S - 2 = 38$ degrees of freedom. That is, for each treatment we take the 20 simulation runs as independent observations. If a reported statistic for the hard-close auction was significantly greater than for the soft-close auction with the same n, this finding is indicated next to the standard errors reported for the *hard-close* auctions; significance at the 1% level is indicated by ***; at the 5% level by **. (Unless otherwise noted, we apply the same two-sample, one-sided t-test for the other statistics reported in all Tables).[21] When $n = 5$, the equality of the two means cannot be rejected using a two-sided t-test with a p-value of 0.70.

[19] For the frequency of bidders *attempting* a late bid, we could divide these percentages roughly by $\rho = 0.9$, the probability of registering a successful bid in period T.

[20] Recall that we consider late bids only in the "first late" period, i.e., period T of soft-close auctions.

[21] We do not report one-sided t-test results with significance levels between 5 and 10% as such levels would not allow us to reject the null hypothesis that the two means are equal in a *two*-sided test at 10% significance level.

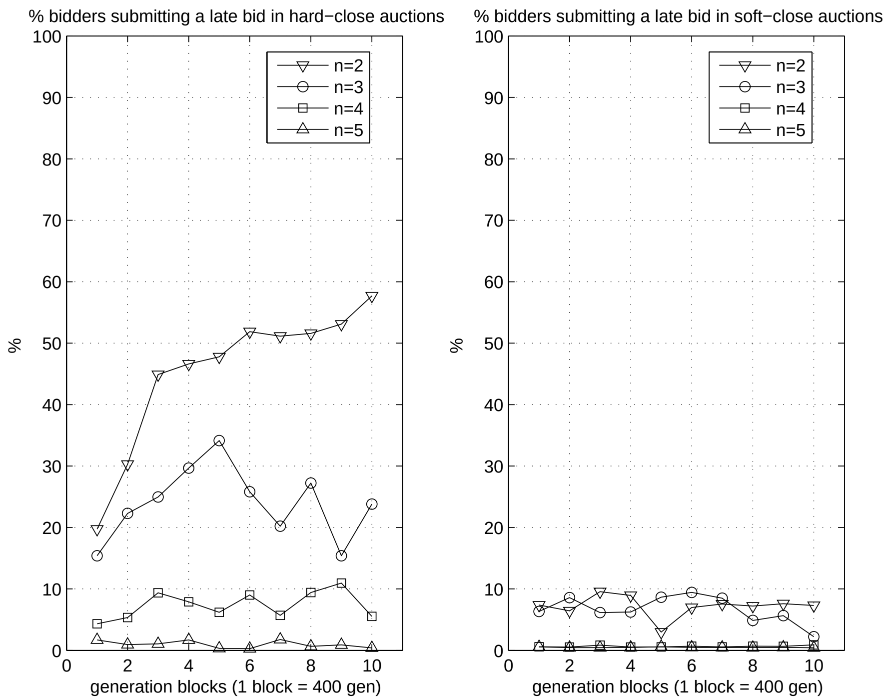

Figure 1: The average percentage of bidders successfully submitting a late bid in hard-close and soft-close auctions with 2,3,4 and 5 bidders.

What accounts for the difference in the percentage of late bidding between the two auction formats? It appears to be due to both a hard close to the auction and some last-minute congestion, i.e., $\rho < 1$. Indeed, a sensitivity analysis reported below for the case of two and three bidders shows that large differences in the percentage of late bidding between auction formats emerge only when $\rho < 1$; when $\rho = 1$, there is little difference in the percentage of late bidding between the two auction formats (though, it is still significant).

For our sensitivity analysis, we change *one and only one* model parameter at a time keeping all other parameter choices at their baseline values and focusing on auctions with 2 or 3 bidders. The results of these exercises are reported in Tables 2 - 9. Specifically, in Table 2 for $n = 2$ and in Table 6 for $n = 3$, we change (1) the number of permissible fractions of private valuations that can be bid from $f = 3$ in the benchmark simulations to 10 or 30; Setting f back to the baseline value of 3, we change (2) the probability of successful last minute bid registration from the benchmark value of $\rho = 0.9$ to 0.8 or 1. In Table 3 for $n = 2$ and in Table 7 for $n = 3$, we change (3) the number of auction stages in a repeated game block from $R = 20$ in the benchmark simulations to $R = 1$ and then $R = 40$. As exercise (4) reported in Tables 3 and 7, (under the heading "incr. late bids,") we examine a modification to the extension period strategies of soft-close auctions; in the benchmark simulation, bidders can only bid their full valuation in extension periods but in this exercise we permit bidders to submit bids other than their full, private value in soft-close extension periods. In particular, they are now allowed to bid the smallest permissible fraction of private value that exceeds their latest bid (i.e., in extension periods, bidders become incremental bidders).[22] In Table 4 for $n = 2$ and in Table 8 for $n = 3$, we change (5) the spread of the value distribution from $\epsilon = 1000$ in the benchmark simulations to 50, 10^5 and 10^6. In Table 5 for $n = 2$ and in Table 9 for $n = 3$, we change (6) the mean of the value distribution from $m = 10^6$ in the benchmark simulations to 1000 and then to 10^9; and finally, (7) the number of possible values in the support from $n_V = 101$ in the benchmark simulations to 6 and then to 10001. For each exercise with 2 or 3 bidders, we conduct 20 simulation runs. With a few exceptions, the results found for the benchmark simulations do not change.

Consider the frequency of late bidding in hard-close auctions when we change ρ from the benchmark $\rho = 0.9$ to $\rho = 0.8$ or $\rho = 1$. In the $n = 2$ case, we observe that the frequency of late bidding when $\rho = 0.9$ is significantly higher than when $\rho = 0.8$ or $\rho = 1$ with one-sided two-sample t-test p-values of 0.017 and 3.68×10^{-5}. For $n = 3$, we cannot reject the equivalence of the late-bidding frequencies in the $\rho = 0.8$ and the benchmark $\rho = 0.9$ cases, with a two-sided p-value of 0.16, while the benchmark $\rho = 0.9$ case leads to more frequent late bidding than the $\rho = 1$ case with a one-sided p-value of 1.73×10^{-5}. As for the frequency of late bidding in soft-close auctions, the null of no difference between the benchmark $\rho = 0.9$ and the $\rho = 0.8$ case cannot be rejected with a two-sided p-value of 0.18 when $n = 2$ and a two-sided p-value of 0.16 when $n = 3$; the late-bidding frequency in the benchmark case is significantly lower than

[22] Incremental bidding instead of whole value bidding in extension periods of a soft-close auction induces *only slightly* higher late bidding and decreases revenue of the seller, since simultaneous whole values are not *successfully* submitted as bids as frequently as in the benchmark simulations.

HARD-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$f = 10$		$f = 30$		$\rho = 0.8$		$\rho = 1$	
Number of early bids per bidder	2.91	(0.052)***	3.09	(0.071)***	1.78	(0.037)***	1.71	(0.035)***
Fraction submitting late bid	0.63	(0.021)***	0.59	(0.025)***	0.34	(0.025)***	0.26	(0.031)**
Fraction learning to bid full value	0.9995	(0.00029)***	0.999	(0.00043)	0.998	(0.00073)***	0.956	(0.019)
Revenue of the seller/mean value	0.986	(0.00046)***	0.995	(0.00020)***	0.95	(0.00038)***	0.98	(0.0067)***
Payoff of bidders/mean value	0.0070	(0.00023)***	0.0024	(0.000098)***	0.027	(0.00019)***	0.0081	(0.0034)***
Freq. of early-bidding automata	0.11	(0.015)***	0.15	(0.019)***	0.37	(0.046)***	0.57	(0.069)
Freq. of late-bidding automata	0.061	(0.0091)***	0.022	(0.0036)***	0.011	(0.0016)**	0.085	(0.011)**
Freq. of cond.-bidding automata	0.83	(0.10)***	0.82	(0.10)***	0.62	(0.075)***	0.35	(0.042)***
SOFT-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$f = 10$		$f = 30$		$\rho = 0.8$		$\rho = 1$	
Number of early bids per bidder	2.23	(0.068)	2.25	(0.059)	1.42	(0.033)	1.48	(0.0092)
Fraction submitting late bid	0.064	(0.018)	0.04	(0.014)	0.055	(0.013)	0.19	(0.0066)
Fraction learning to bid full value	0.976	(0.0079)	0.994	(0.0034)	0.974	(0.0084)	0.961	(0.012)
Revenue of the seller/mean value	0.993	(0.0020)	0.999	(0.00027)	0.98	(0.0067)	0.83	(0.0056)
Payoff of bidders/mean value	0.0034	(0.00098)	0.00061	(0.00013)	0.012	(0.0033)	0.086	(0.0027)
Freq. of early-bidding automata	0.92	(0.11)	0.94	(0.11)	0.88	(0.11)	0.73	(0.089)
Freq. of late-bidding automata	0.027	(0.0043)	0.017	(0.0034)	0.016	(0.0024)	0.06	(0.0074)
Freq. of cond.-bidding automata	0.054	(0.0071)	0.043	(0.0071)	0.11	(0.012)	0.21	(0.024)

Table 2: Sensitivity Analysis 1 for $n = 2$ (Averages over the last 100 generations)

HARD-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$R = 1$		$R = 40$			
Number of early bids per bidder	1.55	(0.049)	1.92	(0.033)***		
Fraction submitting late bid	0.014	(0.0062)	0.65	(0.016)***		
Fraction learning to bid full value	0.997	(0.00121)	0.999	(0.00048)***		
Revenue of seller / mean value	0.996	(0.0013)	0.95	(0.0012)***		
Payoff of bidders / mean value	0.0021	(0.00064)	0.024	(0.00062)***		
Freq. of early-bidding automata	0.84	(0.10)	0.063	(0.0078)***		
Freq. of late-bidding automata	0.013	(0.0021)***	0.0074	(0.0010)***		
Freq. of cond.-bidding automata	0.15	(0.017)	0.93	(0.11)***		
SOFT-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$R = 1$		$R = 40$		incr. late bids	
Number of early bids per bidder	1.54	(0.049)	1.37	(0.042)	1.46	(0.0088)
Fraction submitting late bid	0.0044	(0.0018)	0.11	(0.024)	0.17	(0.006)
Fraction learning to bid full value	0.997	(0.0014)	0.97	(0.0092)	0.78	(0.0068)
Revenue of seller / mean value	0.996	(0.0015)	0.97	(0.0070)	0.80	(0.0057)
Payoff of bidders / mean value	0.0037	(0.0015)	0.014	(0.0035)	0.099	(0.0028)
Freq. of early-bidding automata	0.83	(0.10)	0.88	(0.11)	0.73	(0.089)
Freq. of late-bidding automata	0.0020	(0.00048)	0.027	(0.0037)	0.061	(0.0074)
Freq. of cond.-bidding automata	0.17	(0.019)	0.094	(0.010)	0.21	(0.025)

Table 3: Sensitivity Analysis 2 for $n = 2$ (Averages over the last 100 generations)

HARD-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$\epsilon = 50$		$\epsilon = 10^5$		$\epsilon = 10^6$	
Number of early bids per bidder	1.90	(0.042)***	1.74	(0.044)	1.50	(0.042)
Fraction submitting late bid	0.56	(0.026)***	0.32	(0.034)***	0.032	(0.011)
Fraction learning to bid full value	0.997	(0.0011)***	0.997	(0.0012)**	0.997	(0.00095)
Revenue of seller / mean value	0.95	(0.0024)***	0.96	(0.0025)**	0.83	(0.00066)***
Payoff of bidders / mean value	0.023	(0.0012)***	0.029	(0.0012)**	0.17	(0.00027)
Freq. of early-bidding automata	0.18	(0.022)***	0.53	(0.065)***	0.93	(0.11)
Freq. of late-bidding automata	0.0077	(0.0011)***	0.024	(0.0038)**	0.0039	(0.00064)**
Freq. of cond.-bidding automata	0.81	(0.098)***	0.45	(0.053)***	0.066	(0.0068)***
SOFT-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$\epsilon = 50$		$\epsilon = 10^5$		$\epsilon = 10^6$	
Number of early bids per bidder	1.46	(0.0088)	1.64	(0.043)	1.43	(0.036)
Fraction submitting late bid	0.17	(0.006)	0.050	(0.015)	0.020	(0.0071)
Fraction learning to bid full value	0.986	(0.0077)	0.98	(0.0079)	0.99	(0.0039)
Revenue of seller / mean value	**0.81**	(0.0055)	0.97	(0.0048)	0.82	(0.0018)
Payoff of bidders / mean value	0.096	(0.0027)	0.023	(0.0023)	0.17	(0.00068)
Freq. of early-bidding automata	0.73	(0.089)	0.94	(0.11)	0.96	(0.12)
Freq. of late-bidding automata	0.06	(0.0074)	0.015	(0.0029)	0.0074	(0.0015)
Freq. of cond.-bidding automata	0.21	(0.024)	0.049	(0.0065)	0.029	(0.0040)

Table 4: Sensitivity Analysis 3 for $n = 2$ (Averages over the last 100 generations)

HARD-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$m = 1000$		$m = 10^9$		$n_V = 6$		$n_V = 10001$	
Number of early bids per bidder	1.57	(0.038)***	1.88	(0.044)***	1.91	(0.036)***	2.00	(0.041)***
Fraction submitting late bid	0.023	(0.0077)	0.57	(0.026)***	0.64	(0.018)***	0.59	(0.023)***
Fraction learning to bid full value	0.998	(0.00068)	0.998	(0.0011)***	0.998	(0.00079)	0.998	(0.00075)**
Revenue of seller / mean value	0.83	(0.00061)	0.96	(0.0019)	0.95	(0.0015)***	0.955	(0.0018)***
Payoff of bidders / mean value	0.17	(0.00025)	0.021	(0.00096)	0.024	(0.00074)***	0.022	(0.00091)***
Freq. of early-bidding automata	0.94	(0.11)	0.17	(0.021)***	0.082	(0.010)***	0.14	(0.018)***
Freq. of late-bidding automata	0.0039	(0.00076)**	0.013	(0.0022)***	0.015	(0.0023)	0.021	(0.0032)**
Freq. of cond.-bidding automata	0.052	(0.0055)***	0.82	(0.099)***	0.90	(0.11)***	0.84	(0.10***
SOFT-CLOSE with 2 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$m = 1000$		$m = 10^9$		$n_V = 6$		$n_V = 10001$	
Number of early bids per bidder	1.42	(0.034)	1.60	(0.042)	1.54	(0.052)	1.39	(0.037)
Fraction submitting late bid	0.024	(0.0063)	0.14	(0.025)	0.083	(0.022)	0.047	(0.017)
Fraction learning to bid full value	0.997	(0.0010)	0.964	(0.011)	0.991	(0.0053)	0.987	(0.0061)
Revenue of seller / mean value	0.83	(0.00083)	0.97	(0.0071)	0.988	(0.0038)	0.988	(0.0045)
Payoff of bidders / mean value	0.17	(0.00034)	0.015	(0.0035)	0.0062	(0.0019)	0.0059	(0.0023)
Freq. of early-bidding automata	0.92	(0.11)	0.78	(0.095)	0.88	(0.11)	0.94	(0.11)
Freq. of late-bidding automata	0.0022	(0.00044)	0.060	(0.0078)	0.018	(0.0034)	0.031	(0.0052)
Freq. of cond.-bidding automata	0.080	(0.0092)	0.16	(0.019)	0.099	(0.011)	0.031	(0.0054)

Table 5: Sensitivity Analysis 4 for $n = 2$ (Averages over the last 100 generations)

HARD-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$f = 10$		$f = 30$		$\rho = 0.8$		$\rho = 1$	
Number of early bids per bidder	2.62	(0.085)***	2.58	(0.079)	1.71	(0.047)***	1.51	(0.037)
Fraction submitting late bid	0.30	(0.039)***	0.25	(0.036)***	0.25	(0.032)***	0.11	(0.025)**
Fraction learning to bid full value	0.998	(0.00071)	0.998	(0.0016)	0.998	(0.0011)***	0.998	(0.0012)
Revenue of seller / mean value	0.99	(0.0020)**	1.00	(0.000067)***	0.99	(0.0014)***	1.00	(0.00026)
Payoff of bidders / mean value	0.0018	(0.00065)**	0.00019	(0.000022)***	0.0035	(0.00045)***	0.00015	(0.000086)
Freq. of early-bidding automata	0.60	(0.073)**	0.68	(0.083)**	0.64	(0.077)***	0.82	(0.10)
Freq. of late-bidding automata	0.039	(0.0051)***	0.074	(0.010)***	0.034	(0.0050)***	0.050	(0.0064)**
Freq. of cond.-bidding automata	0.36	(0.044)***	0.25	(0.031)***	0.33	(0.039)***	0.13	(0.014)***
SOFT-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$f = 10$		$f = 30$		$\rho = 0.8$		$\rho = 1$	
Number of early bids per bidder	2.03	(0.057)	2.41	(0.072)	1.29	(0.0071)	1.47	(0.045)
Fraction submitting late bid	0.067	(0.021)	0.036	(0.014)	0.12	(0.0043)	0.18	(0.027)
Fraction learning to bid full value	0.993	(0.0043)	0.998	(0.00071)	0.78	(0.0066)	0.993	(0.0042)
Revenue of seller / mean value	1.00	(0.00079)	1.00	(0.0000086)	0.93	(0.0028)	1.00	(0.0012)
Payoff of bidders / mean value	0.00040	(0.00026)	0.000091	(0.0000028)	0.022	(0.00094)	0.00055	(0.00040)
Freq. of early-bidding automata	0.90	(0.11)	0.93	(0.11)	0.73	(0.089)	0.64	(0.079)
Freq. of late-bidding automata	0.0057	(0.00094)	0.0035	(0.00060)	0.061	(0.0074)	0.069	(0.0092)
Freq. of cond.-bidding automata	0.089	(0.011)	0.064	(0.0083)	0.21	(0.024)	0.29	(0.035)

Table 6: Sensitivity Analysis 1 for $n = 3$ (Averages over the last 100 generations)

HARD-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$R = 1$		$R = 40$			
Number of early bids per bidder	1.40	(0.047)	1.66	(0.046)		
Fraction submitting late bid	0.0015	(0.00073)	0.24	(0.037)***		
Fraction learning to bid full value	0.997	(0.0017)	0.96	(0.017)		
Revenue of seller / mean value	1.00	(0.00042)	0.98	(0.0064)		
Payoff of bidders / mean value	0.00017	(0.00014)	0.0056	(0.0021)		
Freq. of early-bidding automata	0.93	(0.11)	0.69	(0.083)		
Freq. of late-bidding automata	0.0013	(0.00028)	0.090	(0.011)***		
Freq. of cond.-bidding automata	0.065	(0.0073)	0.22	(0.026)		
SOFT-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$R = 1$		$R = 40$		incr. late bids	
Number of early bids per bidder	1.34	(0.049)	1.71	(0.042)	1.30	(0.0072)
Fraction submitting late bid	0.0023	(0.0016)	0.13	(0.025)	0.016	(0.0015)
Fraction learning to bid full value	0.998	(0.00096)	0.984	(0.0079)	0.77	(0.0068)
Revenue of seller / mean value	1.00	(0.000088)	0.99	(0.0030)	0.94	(0.00028)
Payoff of bidders / mean value	0.00013	(0.000029)	0.0020	(0.0010)	0.020	(0.00095)
Freq. of early-bidding automata	0.84	(0.10)	0.80	(0.10)	0.73	(0.089)
Freq. of late-bidding automata	0.0021	(0.00055)	0.025	(0.0037)	0.061	(0.0074)
Freq. of cond.-bidding automata	0.15	(0.018)	0.18	(0.021)	0.021	(0.025)

Table 7: Sensitivity Analysis 2 for $n = 3$ (Averages over the last 100 generations)

HARD-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$\epsilon = 50$		$\epsilon = 10^5$		$\epsilon = 10^6$	
Number of early bids per bidder	1.77	(0.041)***	1.52	(0.046)	1.45	(0.043)
Fraction submitting late bid	0.32	(0.038)***	0.12	(0.027)***	0.025	(0.0089)**
Fraction learning to bid full value	0.999	(0.0013)**	0.95	(0.019)***	0.998	(0.00075)
Revenue of seller / mean value	1.00	(0.00048)	0.98	(0.0068)***	0.99	(0.00059)
Payoff of bidders / mean value	0.0011	(0.00016)	0.014	(0.0022)***	0.084	(0.00014)
Freq. of early-bidding automata	0.56	(0.069)***	0.84	(0.10)	0.93	(0.11)
Freq. of late-bidding automata	0.11	(0.014)***	0.022	(0.0036)***	0.0084	(0.0018)***
Freq. of cond.-bidding automata	0.33	(0.039)***	0.14	(0.015)***	0.059	(0.0064)
SOFT-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*					
Statistics	$\epsilon = 50$		$\epsilon = 10^5$		$\epsilon = 10^6$	
Number of early bids per bidder	1.39	(0.039)	1.51	(0.038)	1.36	(0.039)
Fraction submitting late bid	0.092	(0.024)	0.033	(0.014)	0.0064	(0.0020)
Fraction learning to bid full value	0.988	(0.0060)	0.999	(0.00064)	0.998	(0.00089)
Revenue of seller / mean value	1.00	(0.0021)	1.00	(0.00012)	0.99	(0.00049)
Payoff of bidders / mean value	0.0013	(0.00069)	0.0084	(0.000029)	0.084	(0.00012)
Freq. of early-bidding automata	0.89	(0.11)	0.94	(0.11)	0.94	(0.11)
Freq. of late-bidding automata	0.0079	(0.0015)	0.0067	(0.0012)	0.0020	(0.00043)
Freq. of cond.-bidding automata	0.10	(0.011)	0.052	(0.0055)	0.057	(0.0061)

Table 8: Sensitivity Analysis 3 for $n = 3$ (Averages over the last 100 generations)

HARD-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$m = 1000$		$m = 10^9$		$n_V = 6$		$n_V = 10001$	
Number of early bids per bidder	1.43	(0.042)***	1.74	(0.052)	1.66	(0.050)	1.73	(0.051)***
Fraction submitting late bid	0.022	(0.0079)**	0.29	(0.038)***	0.39	(0.040)***	0.29	(0.038)***
Fraction learning to bid full value	0.995	(0.0017)	0.999	(0.00056)**	0.958	(0.018)**	0.958	(0.018)
Revenue of seller / mean value	0.99	(0.00092)	1.00	(0.00045)***	0.98	(0.0065)***	0.98	(0.0065)
Payoff of bidders / mean value	0.084	(0.00020)	0.0011	(0.00015)***	0.0063	(0.0022)***	0.0059	(0.0022)
Freq. of early-bidding automata	0.95	(0.12)	0.61	(0.075)**	0.52	(0.064)**	0.63	(0.076)
Freq. of late-bidding automata	0.011	(0.0018)***	0.068	(0.0093)***	0.065	(0.0085)***	0.090	(0.012)***
Freq. of cond.-bidding automata	0.042	(0.0055)	0.32	(0.04)***	0.41	(0.049)***	0.28	(0.034)***
SOFT-CLOSE with 3 bidders	*Varying Parameters in the Benchmark Simulations*							
Statistics	$m = 1000$		$m = 10^9$		$n_V = 6$		$n_V = 10001$	
Number of early bids per bidder	1.23	(0.031)	1.45	(0.045)	1.59	(0.046)	1.38	(0.040)
Fraction submitting late bid	0.0066	(0.0023)	0.081	(0.022)	0.12	(0.026)	0.11	(0.026)
Fraction learning to bid full value	0.997	(0.0011)	0.995	(0.0023)	0.999	(0.00067)	0.982	(0.0094)
Revenue of seller / mean value	0.99	(0.00051)	1.00	(0.00057)	1.00	(0.00014)	0.99	(0.0037)
Payoff of bidders / mean value	0.083	(0.00012)	0.00035	(0.00019)	0.00031	(0.000048)	0.0023	(0.0012)
Freq. of early-bidding automata	0.97	(0.12)	0.87	(0.11)	0.80	(0.10)	0.83	(0.10)
Freq. of late-bidding automata	0.00078	(0.00024)	0.012	(0.0019)	0.011	(0.0018)	0.0089	(0.0013)
Freq. of cond.-bidding automata	0.031	(0.0038)	0.12	(0.014)	0.19	(0.022)	0.16	(0.018)

Table 9: Sensitivity Analysis 4 for $n = 3$ (Averages over the last 100 generations)

in the $\rho = 1$ case with a one-sided p-value of 2.07×10^{-5} when $n = 2$ and a one-side p-value of 1.73×10^{-5} when $n = 3$.

With a hard close and $\rho < 1$ (congestion effects) bidders may wish to delay bidding their full valuation in a tacitly collusive arrangement that raises the expected profits accruing to all bidders. Expected profits rise because the delay in bidding full valuations until the last period provides some chance that rival bidders are unsuccessful in placing a final bid allowing the successful late bidder(s) to potentially profit. By contrast, in soft-close auctions the gains to this collusive strategy are reduced by the fact that the auction will continue if just one late bidder succeeds in bidding late. This logic follows from Ockenfels and Roth's (2006) theoretical analysis. However, Ockenfels and Roth's hard-close auction model with two fully rational bidders, has many other Bayesian Nash equilibria than the one they highlight, where all bid late; for instance, all bidding early is another equilibrium of their model. Note that our computational model is different from Ockenfels and Roth's in the sense that (1) we do not have continuous bidding periods (which underlies the result of Ockenfels and Roth), (2) we permit repeated game strategies, and (3) we consider automata as repeated game strategies. So, the value–added of our approach is that our boundedly rational (and more cognitively plausible) adaptive search algorithm serves as a kind of equilibrium selection device –evolving strategies consistent with late bidding in certain hard-close auctions and early bidding in all soft-close auctions – a result that is consistent with the empirical evidence. Note that for Result 1 to hold ρ should neither be too low nor too high. As ρ decreases, the chance of successful registration of a bidder's own late bid falls, and this effect will eventually outweigh the expected collusive gains to waiting to bid at the last minute. On the other hand as ρ approaches 1, those collusive gains evaporate and there is no cause for waiting to bid until the last period.

Notice also that for the $\rho = 1$ case, we observe an 11-26% frequency of late bidding both in hard-close and in soft-close auctions, and this amount of late bidding in the soft-close auction is significantly greater than in soft-close auctions with $\rho < 1$. The reason for this amount of late bidding is that bidders in the $\rho = 1$ are indifferent between bidding early or late and so bidding their full value at any time period is a best response. Later in Section 6 we show that this result may change with the addition of naive incremental bidders.

Result 1 is robust for a wide range of parameters. In addition to the hard-close auction format, the number of bidders, n, and the last-minute congestion parameter, ρ, two other parameters appear to play a critical role in establishing that result. First, an increase in the number of stages, R, in a repeated-auction block leads to an increase in the frequency of late bidding in hard-close auctions. There is no late bidding when $R = 1$. Both of these results are statistically significant: With $n = 2$ bidders and $R = 40$, the frequency of late bidding is significantly greater than in the benchmark case, where $R = 20$, with a one-sided p-value of 0.010. Similarly when $n = 2$ and $R = 1$, the frequency of late bidding is significantly lower than in the benchmark case, with a one-sided p-value of 4.24×10^{-24} and no different from that found in the soft-close case.[23] This result is not surprising. As our theoretical results predict, in a single-shot ($R = 1$) auction, early bidding of full value is a dominant strategy. On the other hand, increasing the number of stage games in a repeated-auction increases the expected surplus generated for all bidders who play a collusive, late-bidding strategy. Second, the mean-to-spread ratio for the distribution of values, $\frac{m}{\epsilon}$, should be high enough to sustain late bidding in hard-close auctions. For example, with $n = 2$ or 3 and our baseline value of $m = 10^6$, when ϵ is increased from its baseline value of 1000 to 10^5 so that $\frac{m}{\epsilon}$ decreases from the baseline ratio of 1000 to 10, there is 32% late bidding in the $n = 2$ and 12% late bidding in the $n = 3$, hard-close auctions. When $m = 1000$ and ϵ is set at the baseline value of 1000, so that $\frac{m}{\epsilon} = 1$, there is less than 3% late bidding in hard-close auctions in both the $n = 2$ and 3 cases.

Since our results appear sensitive to the ratio $\frac{m}{\epsilon}$, we conducted a more careful analysis of the impact of this ratio on our findings. We set $\epsilon = 100$, 1000, and 10000, and set $m = k\epsilon$ for all $k \in \{2, 3, 4, 5, 6, 7, 8, 9\}$. We report the fraction of bidders submitting a late bid in hard- and soft-close auctions in the $n = 2$ case in Table 10. We also check, as before, the difference between the hard-/soft-close treatment using t-tests. We observe that, so long as $\frac{m}{\epsilon} \geq 3$ (a few exceptions withstanding), there is significantly more late bidding under the hard-close format than under the soft-close format.

How reasonable are such distributions of buyers' valuations? This is a difficult question to address as we do not have data on private valuations. Further, we don't specify how units of private valuation are converted into monetary terms; depending on the conversion rate, we could have either tightly or widely dispersed distributions of valuations in monetary terms. We note that late bidding is observed even for goods where the distribution of private valuations is likely to be tightly centered about its mean; an example is eBay auctions of computer monitors, as studied by Gonzalez et al. (2004). In the case of a monitor, the mean private valuation may be well-proxied by posted-prices for the same monitor. As such posted-price information is public (e.g., via seller websites), one can imagine that private valuations are tightly centered about the mean. Still, Gonzalez et al. report evidence of sniping in eBay monitor auctions, especially by experienced eBay users, who do not appear to be purchasing for resale purposes. Given this example, and our other observations, we believe the distributions of private valuations we require to obtain late bidding may be quite reasonable.

Our next result concerns how the frequency of late bidding is affected by the number of bidders.

Result 2: The frequency of late bidding in hard-close auctions decreases as the number of bidders increases.

Support for this finding is found in Figure 1a and in Table 1. The statistical comparison *between* auction formats holding n constant is reported in Table 1. Using the standard errors of that table, we

[23] With $n = 3$ bidders, the difference between the benchmark and $R = 40$ case is not significant with a two-sided p-value of 0.13, and but the difference between the benchmark case and the $R = 1$, remains significant with a one-sided p-value of 1.82×10^{-10}.

2 bidders		$m = 2\epsilon$	$m = 3\epsilon$	$m = 4\epsilon$	$m = 5\epsilon$
$\epsilon = 100$	**Hard Close:**	0.018 (0.0073)	0.040 (0.013)***	0.13 (0.026)**	0.099 (0.021)***
	Soft Close:	0.006 (0.002)	0.0082 (0.0198)	0.050 (0.013)	0.043 (0.013)
$\epsilon = 1000$	**Hard Close:**	0.019 (0.008)	0.059 (0.016)***	0.057 (0.015)	0.17 (0.018)***
	Soft Close:	0.052 (0.016)	0.0098 (0.0036)	0.075 (0.02)	0.047 (0.013)
$\epsilon = 10000$	**Hard Close:**	0.021 (0.0088)	0.016 (0.0067)	0.11 (0.025)***	0.14 (0.025)***
	Soft Close:	0.011 (0.0050)	0.026 (0.0096)	0.010 (0.0043)	0.034 (0.011)

2 bidders		$m = 6\epsilon$	$m = 7\epsilon$	$m = 8\epsilon$	$m = 9\epsilon$
$\epsilon = 100$	**Hard Close:**	0.13 (0.026)***	0.26 (0.032)***	0.29 (0.032)***	0.17 (0.028)***
	Soft Close:	0.020 (0.0069)	0.030 (0.011)	0.043 (0.014)	0.042 (0.016)
$\epsilon = 1000$	**Hard Close:**	0.15 (0.027)***	0.26 (0.032)***	0.44 (0.032)***	0.35 (0.034)***
	Soft Close:	0.052 (0.015)	0.027 (0.011)	0.059 (0.017)	0.019 (0.0074)
$\epsilon = 10000$	**Hard Close:**	0.18 (0.028)***	0.076 (0.020)	0.34 (0.034)***	0.31 (0.033)***
	Soft Close:	0.016 (0.0056)	0.050 (0.015)	0.075 (0.019)	0.0064 (0.0020)

Table 10: Sensitivity Analysis: Variation in the Ratio m/ϵ and Fraction of Bidders Submitting a Late Bid

can provide a statistical analysis of the effect of changes in the number of bidders *within* the same auction format. For the hard-close auction, the difference between late-bidding percentages with two bidders and three bidders is positive and significant with a one-sided p-value of 5.62×10^{-7}; between three and four bidders, it is significantly positive with a one-sided p-value of 8.05×10^{-7}; between four and five bidders, it is significantly positive with a one-sided p-value of 0.0019. Similarly for the soft-close auctions, the respective p-values are 0.019, 0.44, and 0.049 (where the second "insignificant" p-value is two-sided and the others are one-sided). Though some differences are significant in the soft-close auctions, we observe only 0.26% to 8.8% late bidding in soft-close auctions. The intuition for Result 2 is straightforward. The expected gains from the collusive strategy diminish as the number of bidders increases; these gains are proportional to $(1-\rho)^{n-1}$, the chance that the other $n-1$ bidders' bids fail to register in the last period, and this factor is decreasing in n.

Result 3: For $n = 2$ or 3, the average revenue of sellers is significantly higher in soft-close than in hard-close auctions and the average revenue of bidders is significantly lower in soft-close than in hard-close auctions.

Support for this finding is again found in Table 1. We report here the p-values from that Table's pairwise comparisons between auction formats. Average seller revenue is higher in the soft-close than in the hard-close auctions with a one-sided p-value of 0.012 in the $n = 2$ case and a one-sided p-value of 4.5×10^{-10} in the $n = 3$ case. The null hypothesis of no difference in seller revenue between auction formats cannot be rejected in the case of $n = 4$ or 5 bidders. Similarly average bidder payoffs are significantly lower in the $n = 2$ and 3 cases with one-sided p-values of 0.011 and 3.9×10^{-9} respectively, and there are no significant differences in bidder payoffs in the $n = 4$ and 5 comparisons across auction formats. The intuition for this result follows from the greater gains buyers can extract from playing the collusive strategy in hard-close auctions. As the number of bidders increases, the collusive gains diminish in hard-close auctions and it becomes more difficult to sustain the coordination on late bidding necessary for the collusive strategy to be profitable.

These results are fairly robust to many different parameter changes, conditional on the hard-close auction leading to significantly more frequent late bidding (the conditions for which were summarized above in detail) than in the soft-close auction. A few exceptions to that rule are found in the cases where $n_V = 10001$ and $n = 3$; $R = 40$ and $n = 3$; and $\epsilon = 50$ and $n = 3$. In these cases, there is significantly higher late bidding in the hard-close-format but the two auction formats are not significantly different from each other in terms of bidder payoffs and seller revenue.

We next consider whether there are differences in the amounts that bidders are bidding in hard- and soft-close auctions.

Result 4: Bidders *nearly always* learn to bid their full value in both soft- and hard-close auctions for all numbers of bidders, i.e., *bid shading is rarely observed.*

Support for this result can be found in Table 1 (and the subsequent Tables reporting on the sensitivity analysis), where the percentage of bidders who *attempt* to bid their full value in the auctions is given. Result 4 suggests that bidders are behaving rationally regardless of the auction closing rule in the sense that they bid their full valuation by the last period of the auction consistent with theoretical predictions for second-price auction formats. The findings of the sensitivity analysis support this result for a wide range of parameters. In all but three treatments, the percentages of bidders learning to bid full value is not significantly less than 98% using a one-sided t-test at the .05 level of significance.[24] The three treatments with low percentages of bidders bidding full value are all soft-close auctions: the $n = 2$ and $n = 3$ soft-close auctions when the extension periods allow incremental late bids and the $n = 3$ soft-close auction when $\rho = 0.8$. In these three treatments, only 77-78% of bidders submit late bids. These percentages are significantly lower than all other treatments at the 1% significance level with a one-sided t-test. In many treatments, the hard-close format results in larger percentages of full value bidding than does the comparable soft-close format, although the difference is only sometimes significant (as seen in Table 1).

Finally, we explore whether there are differences in the frequencies of phenotypes observed in hard- and soft-close auctions. Table 1 reports some cumulative, aggregate frequencies with which early-, late- and conditional-bidding automata are observed across hard- or soft-close auctions in our baseline case with $n = 2$, 3, 4, or 5 bidders.[25] Table 11 provides some further disaggregation for the baseline case – specifically the average frequencies of various "phenotypes" that exceed a small threshold, $\frac{1}{30}$ (i.e., 1 in every generation). We measure the *diversity* of the phenotypes according to the number of phenotypes occurring more frequently than $\frac{1}{30}$. The main finding from our analysis of these phenotypes is:

[24] We chose 98% as an *ad-hoc* baseline.

[25] This average is found by taking the average over last 100 generations in 20 simulations for each treatment.

Result 5: Strategies are more diverse in hard-close auctions than in soft-close auctions. Further, when n is small (e.g., $n = 2$ or 3) there is a large fraction of "conditional" phenotypes in hard-close auctions.

Support for Result 5 is found in Tables 1 and 11. As these tables reveal, in soft-close auctions, more than 90 percent of all evolving strategies are characterized as "E" (unconditional early bidding) phenotypes. In hard-close auctions, the tables reveal that the frequency of "E" phenotypes ranges from 15 to 97 percent depending on the number of bidders. In two-bidder hard-close auctions, the phenotype "L 1 2 E 2 1" is the one most commonly observed with a frequency of 42 percent. Strategies in this phenotype tell the bidder to bid late as long as the rival bidder also bids late. Otherwise, an early final bid is placed. If the rival only submits early bids, this strategy switches back to late bidding. Otherwise, early bidding is played again. The second most common phenotype in two bidder hard-close auctions is "E 1 2 L 2 1" with a frequency of 36 percent. This phenotype is almost identical to the earlier one except for the initial strategy which involves early bidding. In hard-close auctions with 3, 4 and 5 bidders, we observe the "E" phenotype emerge with increasing frequency from 54 percent when $n = 3$ up to 97 percent when $n = 5$. Nevertheless with 3 bidders other phenotypes such as "E 1 2 L 2 1", "E 2 2 L 2 2", "L", "L 1 2 E 2 1", "L 2 1 E 2 2", "L 1 2 E 1 1", and "L 2 1 E 1 2" are also observed with non-negligible frequencies. To better measure the diversity of phenotypes in the different auction environments, we calculate Simpson's diversity index for each case and report this index value at the bottom of Table 11. A value of 0 implies perfect diversity of phenotypes across the 22 possibilities while a value of 1 indicates no diversity.[26] The diversity index values in Table 11 indicate greater diversity of phenotypes in the hard-close as opposed to the soft-close auctions when $n < 5$.

Recall that the strategies with conditional phenotypes tell the bidder to bid early or late depending on the history of rival bids in the previous auction. We observe that conditional phenotypes are observed more frequently in hard-close auctions than in soft-close auctions. On the other hand, the frequency of unconditional early-bid phenotypes increases as the number of bidders increases in hard-close auctions. The intuition for this result is the same underlying Result 2; the gains from the collusive strategy decrease as the number of bidders increases. Still, the frequency of unconditional early bidding is always less in hard-close auctions than in soft-close auctions for 2, 3 and 4 bidders. In the sensitivity analysis reported on in Tables 2-9, Result 5 continues to hold conditional on the percentage of late bidding being greater in hard-close auctions than in soft-close auctions.

[26] Simpson's index of diversity is given by $D = \sum_i n_i(n_i - 1))/(N(N - 1))$, where n_i is the number of observations of phenotype i and N is the total number of phenotypes in the sample. This index can be viewed as the probability that two phenotypes, drawn at random from a sample, are the same. See, e.g., Miller (2003). We calculate this statistic using data from the last 100 generations of the 20 simulations of each treatment.

HARD-CLOSE	*Number of Bidders*			
Phenotypes	$n = 2$	$n = 3$	$n = 4$	$n = 5$
E	0.156	0.546	0.887	0.972
E 1 2 L 2 1	0.360	0.0831	0.0267	0.000433
E 2 2 L 2 2	0.00037	0.0742	0.0107	0.000450
L	0.00820	0.0532	0.00037	0.00118
L 1 2 E 1 1	0.021	0.0364	0.0117	0.000367
L 1 2 E 2 1	0.418	0.0455	0.000967	0.000333
L 1 2 E 2 2	0.00587	0.0347	0.0288	0.000983
L 2 1 E 2 2	0	0.0436	0.0000333	0.0000667
Remaining:	0.0306	0.0833	0.0337	0.0242
Diversity Index	0.33016	0.321469	0.788014	0.944513
SOFT-CLOSE	*Number of Bidders*			
Phenotypes	$n = 2$	$n = 3$	$n = 4$	$n = 5$
E	0.901	0.913	0.934	0.968
E 2 1 L 1 1	0.00095	0.0370	0.00675	0.0142
Remaining:	0.0981	0.0500	0.0593	0.0178
Diversity Index	0.81313	0.83549	0.87367	0.936566

Table 11: Frequency of surviving automata phenotypes in benchmark simulations (Averages over the last 100 generations surpassing a frequency of 1/30)

6 Simulations with Adaptive and Naive Incremental Bidders

Roth and Ockenfels' (2002) empirical findings suggest that there is a significant amount of "inexperienced" bidders participating in internet auctions, as revealed by the number of their feedback points. Such bidders often use an "naive incremental bidding stage-game strategy." This involves increasing the current auction price by bidding incrementally higher than the current second bid so long as the incremental bidder is not the current high bidder until the incremental bidder achieves high bidder status. We assume here that this incremental bidding strategy proceeds so long as the current price is lower than the incremental bidder's value. Such a strategy is a dominant strategy in an English auction. However, it is not dominant in hard-close, second-price internet auctions.

In this section, we introduce one such naive incremental bidder to each internet auction. These bidders do not evolve and should be viewed as "one–time bidders" indeed, they are replaced by a different naive bidder at the start of each auction. These naive bidders use a simple incremental bidding strategy: they only bid whenever they are not the current high bidder; in that event, they only bid the lowest fraction of their value, $\frac{1}{3}$, $\frac{2}{3}$, or 1, that is greater than or equal to the current price.

We re-ran all of our benchmark simulations replacing just one of the n bidders with a naive incremental bidder. Figure 2, which is comparable to Figure 1, shows the frequency over time with which the $n-1$ adaptive bidders submit late bids in hard- and soft- close auctions in this case where one bidder is a naive incremental bidder. Our main finding from this exercise may be summarized as follows:

Result 6: The addition of incremental bidders increases the percentage of late bidding in hard-close auctions relative to the environment without such bidders but leads to almost no change or less late bidding in soft-close auctions.

Support for this finding is found in comparisons of Figures 2 and 1, where we observe that the frequency of late bidding is greater in hard-close auctions with 1-3 adaptive bidders (we exclude the naive bidder) than without the incremental bidder. This percentage reaches and stays above 80 percent with 1 and 2 adaptive bidders, and 40 percent with 3 adaptive bidders. We formally test the differences between the benchmark simulations and simulations with naive bidders over the last 100 generations. For hard-close auctions, we observe that simluations with a naive bidder lead to significantly more late bidding by adaptive agents for $n = 2$ with a p-value of 9.8×10^{-16}, for $n = 3$ with a p-value of 4.55×10^{-15}, for $n = 4$ with a p-value of 9.08×10^{-9}, and for $n = 5$ with a p-value of 0.0026. On the other hand, for the soft-close auctions, the benchmark simulations (without the naive bidder) result in significantly greater late bidding by adaptive agents for $n = 3$ with a p-value of 0.0043 and for $n = 4$ with a p-value of 0.042. There is no significant difference between the two treatments when $n = 2$ with a two-sided p-value of 0.94 and when $n = 5$ with a two-sided p-value of 0.60. The intuition for this finding is that the presence of the naive incremental bidder provides a different, but nonetheless strong incentive for the adaptive bidder to delay bidding until the end: avoiding a bidding war. As the naive incremental bidders don't evolve in our set-up (alternatively there are many new incremental bidders arriving each period), these bidders never learn to bid late. Therefore the gains to the adaptive bidders from a late-bidding strategy may be even greater than in the case where all bidders are adaptive.

Ariely, Ockenfels and Roth (2005) report on a laboratory experiment with human subjects who play either hard- or soft- close auctions. In their experimental design, there are just 2 bidders in each auction and both play 18 auction games repeatedly. They model hard- and soft- close auctions differently than we do, but in their hard-close auctions, they do adopt a ρ value less than or equal to 1 as we do. They report that experimental subjects engage in significant late bidding in hard-close auctions and generally learn to bid early in soft-close auctions. One of the striking findings of the Ariely et al. study is that in hard-close auctions, the percentage of late bidding is higher when $\rho = 1$ than when $\rho < 1$. This finding is at odds with the tacit cooperation hypothesis that Ockenfels and Roth (2006) use to justify late bidding as an equilibrium strategy in hard-close auctions which requires that $\rho < 1$. Consequently, Ariely et al. pursue the hypothesis that late bidding is a best response to the presence of naive incremental bidders.

To perform a similar comparison with naive incremental bidders, we modify our benchmark simulation setup somewhat so that it is more closely, though not perfectly aligned with that of Ariely et al. (2005). Specifically, we run simulations with 1 adaptive and 1 naive bidder where the adaptive bidder updates his strategies after each internet auction game as in Ariely et al.'s laboratory study. That is, for comparison purposes, we set $R = 1$. We also consider two different values for ρ – 0.8 and 1 – the same values adopted by Ariely et al. (2005). All other parameters are as in our benchmark simulations. The findings from this simulation exercise are reported in Figure 3. We summarize the main finding as:

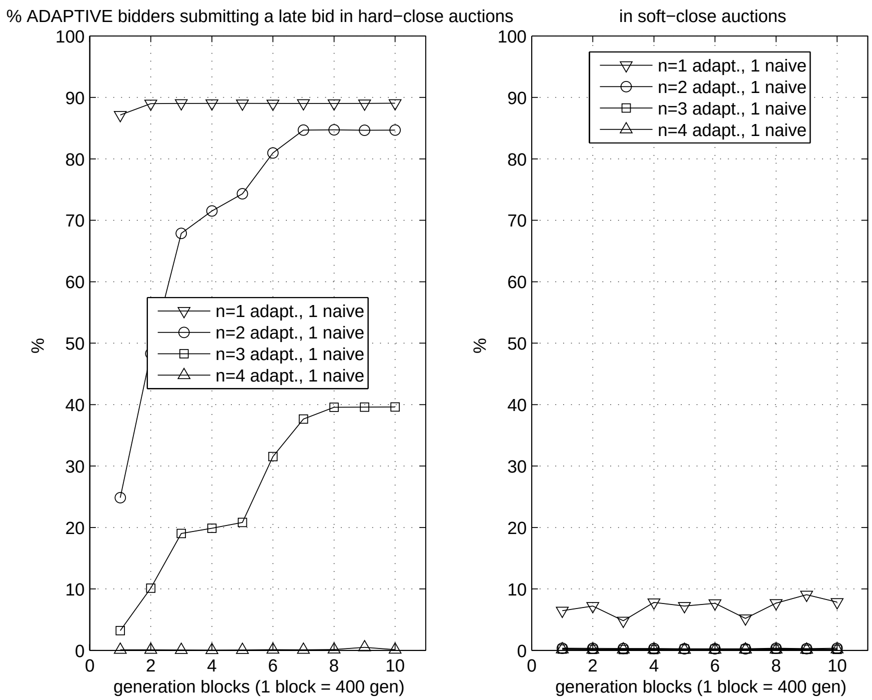

Figure 2: Simulations with 1 naive incremental bidder and 1, 2, 3 or 4 adaptive bidders. The average percentage of adaptive bidders successfully submitting a late bid in hard-close and soft-close auctions.

Result 7: In the presence of naive incremental bidding behavior, the percentage of late bidding in 2-bidder hard-close auctions is greater than in 2-bidder soft-close auctions irrespective of whether there are congestion effects ($\rho = .8$ or $\rho = 1$).

Support for this result is found in Figure 3a. The differences between the percentages of late bidding in hard- versus soft-close formats are large and significant. Like Ariely et al., we observe a somewhat higher frequency of late bidding by adaptive bidders in hard-close auctions when $\rho = 1$ (nearly 85 percent) than when $\rho = 0.8$ (nearly 80 percent). Also, bidders submit significantly more early bids per bidder in soft-close auctions than they do in the comparable hard-close auction (with the same ρ value) as seen in Figure 3b). We formally test these results using averages over the last 100 generations as follows: The difference in late bidding between hard-close and soft-close auctions with $n = 2$ is significant for $\rho = 0.8$ with a p-value of 2.45×10^{-16} and for $\rho = 1$ with a p-value of 1.18×10^{-23}. The difference in the number of bids per adaptive bidder between soft-close and hard-close auctions is significant with $n = 2$ for $\rho = 0.8$ with a p-value of 0.0048 and for $\rho = 1$ with a p-value of 1.62×10^{-17}. The intuition for this finding is simply that naive incremental bidders do not learn to bid their full valuation in the final period. Hence, *congestion effects are not necessary for there to be gains to following a late-bidding strategy in hard-close auctions.*

7 Conclusions and Implications for Market Design

We have developed a simple, agent-based model with the aim of understanding bidding behavior in internet auctions. Simulations of our model yield the main finding of significantly higher late bidding in hard-close auctions as compared with soft-close auctions. This result is consistent with the empirical evidence of significantly greater late bidding in hard-close as opposed to soft-close internet auctions as reported on by Roth and Ockenfels (2002) using field data and by Ariely, Ockenfels and Roth (2005) using experimental data. This external, empirical validation of our agent-based model findings gives us some degree of confidence that our model might serve as an aid in understanding other aspects of internet auctions. Furthermore, the results we report for our model appear to be robust to a wide variety of different parameterizations. Indeed, we are able to provide insight into parameterizations that will reliably generate our main results and we can explain why such parameterizations are necessary.

The intuition for our different findings between auction formats lies in the possibility that end-of-hard-close auction congestion ($\rho < 1$) provides potential gains to bidders if they all collusively adopt late-bidding strategies. Such gains are greatly reduced in soft-close actions where successful late bidding only serves to extend the duration of the auction, and so bidders have less incentive to delay bidding their full valuations. Whether agents in hard-close auctions take advantage of the potential gains or not (all bidding early) is an equilibrium selection question that we seek to address with our agent-based model.

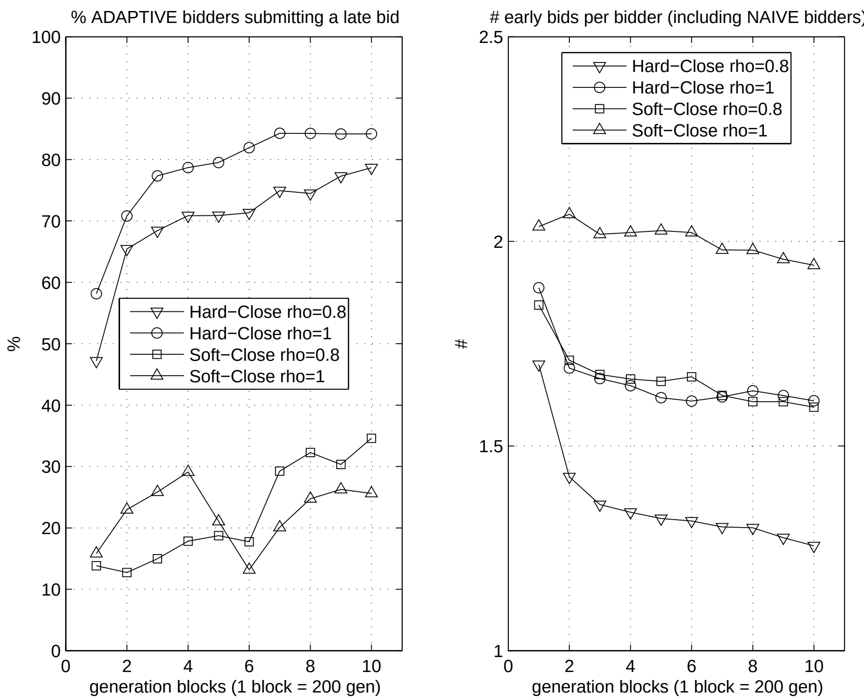

Figure 3: Simulations with 1 naive and 1 adaptive bidder with $R = 1$.The average percentage of adaptive bidders successfully submitting a late bid in hard- or soft-close auctions with $\rho = 0.8$ or 1; the number of early bids per bidder.

We find that strategies in hard-close auctions do evolve so as to take advantage of these potential gains from collusive late biding. However, the gains to late bidding in the hard-close auction should diminish as the number of bidders increases, and, indeed, the percentage of late bidding falls as n increases.

This finding suggests that hard-close auctions will result in *lower* revenue for sellers than soft-close auctions, with consequent gains to bidders, and this result also emerges from our simulations. These findings bear additional emphasis from a market design perspective.[27] Since internet auction web-sites view themselves as clearinghouses or intermediaries for market transactions, their interests are not clearly aligned with either sellers or buyers. However, information on *which* auction format favors sellers or buyers is of obvious use to such players, as they may be able to choose the auction format they participate in. Therefore, it may be natural to see both formats surviving side by side, as is currently the case.[28] The points raised by our study set the stage for further investigation on the evolution of different market designs for internet auctions.

In our benchmark simulations all bidders are "adaptive" learners but they eventually do learn to use "good" strategies, i.e., ones that have them bid their full valuations by the final period, T. The strategies themselves appear to be mainly unconditional early-bidding strategies in soft-close auctions, but there is greater diversity of both conditional- and unconditional-bidding strategies in hard-close auctions especially when n is small. The greater presence of conditional-bidding strategies in hard-close auctions is what underlies the greater frequency of strategic late bidding in those auctions.

With the addition of "naive" non-learning, incremental bidders (often observed in internet auctions), we find an even greater contrast in the frequency of late bidding by the adaptive bidders between the two auction formats, even in the absence of congestion effects and repeated game format. This finding is not so surprising; the presence of naive incremental bidders encourages the more sophisticated (but adaptive!) bidders to delay their bidding so as to increase their likelihood of achieving a higher surplus, and as the naive incremental bidders never learn to bid late, this strategy will often be profitable. In a soft-close auction, there are no gains to such a delay because any advantages to last-minute bidding (collusion or avoidance of incremental bidders) is removed.

Like every other modeling approach, our auction models lack some of the properties of the real-life auction formats. On the other hand, we modeled the features of the auctions that would effect our results as closely as possible. Also, we had to make modeling choices because of our simulation-based approach and other scientific reasons of secondary order of importance. For example, we used a discrete period format because of the limitations of the simulation-based approach. Another example is the extension periods of soft-close auctions. In reality extension periods are 10 minutes long in many soft-close auction sites. In the paper, we modeled this as a single period with a probability ρ of bid registration. We could

[27] It would be interesting to verify this prediction using field data from internet auctions, though this would require investigating auctions involving the same good under two different closing rules, and controlling for other factors including the number and demography of bidders, etc.

[28] Yahoo! auctions provide a third-way, allowing the seller to choose whether to have a hard- or soft-close to the auction.

have also modeled this as two periods, the last one being the final minute of the extension period and we could have adopted different bid registration probabilities for each period. These kinds of modeling choices only have secondary effects on our results and do not change fundamentally what we find in this paper. In order to have similar strategy formats in both hard- and soft-close auctions, we had to make these modeling choices.

Bidding in internet auctions is a particularly interesting topic for economists working on market design. Agent-based computational economics can be used as an important tool in testing alternative designs of market clearinghouses. As we show in this study, these techniques can successfully generate many of the empirical phenomena observed in real internet auctions and can therefore be used as a tool for effectively deciding which auction formats to adopt in applications or participate in as buyers or seller.

Appendix: Proof of the Theorem

We prove the Theorem separately for hard-close and soft-close auctions. Let the set of bidders be $N = \{1, 2\}$ and the increment be $\Delta = 0$. We will show that a strategy which involves bidding fraction 1 before period T weakly dominates any other strategy which does not involve bidding fraction 1 early in a stage auction.

1. First, we consider a stage hard-close auction. Let σ^1 be a stage game strategy of bidder 1 with the highest fraction $\alpha \in \left\{0, \frac{1}{3}, \frac{2}{3}\right\}$ submitted in one of the first $T-1$ periods and fraction $\alpha' \in \left\{0, \frac{1}{3}, \frac{2}{3}, 1\right\}$ such that $\alpha' \geq \alpha$ submitted in period T. Note that $\alpha' = \alpha$ means that bidder 1 does not submit a late bid in period T. Let σ^2 be a strategy of bidder 2 with the highest fraction $\beta \in \left\{0, \frac{1}{3}, \frac{2}{3}, 1\right\}$ submitted in one of the first $T-1$ periods and fraction $\beta' \in \left\{0, \frac{1}{3}, \frac{2}{3}, 1\right\}$ such that $\beta' \geq \beta$ submitted in period T. Let $\sigma = (\sigma^1, \sigma^2)$. Also consider a strategy of bidder 1 such that he bids fraction 1 in one of the first $T-1$ periods. Let σ'^1 be this strategy. Let $\sigma' = (\sigma'^1, \sigma^2)$.

   After the stage game under σ, av_1 will be the highest bid of bidder 1 for some $a \in \{\alpha, \alpha'\}$. Under σ', v_1 will be the highest bid of bidder 1. In both cases, it is equally likely that bidder 2 will have the highest bid, bv_2, for each $b \in \{\beta, \beta'\}$. We will consider 5 cases:

   (a) $v_1 \geq av_1 > bv_2$: Bidder 1 wins the auction under both σ and σ' with the ex-post payoff $v_1 - bv_2 - \Delta = v_1 - bv_2 > 0$.

   (b) $v_1 = av_1 = bv_2$: Bidder 1 may or may not win the auction under both σ and σ'. His payoff is 0 whether he wins or not.

   (c) $v_1 > av_1 = bv_2$: Bidder 1 may or may not win the auction under σ depending on the arrival time of his bid. If he wins, his payoff is $v_1 - bv_2 > 0$. Otherwise, his payoff is 0. Bidder 1 always wins under σ' with payoff $v_1 - bv_2 - \Delta = v_1 - bv_2 > 0$.

(d) $v_1 > bv_2 > av_1$: Bidder 1 does not win under σ'. His payoff is 0 in this case. He wins under σ' with payoff $v_1 - bv_2 - \Delta = v_1 - bv_2 > 0$.

(e) $bv_2 \geq v_1 > av_1$: Bidder 1 does not win under σ. He may win or lose under σ'. His payoff is 0 under both strategies.

We showed that for every highest bid submitted by bidder 2, it is a weakly best response for bidder 1 to use σ'^1 over σ^1.

2. Next, we consider a stage soft-close strategy. Let σ^1, σ^2 and σ'^1 be defined as above for the first T periods of the soft-close auction. In the extension periods, strategies σ^1 and σ^2 can involve bidding fraction 1 or 0 only. Let $\sigma = (\sigma^1, \sigma^2)$ and $\sigma' = (\sigma'^1, \sigma^2)$. Two cases are possible:

   (a) Under σ both agents do not bid in the extension periods, so the auction reduces to a hard-close auction. By the proof in part 1, σ'^1 weakly ex-post dominates σ^1.

   (b) Under σ bidder 1 or bidder 2 bids in the extension periods: After the stage game under σ, av_1 will be the highest bid of bidder 1 for some $a \in \{\alpha, \alpha', 1\}$. Under σ', v_1 will be the highest bid of bidder 1. Under both cases, bidder 2 will have the highest bid bv_2 for some $b \in \{\beta, \beta', 1\}$. Cases (a) to (e) outlined in the first part of the proof still hold. However, the events are not equally likely to occur under σ and σ'. If we can show that bidder 1's registered highest bid is more likely to be higher under σ' and bidder 2's registered highest bid is more likely to be lower under σ', then the proof will be complete.

   Bidder 1's highest fraction can be α, α' or 1 under σ. His highest bid is fraction 1 under σ'. Therefore, the probability distribution of bidder 1's highest bid under σ' weakly first-order stochastically dominates the distribution of bidder 1's highest bid under σ.

   Bidder 2's highest can be fraction β, β' or 1 under both σ and σ'. His behavior can be observed under three cases:

      i. Bidder 2 does not bid in period T under σ^2: then his highest bid will be fraction β or 1 under σ, since bidder 1 can cause an extension of bidding and bidder 2 can bid in that extension period. On the other hand, there will be no extension period under σ'. Hence, bidder 2's highest bid will be a fraction β under σ'.

      ii. Bidder 2 bids in period T but he does not bid in the extension periods: then his highest registered bid will be fraction β or fraction β' with the same probability under σ and σ'.

      iii. Bidder 2 bids in period T and in the extension periods: then the probability of having an extension period under σ is no smaller than the same probability under σ' since, bidder 1 may be bidding in period T under σ. Bidder 2's fraction 1 registers with no smaller probability under σ than under σ' in the extension periods. On the other hand, bidder 2's highest bid will be fraction β with no larger probability under σ than under σ'. This is true, because more extension periods under σ provide more opportunities for bidder 2 to increase his bid over fraction β.

Cases (i) to (iii) imply that the probability distribution of bidder 2's highest bid under σ weakly first-order stochastically dominates the same distribution under σ'. Recall that the cumulative distribution of bidder 1's highest bid under σ' weakly first-order stochastically dominates the same distribution under σ. Hence, strategy σ'^1 weakly ex-post dominates σ^1.

References

[1] Abreu, D. and A. Rubinstein (1988), "The Structure of Nash Equilibrium in Repeated Games with Finite Automata," *Econometrica*, 56: 1259-1282.

[2] Andreoni, J. and J. H. Miller (1995), "Auctions with Artificial Adaptive Agents," *Games and Economic Behavior*, 10: 39-64.

[3] Ariely, D., A. Ockenfels and A. E. Roth (2005), "An Experimental Analysis of Late-Bidding in Internet Auctions," *RAND Journal of Economics,* 36: 890-907.

[4] Arifovic, J. (2000), "Evolutionary Algorithms in Macroeconomic Models" *Macroeconomic Dynamics,* 4: 373-414

[5] Avery, C. (1998), "Strategic Jump Bidding in English Auctions," *Review of Economic Studies*, 65: 185-210.

[6] Bajari, P. and A. Hortaçsu (2003), "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34: 329-355.

[7] Bajari, P. and A. Hortaçsu (2004), "Economic Insights from Internet Auctions," *Journal of Economic Literature,* 42: 457–486.

[8] Dawid, H. (1999), *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*, Second Edition. Berlin: Springer.

[9] Engle-Warnick, J. and R. Slonim (2006), "Inferring Repeated Game Strategies from Actions: Evidence from Trust Game Experiments," *Economic Theory*, 28: 603-632.

[10] Epstein, J. M. (2006), *Generative Social Science*. Princeton, NJ: Princeton University Press.

[11] Goldberg, D. E. (1989), *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA: Addison-Wesley.

[12] Gonzalez, R., K. Hasker and R. C. Sickles (2004), "An Analysis of Strategic Behavior in eBay Auctions," working paper.

[13] Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems*. Ann-Arbor, MI: The University of Michigan Press.

[14] Lucking-Reiley, D. (1999), "Using Field Experiments to Test Equivalence between Auction Formats: Magic on the Internet," *American Economic Review*, 89: 1063-1080.

[15] Michalewicz, Z. (1994), *Genetic Algorithms + Data Structures = Evolution Programs*, Second Edition. Berlin: Springer-Verlag.

[16] Miller, C. B. (2003), *Biological Oceanography*. Oxford: Blackwell Publishing.

[17] Miller, J. H. (1996), "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *Journal of Economic Behavior and Organization*, 29: 87-112.

[18] Moore, E. (1956), "Gedanken Experiments on Sequential Machines," *Automata Studies*. Princeton, NJ: Princeton University Press.

[19] Ockenfels, A and A. E. Roth (2006), "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction," *Games and Economic Behavior*, 55: 297-320.

[20] Roth, A. E. and A. Ockenfels (2002), "Last Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92: 1093-1103.

[21] Tesfatsion, L. and K. L. Judd (2006), eds., *Handbook of Computational Economics*, *Volume 2 - Agent-Based Computational Economics*. Amsterdam: Elsevier.

[22] Vickrey, W. (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16: 8-37.