

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

John Duffy[†]

University of Pittsburgh

M. Utku Ünver[‡]

University of Pittsburgh

October 7, 2005

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. As an interesting result, 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 Alexander Elbittar and Jeff McKie-Mason for 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. However, this hybrid format is strategically different from the sealed-bid second-price auctions and ascending-bid English auctions.² Indeed, the advent of the internet auction has led to a new theoretical and empirical literature devoted to exploring this new auction format.³

There are two popular ascending-bid, second-price formats used by internet auction sites. The first of these, a “hard-close” auction, closes at the end of a fixed preset time period, typically one or two weeks. The high bidder wins the object by paying the second highest bid plus some small increment.⁴ The second type, a “soft-close” auction, closes at the end of a fixed duration if and only if no bidder submits a “late” bid within a certain interval of time near the closing time (e.g., 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.⁵

An interesting phenomenon, known to participants in hard-close internet auction sites and empirically documented by Roth and Ockenfels (2002) is that of “last-minute” or “late” bidding, which practitioners call “sniping.” Specifically, more bids are submitted 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. Ockenfels and Roth (2002) present a model that can rationalize late bidding as an equilibrium strategy in hard-close auctions under both private-value and common-value auctions.⁶ However, late-bidding is just one equilibrium possibility; all bidders bidding early remains another.

In modelling late-bidding as an equilibrium phenomenon, Ockenfels and Roth rely on the assumption that there is some probability that bids submitted in the final period of the hard-close auction will fail to be properly transmitted to the auction software, due either to internet congestion or to human-related factors such as high monitoring costs. Bidders in hard-close auctions who adopt a mutual, late-bidding

¹For instance, eBay Inc. reports steady growth in its annual gross merchandise volume – the total value of all successfully closed listings on eBay’s trading platforms (primarily auction listings) – from \$95 million in 1997 to \$34.2 billion in 2004.

²See Vickrey (1961) for second-price auctions and English auctions. See Milgrom and Weber (1982) and McAfee and McMillan (1987) for symmetric auctions.

³See Bajari and Hortaçsu (2004) for a survey of the economics literature on the subject.

⁴This format was originally implemented by eBay.

⁵This auction format was originally implemented by Amazon. It has a “going-going-gone...” feature as in English auctions.

⁶The 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.

strategy may therefore occasionally win items at low prices. That is, the end-of-auction congestion creates a potentially large ex-ante surplus for them, and the potential to capture this surplus is what rationalizes their late bidding strategy.⁷

In this paper, we develop a simple model of internet auction behavior with the aim of understanding the evolution of bidding behavior in hard- and soft- close auctions with various numbers of bidders. Our model involves a single seller offering an item for which bidders have independent, private-values. The bidders are the only active players. The bidders play their strategies against each other repeatedly in either hard-close or soft-close multi-period (dynamic) auctions. They have a “selective” message space concerning the history of previous auctions: specifically, they care only about the timing of rival bids in the most recently completed auction. Bidder strategies specify the amount and timing of a bidder’s bids within a dynamic auction and these strategies may or may not condition on the history of rival bids in the most recently completed auction. While our model of bidding behavior is quite simple, it is flexible enough to allow early or late bidding as well as history contingent or unconditional bidding behavior.

As there are multiple bidding strategies that can comprise equilibria in repeated auctions (e.g. all bid early or all bid late as shown by Ockenfels and Roth (2002)), we adopt an agent-based computational approach, with the aim of finding optimal bidding strategies for a given, repeated-auction environment. Specifically we analyze the evolution of bidding behavior by n bidders under hard- and soft- close auctions using a “genetic algorithm” (Holland, 1975), which is a versatile search and optimization tool for large strategy spaces. Genetic algorithms optimize on the efficient boundary between exploiting strategies that have worked well in the past and exploring new strategies (Goldberg, 1989). Alternatively, one can think of a genetic algorithm as a model of social learning (Dawid, 1999) or as a macroeconomy with heterogenous agents (Arifovic, 2000). In a genetic algorithm, the better strategies (as measured by payoffs) of the current generation of players are copied and/or modified and then transmitted for use by future generations.

We report the results of several simulation exercises using artificial bidders who use strategies updated by a genetic algorithm. The format of the strategies and the method by which these strategies evolve is the same in both the hard- and soft-close auctions. Nevertheless, in all of our simulations we consistently find that hard-close auctions lead to much more frequent use of late bidding strategies than do soft-close auctions, even as we vary certain parameters, e.g., the number of bidders is varied from 2 to 5. Further,

⁷Ockenfels and Roth (2002) also claim that late bidding can be caused by the presence of naive incremental bidders. As evidence of this, they show evidence of 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 begin auctioned. Empirically, Hasker, Gonzalez and Sickles (2004) show that the type of late bidding equilibrium introduced by Ockenfels and Roth (2002) 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. A recent experimental paper by Ariely, Ockenfels and Roth (2005) confirms that hard-close auctions are prone to late bidding while soft-close auctions are not. In an interesting study, Borle, Boatwright and Kadane (2005) classifies e-Bay auctions on a scale from private-value to common-values using data that explains whether there was late and multiple bidding in a given auction. There are other empirical papers which explain various aspects of internet auctions. An incomplete list includes Lucking-Reiley (1999, 2000), Melnik and Alm (2001).

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. We perform some sensitivity analysis and show how some of our findings may be dependent on the choice of model parameters such as those governing the distribution of private valuations or the number of bidders. By way of an explanation for our findings, we present evidence of a greater variety of bidding strategies in hard-close versus soft-close auctions; in particular, the use of history contingent bidding strategies is much more common in hard-close auctions than in soft-close auctions. Finally, we inspect the evolution of adaptive bidding strategies in the presence of naive incremental bidding agents. The existence of such incremental bidders is consistent with behavior observed empirically, 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 found in simulations without these incremental bidders. We conclude that agent-based computational economics can be used as a tool for market designers interested in predicting outcomes under various auction formats.

2 Internet Auction Rules

In this section, we explain the general rules of the hard- and soft-close internet auction that we formally introduce in the next section as an auction model. We start with hard-close auctions. A hard-close auction is a dynamic auction. Each seller can post a single item for sale for a certain fixed amount of time on the internet auction website. Prior to the end of the auction (the hard-close), any bidder can submit a bid at any time. At any moment during the auction, the “current bid” or price of the object for sale is defined to be the second highest bid submitted thus far (if any) plus a small increment. Setting a reserve price is also an option. Any new bid lower than the current bid is considered an invalid bid.

A valid internet auction bid of some amount b in excess of the “current bid” or price is called a “proxy bid”. It is more properly viewed as a proxy bidding *rule* that automatically increments the high bidder’s bid as new bids come in to challenge that bid (plus a small increment), enabling a high bidder to retain his high bid position, so long as the new bid amount necessary for the high bidder to retain high bid status does not exceed the maximum amount bid, b (and the owner of the old high bid is a different bidder than the owner of b). Otherwise, a bidder who has bid b is outbid and a new high bidder takes his place (with the same proxy bidding rule working for that bidder). When the predetermined time of the auction runs out, the current high bidder wins the auction at the current bid price.

The rules of soft-close auctions are similar, but with one important difference. Every valid proxy bid submitted within t minutes of the scheduled end of the auction causes that auction to be extended t more minutes starting from this last bid. The auction is concluded t minutes after the last valid proxy bid

is submitted. If a new valid proxy bid is submitted during this time, the auction is extended t more minutes, and so on. When the auction is finally concluded, the high bidder wins the object at the current bid price.

In the next section, we will extensively discuss the formal model we use and how it relates to the real-time implementation of bids in these auctions.

3 The Internet Auction Model

In order to tractably analyze bidding behavior in internet auctions consistent with the description given above, we consider a highly simplified model of an internet auction but one that captures all of the essential features of these auctions. In particular, suppose the internet auction involves a single seller, who is offering an indivisible object without any reserve price. Let $N = \{1, 2, \dots, n\}$ be the set of bidders, each of whom is risk neutral in money. We will sometimes refer to the bidders as “agents,” or “strategies,” since in our analysis bidders are the only active players. Each agent has a valuation for the object; this valuation serves as the agent’s “type”. Each agent’s valuation 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. Each bidder knows his own type and the probability density function g used to determine other agents’ types. 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 $m > \frac{\epsilon}{2}$. Specifically:

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

The internet auction is a “standard auction” in which the high bidder wins the object by paying a fee to the seller. No other bidder pays a fee to the seller and the seller subsidizes no bidders.

We model the internet auction as a dynamic Bayesian game in which there are T consecutive bidding periods, indexed by $t = 1, 2, \dots, T$. The discrete choice for bidding periods is necessary for our agent-based model and is also consistent with the design of complementary experimental studies (discussed later in the paper). In each period t , each bidder can submit “at most” one bid over the internet to a software program. A bid submitted in one of the first $T - 1$ bidding periods is always registered by the software program correctly. However, a bid that is submitted in the final period T is correctly registered by the software program with some fixed probability $\rho \leq 1$. The latter assumption captures congestion effects or bidder timing mistakes in the scheduled end period of an auction.

The dynamic internet auction rules are as follows. When there are no bids submitted, the smallest admissible bid is equal to some fixed and known increment $\Delta > 0$. When the first admissible bid not smaller than Δ arrives, the second bid is set to Δ . The owner of this first admissible bid becomes the “high bidder.”

Consider any bidding period. Let bid b^* be the current high bid, bid b_2 be the current second bid, and bidder i^* be the current high bidder. Bidders only observe the second bid b_2 and the identity of the high bidder i^* . All bidders simultaneously submit their bids in this bidding period. However, the arrival order of these bids to the computer program (internet website) is not simultaneous. The arrival order of bids is determined by a randomly drawn permutation of all n bidders. Each permutation is equally likely to occur.

Suppose bid b is the first correctly registered bid in period t and bidder i be the owner of bid b . Three cases are possible:

1. $b > b^*$: In this case, we compare the identities of bidder i and the high bidder i^* . There are two possibilities:
 - (a) Bidder i is different from bidder i^* . In this case,
 - i. bid b becomes the new high bid,
 - ii. bid b^* becomes the new second bid, and
 - iii. bidder i becomes the new high bidder.
 - (b) Bidder i is the same agent as bidder i^* . In this case, bid b becomes the new high bid. The second bid and the identity of the high bidder do not change.
2. $b^* \geq b \geq b_2 + \Delta$: Bid b becomes the new second bid. The high bid and the identity of the high bidder do not change.
3. $b_2 + \Delta > b$: The high bid, the second bid, and the identity of the high bidder do not change. In this case, bid b is regarded as an invalid bid; it does not get counted in statistics reporting on submitted bids.

Here cases 1 and 2 reflect the proxy feature of the submitted bids. In case (1a), an implicit bidding war among proxy bids b and b^* determines the new high bidder as the owner of b at the price $b^* + \Delta$. In case (1b), the new bid b only reflects an update of the high bidder regarding his older bid without affecting the second bid (current “price”). In case (2), the proxy-bid b^* responds to the proxy-bid b by posting a new bid that is incrementally higher than b .

After bid b , the second arriving bid is processed in a similar fashion using the updated information concerning the high bid, the second bid, and the high bidder. Other bids are processed in a similar manner.

We consider two ending procedures for the internet auction. In a “hard-close auction,” the auction closes after the last bid of period T is processed. The high bidder at the end of period T wins the object by paying the “price.” The price is equal to the second bid plus increment Δ if that is not greater than the high bid. Otherwise, the price is equal to the high bid.

In a “soft-close auction,” the auction closes after period T if and only if a valid bid does not register in period T . Otherwise, the auction is extended one more period where the agents simultaneously submit their bids. The auction closes if and only if a valid bid does not register in the extension period. Otherwise, the auction is extended one more period. After one of the extension periods, the auction ends if and only if no new valid bid registers. The rules for bid registration in any extension period $t > T$ are the same as for period T . Hence, a submitted bid may fail to register with probability $1 - \rho$. The high bidder at the end of the soft-close auction wins the object by paying the current bid price. Note that, by design, a finite end is guaranteed to the soft-close auction by the finite number of agent types; one agent will have the highest valuation, and bidding cannot exceed this value.

An internet auction strategy σ^i for bidder i is a list of bids at each period $t \in \{1, 2, \dots\}$ for every type of bidder i and for every possible history of: (i) bidder i 's bids in periods $1, \dots, t - 1$, (ii) the identity of the high bidders in every period $1, \dots, t - 1$, and (iii) the second bids in every period $1, \dots, t - 1$.

3.1 Repeated Auctions and Modelling Repeated Game Strategies

Each agent plays a block of R consecutive dynamic, internet auctions. Each auction in a block is also referred to as a “round” or “stage game”. The type (valuation) of each agent is redetermined via another draw from g at the beginning of each round. The utility of agent i is the discounted summation of stage game utilities. Let $\delta \in (0, 1]$ be the discount factor. After each auction, bidders can observe the “bid history” of the previous auction. A bid history is a list which shows the bids of each agent at each period of the auction. Let $r \in \{1, 2, \dots, R\}$ index rounds of the repeated auctions. A “history” h_r is defined as a list of bid histories in rounds $1, \dots, r - 1$. History h_1 is the empty set. A repeated game strategy s^i of bidder i is a list of stage game strategies $\sigma^i | h_{r'}$ for each possible history $h_{r'}$ for each round $r' \in \{1, 2, 3, \dots\}$.

This strategy representation is obviously quite complicated. For modelling purposes, it will therefore be useful to limit the space of admissible strategies. We make three main assumptions on the admissible strategy space:

1. A stage game strategy for agent i depends only on the previous bids of agent i and his type. Further, we permit just four possible bids for each bidder in a bidding period. The bid of bidder i with type v_i in any period 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.⁸
2. In extension periods of a soft-close auction $t > T$, we permit just two possible bids: v_i or 0. Further, bidder i 's bid is restricted to be the same in every extension period.

⁸We do not allow bids greater than the bidder's private value simply because bidding above one's private value is always 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.

3. In the repeated auction game, history h_r is “subjective” and “selective” for each bidder. Specifically, each bidder i ’s history, h_r^i is characterized by the timing of the last round ($r - 1$) bids of his rivals, and only two possible histories are permitted. History $h_r^i = \textit{late}$ denotes the state where, in round $r - 1$, a rival bid arrived in period $t \geq T$. History $h_r^i = \textit{early}$ denotes the state where, in round $r - 1$, no rival bid arrived in any period $t \geq T$.

This simplified model captures the essential features and rules of internet auctions, but some of our simplifications have important consequences. Assumption (1), for instance, eliminates incremental bidding 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 such 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 (2) is just a simplification that eliminates cumbersome bidding strategies in the extension periods.⁹ Assumption (3) eliminates the price dimension from the history of previous auctions and focuses on the timing issues that are the focus of our analysis. 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.

In the next subsection, we illustrate the admissible strategies in great detail.

3.2 Finite Automata as Repeated Game Strategies

The repeated auction strategy of each bidder i , s^i , is approximated by a “finite automata” representation (Moore, 1956).¹⁰ An “automaton” is a string of integers that describe a stage game strategy and the next move of the bidder for each history of the game.

In our implementation, there are just two histories, “early” and “late,” upon which bidders can condition their bidding strategy. To allow for history-contingent, “repeated game” strategies, we therefore allow *two* automata for each strategy, one for each possible history.¹¹ Each automaton consists of two parts: (i) the “stage game” (internet auction) strategy and (ii) two indexes. These two indexes determine which

⁹This 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.

¹⁰Finite 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 prisoners’ dilemma. See Engle-Warnick and Slonim (2005) for an application to inference of human strategies from experimental data.

¹¹This is the minimal number of automata necessary for implementation of repeated game strategies. Nevertheless, we believe that our implementation is sufficiently general to characterize a wide range of possible bidding strategies in repeated internet auctions.

automaton will be used after the current stage game strategy in each possible history. A representation of the two automata that comprise the strategy of bidder i , s^i , is given below.

$$s^i = \left(\overbrace{\underbrace{\sigma_1^i}_{\text{1st stage-auction strategy}}, \underbrace{[L_1^i, E_1^i]}_{\text{index of next automaton}}}^{\text{1st automaton}} ; \overbrace{\underbrace{\sigma_2^i}_{\text{2nd stage-auction strategy}}, \underbrace{[L_2^i, E_2^i]}_{\text{index of next automaton}}}^{\text{2nd automaton}} \right)$$

Consider automaton $j \in \{1, 2\}$. A stage game strategy σ_j^i is a list of 3 integers for a hard-close auction or 4 integers for a soft-close auction. (The particular ending-rule of the auction is assumed to be known in advance). The automaton σ_j^i that is part of bidder i 's strategy in a hard-close auction is illustrated below:

Stage game strategy of a hard-close auction:

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

Integer σ_{jk}^i for every $k \in \{1, 2, 3\}$ is a period number, that is $\sigma_{jk}^i \in \{0, 1, 2, \dots, T\}$. 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 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.¹²

The automaton σ_j^i that is part of bidder i 's strategy in a soft-close auction is illustrated below:

Stage game strategy of a soft-close auction:

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

Integer σ_{jk}^i for every $k \in \{1, 2, 3\}$ is a period number, that is $\sigma_{jk}^i \in \{0, 1, 2, \dots, T\}$. The first three integers have identical roles to their roles in a hard-close stage auction strategy. The fourth integer, σ_{j4}^i , is binary, i.e., it lies in $\{0, 1\}$. When $\sigma_{j4}^i = 1$, bidder i will bid his entire value in every extension period. When $\sigma_{j4}^i = 0$, bidder i will not bid anything in any extension period.

¹²If 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.

The next two integers, L_j^i and E_j^i in either the hard- or soft-close automaton are indexes in $\{1, 2\}$, as illustrated below:

$$\left[\begin{array}{c} \text{index of the next stage-auction strategy} \\ \overbrace{L_j^i \quad , \quad E_j^i} \\ \underbrace{\hspace{1.5cm}} \\ \text{when state is late} \quad \text{when state is early} \end{array} \right]$$

These indicate which automaton will be selected for the next auction (stage game). When the subjective history is late, automaton L_j^i will be selected next. In this case, bidder i will play stage game strategy $\sigma_{L_j^i}^i$. When the subjective history is early, bidder i will select automaton E_j^i and play stage game strategy $\sigma_{E_j^i}^i$. In the very first round of the repeated internet auction, we assume that bidder i always selects the first automaton in his strategy s^i . Following the first round, the transitions between the two automata are dictated by the subjective history and transition index values of the automata.

To help clarify the finite automata representation of bidder strategies, we give an example below:


Example 1: Let $T = 8$. Consider the following soft-close auction strategy for bidder i , s^i :

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

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 2nd digit) and v_i in every extension period (as indicated by the 1 in the 4th digit). Notice that the strategy also tells the bidder to bid $\frac{1}{3}v_i$ in period 3 (the 3 in the 3rd 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 first stage game strategy call for bidder i to repeat use of this first stage game strategy in the next auction (as indicated by the 1 in the 5th digit). Otherwise, if no rival submits a bid in period 8 or an extension period, the history is early, and the second stage game strategy is used in the next auction (as indicated by the 2 in the 6th digit).

The second stage game strategy, which starts at the 7th digit, tells bidder i to bid $\frac{2}{3}v_i$ in period 4 (as indicated by the 8th digit) and to bid v_i in period 8 (as indicated by the 7th digit). If the auction is extended, no bid will be submitted (as indicated by the 10th digit). If a rival bid arrives in or later than period 8, then the history is “late” and this same strategy will be played again (as indicated by by the 11th digit), otherwise the history is “early” and the first stage game strategy will be played next (as indicated by the 12th digit).

The example strategy presented above can also be characterized using an alternative finite-state representation known as a Moore machine, as illustrated below:



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

4 Theoretical Analysis

Before proceeding to a simulation analysis of bidding behavior in the internet auction it will be useful to establish a theoretical result regarding the timing of bidding behavior in a single play of the internet auction. 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: 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 ex-post (and strictly ex-ante) 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 $\Delta = 0$.¹³

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

¹³See Appendix A for the proof.

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 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.

5 Simulations with Artificial Adaptive Agents

5.1 Some motivation

We adopt an adaptive approach in our analysis. Since adaptation requires repetition, we allow our artificial adaptive agents to gain experience by playing a repeated-game version of the dynamic internet auction. These artificial agents initially use randomly generated strategies. They adaptively learn to experiment and make use of better strategies over time in a trial-and-error learning process. Specifically, we use a model of adaptation known as “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 and have been shown to optimize on the trade-off between exploring new strategies and exploiting strategies that have performed well in the past. (Holland, 1975). 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. Andreoni and Miller did not consider dynamic auctions with the bidding rules of internet auctions as such auction formats were only invented after their paper was published. Still, our approach has much of the flavor of their study: the optimal bidding strategies in repeated auction formats are difficult to characterize analytically and so a numerical search is a reasonable and promising alternative. That is, an algorithm that has been shown to find optimal solutions in highly complex environments—the genetic algorithm—would seem to be an excellent candidate for characterizing bidding strategies in repeated dynamic internet auctions.

5.2 Algorithmic details

We suppose there is a population of strategies (finite automata) of the type described above. The size of this population is fixed at N . The automata in this population are initially generated randomly subject to constraints on integer values, e.g. 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 N 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 repeated internet auctions. Specifically, the genetic operators of the genetic algorithm are called on after a fixed number of blocks (a “tournament”) has been played.

Each block proceeds as follows. First, a set of n finite automata (bidders) are randomly chosen from the N -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 consecutive dynamic, internet auctions, each lasting T 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, taking into account 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 strategies.

1. Selection: Some number $M < N$ of the best strategies of the current generation, as determined by fitness levels, are reproduced (copied intact), to be included among the set of N "offspring" strategies that comprise the next generation of N strategies. The remaining $N - M$ next generation, offspring strategies are obtained using the crossover operation described next.
2. Crossover: Parts of the better strategies of the current generation are recombined to form the remaining members of the next generation of strategies. There are various crossover operators used in the literature. We adopt the "linear crossover" operator, described in the following steps.
 - (a) Two parent strategies are selected randomly 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 $\ell \in \{1, \dots, L - 1\}$ 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) The crossover operation is repeated until there are N new offspring strategies for the next generation, the M strategies obtained via selection and the $N - M$ strategies obtained via crossover or recombination.

3. Mutation: The mutation operation applies to all of the N offspring strategies created via selection and crossover. Specifically, each integer of each offspring strategy is randomly changed to another admissible integer value with a small fixed probability.

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 a number of generations. Further, we run a number of simulations for the same treatment (number of bidders n ; hard- or soft- close auction format) with different random seed values to obtain Monte Carlo estimates of different statistics. A pseudo-code description of our algorithm, including our specific parametric choices, is given below:

For $i = 1$ to 20 (the total number of simulations)

1. Randomly generate $N = 30$ (size of a generation) strategies for the initial strategy pool.

2. For $G = 1$ to 4000 (total number of generations)

(a) Conduct a tournament consisting of 300 repeated game blocks (internet auctions).

i. For each of these repeated game blocks, randomly match n bidders (strategies) from the strategy pool consisting of $N = 30$ strategies.

(b) The fitness of each strategy is the average of the payoffs that strategy earns in all the auctions it is selected to participate in during the tournament.

(c) For $k = 1$ to 6, select the highest fitness strategy not selected yet as an offspring to be included in the next generation of strategies.

(d) For $k = 7$ to 30, cross strategies over to generate additional new offspring.

i. Choose two “parent” strategies using a “biased-random-wheel” selection: The “propensity” of a strategy is defined as the fitness of that strategy minus three fourths of the smallest fitness in the generation. The probability of choosing a strategy as a parent is its propensity over the sum of the propensities of all strategies in the generation.

ii. Cross the two parents over linearly to generate two offspring strategies.

- (e) Mutate each digit in the offspring strategies with probability 0.01.
 - i. Use a non-uniform unimodal probability distribution for mutation around the value of the digit.
- (f) Form the new strategy pool using the offspring strategies.

As noted above, we conduct benchmark simulations with either $n = 2, 3, 4$ and 5 bidders, and under either a hard— or soft—close format.¹⁴ In these simulations, we set the expected value of the probability density function g used to draw valuations, at $m = 10^6$, and the spread of the support interval, $\epsilon = 40$. We choose the number of mass points in the interval $[10^6 - 20, 10^6 + 20]$ as $n_V = 6$.¹⁵ In the sensitivity analysis reported on in Appendix B, we consider other mean values. 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 set the discount factor $\delta = 1$. We set $\rho = 0.9$ as the probability of bid registration in each period $t \geq 8$. Finally, we set the bid increment, $\Delta = 1$. We also conduct supporting simulations with different model parameters. In these sensitivity analyses, described in Appendix B, we change one parameter at a time, and then compare new Monte Carlo estimates with the original ones.

In the next subsection, we introduce a method for classifying strategies that aids in our presentation of the simulation findings. We use this method to give summary statistics on the evolving strategies.

5.3 Classification of strategies

We introduce a simple, strategy classification method for use in interpreting the evolving strategies in our simulation exercises. Each category in this classification is called a “phenotype.”¹⁶ The phenotype of a repeated game strategy (or bidder) 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 normal duration of the game, i.e. within 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.

¹⁴The 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).

¹⁵The mean value $m = 10^6$ is a pedagogical value only; it could represent any number of units of money, e.g., cents.

¹⁶This term is inspired by a similar term for classifying genes in biological evolution.

We can perhaps best describe the idea of a phenotype via some examples. Phenotype “E 2 1 L 1 1” characterizes a strategy where, in the initial, first strategy (automaton) the bidder’s final bid (any fraction of his value) is made early, in some period before period T ; hence the “E” in the first position. His second strategy (automaton) involves placing a bid (of any fraction) in the final period T and is therefore labeled as “L” 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 all rival bids also arrive early, so that the history is E, the integer 1 in position 3 of this phenotype indicates a return to the first strategy of early bidding. If any rival bid arrives late so that the history is L, the integer 2 in position 2 indicates that this bidder will move to the late bidding strategy 2 for one auction and will then always switch back to the early bidding strategy 1, regardless of the history of play in the auction where he uses the late bidding strategy; hence the integers 1 in positions 5 and 6. Phenotype “E 2 1 L 2 1” characterizes a strategy where the bidder initially bids early and later imitates the timing of his rivals’ final bids in the first T periods. This phenotype characterizes “tit-for-tat” strategies. Phenotype “L 1 2 E 1 2” characterizes strategies where bidding is initially late but moves to imitation of the timing of rivals’ final bids. This is another kind of tit-for-tat strategy. We note that there are just 22 such 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 strategies into phenotypes is 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 $((\mathbf{0}, \mathbf{8}, \mathbf{2}, \mathbf{0}), [2, 1]; (\mathbf{8}, \mathbf{1}, \mathbf{2}, \mathbf{0}), [1, 1])$

hard-close auction strategy $((\mathbf{8}, \mathbf{1}, \mathbf{5}), [1, 2]; (\mathbf{8}, \mathbf{0}, \mathbf{0}), [2, 1])$

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 $((\mathbf{0}, \mathbf{4}, \mathbf{2}), [2, 1]; (\mathbf{5}, \mathbf{8}, \mathbf{2}), [1, 2])$

soft-close auction strategy $((\mathbf{7}, \mathbf{6}, \mathbf{5}, \mathbf{1}), [2, 2]; (\mathbf{5}, \mathbf{8}, \mathbf{7}, \mathbf{0}), [1, 1])$

Example 5: The following strategy is a grim-trigger strategy in a hard-close auction with phenotype “L 1 2 E 2 2”:

$$((\mathbf{8}, \mathbf{6}, \mathbf{1}), [1, 2]; (\mathbf{2}, \mathbf{8}, \mathbf{0}), [2, 2])$$

Example 6: The following is a tit-for-tat strategy in a soft-close auction with phenotype “L 1 2 E 1 2”:

$$((\mathbf{0}, \mathbf{8}, \mathbf{4}, \mathbf{1}), [1, 2]; (\mathbf{2}, \mathbf{8}, \mathbf{5}, \mathbf{0}), [1, 2])$$

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

6 Results

In this section and the next we summarize the main findings from our simulation exercises as a number of different results.

Result 1: The percentage of bidders submitting late bids (bids in period T) is significantly greater in hard-close auctions than in soft-close auctions.

Support for Result 1 is found in Table 1, where we observe that the fraction of bidders submitting a late bid is always greater in hard-close auctions than in soft-close auctions with the same number of bidders, $n=2,3,4$, or 5.¹⁷ For example, in hard-close auctions with just 2 bidders, 62.3 percent of bidders are late bidders, while in soft-close auctions with 2 bidders only 2 percent of bidders are late bidders.^{18, 19} Similar differences in the frequency of late bidding are observed in comparisons of hard- and soft-close auctions involving 3, 4 and 5 bidders.²⁰ Figure 1a shows the frequency of bidders submitting late bids over time in hard-close auctions and Figure 1b shows the comparable frequency of late bidding over time in the soft-close auctions.

¹⁷For 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 .

¹⁸These percentages and the ones reported later in the paper are obtained by taking averages over the last 100 generations of all 20 simulations run for each treatment.

¹⁹Recall that we consider late bids only in the “first late” period, i.e., period T of soft-close auctions.

²⁰The reported differences in Table 1 (and further differences in bidding behavior reported below) between the two auction formats are always significant at the 5 percent level using a two-sample t -test with 38 degrees of freedom. In these tests, we use the average frequencies from each of the 20 simulation runs of the hard- and soft-close auction formats as independent observations.

HARD-CLOSE	Number of Bidders								
	Statistics		n=2		n=3		n=4		n=5
Number of early bids per bidder	1.910	(0.0252)	1.907	(0.0296)	2.089	(0.0524)	1.889	(0.0435)	
Fraction submitting late bid	0.623	(0.0182)	0.461	(0.0387)	0.326	(0.0392)	0.123	(0.0254)	
Revenue of the seller / mean value	0.954	(0.00140)	0.995	(0.000493)	0.9995	(0.0000983)	0.99998	(0.000007)	
Payoff of bidders / mean value	0.0229	(0.00699)	0.00174	(0.000164)	0.000130	(0.0000248)	0.000005	(0.000001)	
Freq. of early bidding automata	0.0893	(0.0113)	0.429	(0.0525)	0.557	(0.0682)	0.716	(0.0871)	
Freq. of late bidding automata	0.0270	(0.0042)	0.157	(0.0194)	0.104	(0.0145)	0.103	(0.0134)	
Freq. of cond. bidding automata	0.884	(0.1074)	0.414	(0.0508)	0.339	(0.0417)	0.182	(0.0224)	

SOFT-CLOSE	Number of Bidders								
	Statistics		n=2		n=3		n=4		n=5
Number of early bids per bidder	1.66198	(0.0344)	1.842	(0.0481)	1.819	(0.0400)	1.711	(0.0440)	
Fraction submitting a late bid	0.01991	(0.0084)	0.140	(0.0290)	0.0438	(0.0144)	0.0355	(0.0113)	
Revenue of the seller / mean value	0.995	(0.00215)	0.999	(0.000175)	0.99998	(0.0000131)	1.000001	(0.000003)	
Payoff of bidders / mean value	0.00226	(0.00108)	0.00256	(0.0000585)	0.000007	(0.000003)	0.000002	(0.0000007)	
Freq. of early bidding automata	0.972	(0.1181)	0.8054	(0.0979)	0.887	(0.108)	0.868	(0.106)	
Freq. of late bidding automata	0.00857	(0.0019)	0.0164	(0.0028)	0.0234	(0.0034)	0.0254	(0.0035)	
Freq. of cond. bidding automata	0.0194	(0.0031)	0.178	(0.0219)	0.0894	(0.0115)	0.106	(0.0131)	

Table 1: Benchmark Simulation Results (Averages and Standard Deviation of Averages over the Last 100 Generations)

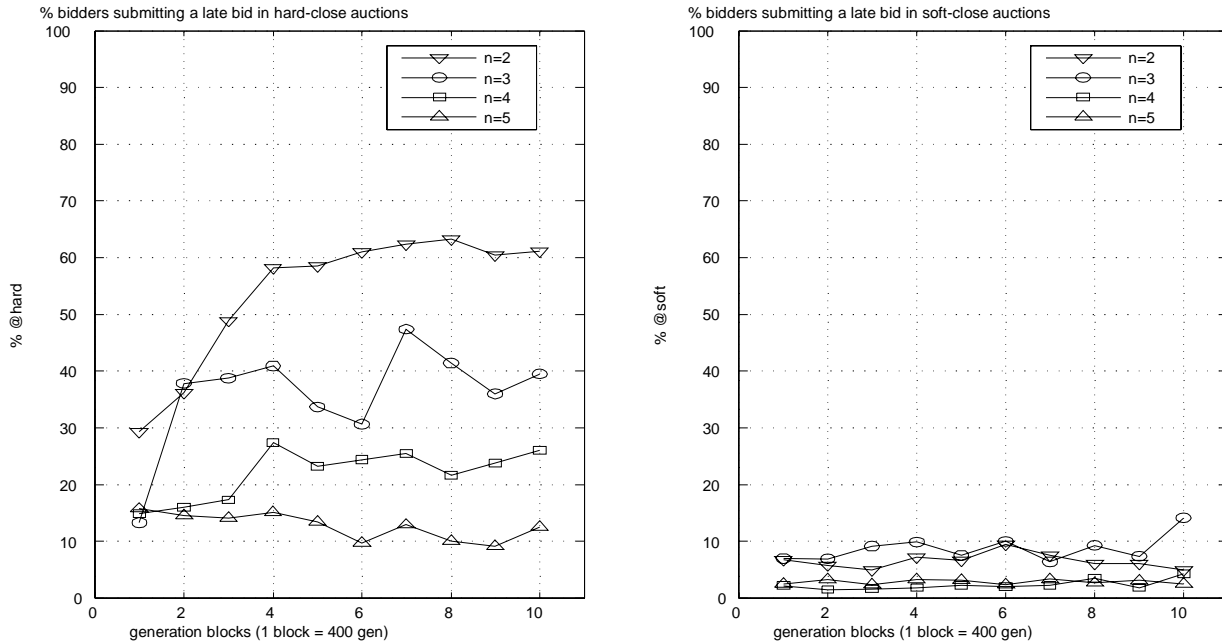


Figure 1: The average percentage of bidders successfully submitting a late bid in hard-close and soft-close auctions with 2,3,4 and 5 bidders.

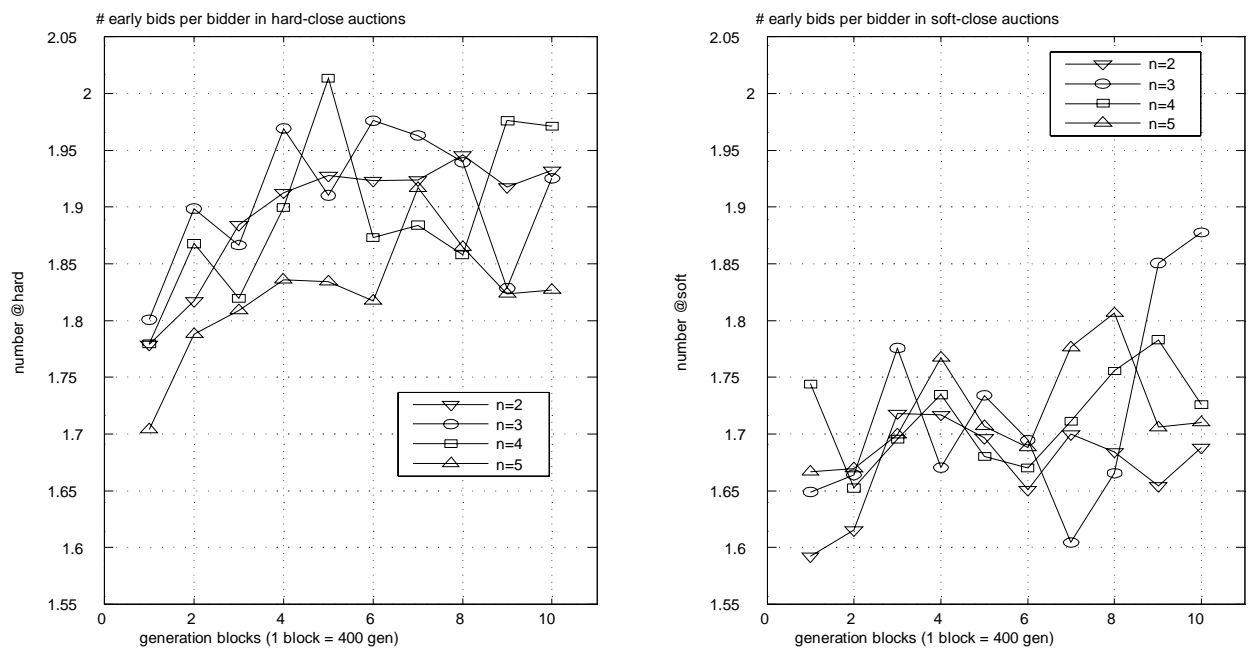


Figure 2: The average number of valid early bids per bidder in hard-close and soft-close auctions for $n=2,3,4$, and 5 bidders.

Result 2: The frequency of late-bidding decreases as the number of bidders increases in hard-close auctions. However, late bidding remains more prominent in hard-close auctions than in soft-close auctions with the same number of bidders.

Support for this finding is found in Figure 1a and in Table 1. Note that in all cases, the number of bids per bidder is significantly larger than one. Multiple bidding is prominent in all auctions. The evolution of early bids per bidders are given in Figure 2.

Result 3: The average revenue of sellers is significantly higher in soft-close auctions than in hard-close auctions for each number of bidders, and the average revenue of bidders is significantly lower in soft-close auctions as compared with hard-close auctions for each number of bidders.

Support for this finding is again found in Table 1. Results 1-3 suggest that the evolving strategies in the hard-close auctions should be different from the ones in other hard-close auctions.

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

Result 4: Bidders *always* learn to bid their full value as their final bid in both soft-close and hard-close auctions for all numbers of bidders, i.e., *bid-shaving is not observed*. More precisely, the percentage of

evolving automata that bid full value by period T is very close to 100 percent; it is never exactly 100 percent due to ongoing mutation.

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.

We now 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 with $n = 2, 3, 4$ or 5 bidders.²¹ Table 2 provides some further disaggregation – specifically the average frequencies of various “phenotypes” that exceed a small threshold, $\frac{1}{30}$ (i.e., 1 in every generation). The main finding from our analysis of these phenotypes is:

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

Support for Result 5 is found in Tables 1-2 and Figure 3. 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 more conditional phenotypes are observed 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. Still, this frequency is always less in hard-close auctions than in soft-close auctions for 2, 3, 4 and 5 bidders.

In soft-close auctions, more than 80 percent of all evolving strategies are characterized as “E” (unconditional early bidding) phenotypes (see Table 2). In hard-close auctions, the frequency of “E” phenotype ranges from 9 to 72 percent depending on the number of bidders (see again Table 2). In the two-bidder hard-close auctions, the phenotype “L 1 2 E 2 1” is observed with a frequency of 57 percent, higher than any other phenotype. Strategies in this phenotype tell the bidder to bid late as long as the rival bidders also bid late. Otherwise, an early final bid is placed. If the rival only submits early bids, this strategy switches back to late bidding. Otherwise, the early bidding automaton is played again. The other most common phenotype in two bidder hard-close auctions is “E 1 2 L 2 1” with a frequency of 28 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 or more bidders, we observe the “E” phenotype in very high frequencies ranging from 43 to 72 percent. The second most common phenotype is “L” with significantly lower frequencies ranging from 10 to 16 percent. There are other commonly observed phenotypes such as “L 1 2 E 1 1”, “L 2 1 E 1 2” and “E 1 2 L 2 1”. With 3 or more bidders, it appears to be more difficult to coordinate on late bidding in hard-close auctions.

²¹This average is found by taking the average over last 100 generations in 20 simulations for each treatment.

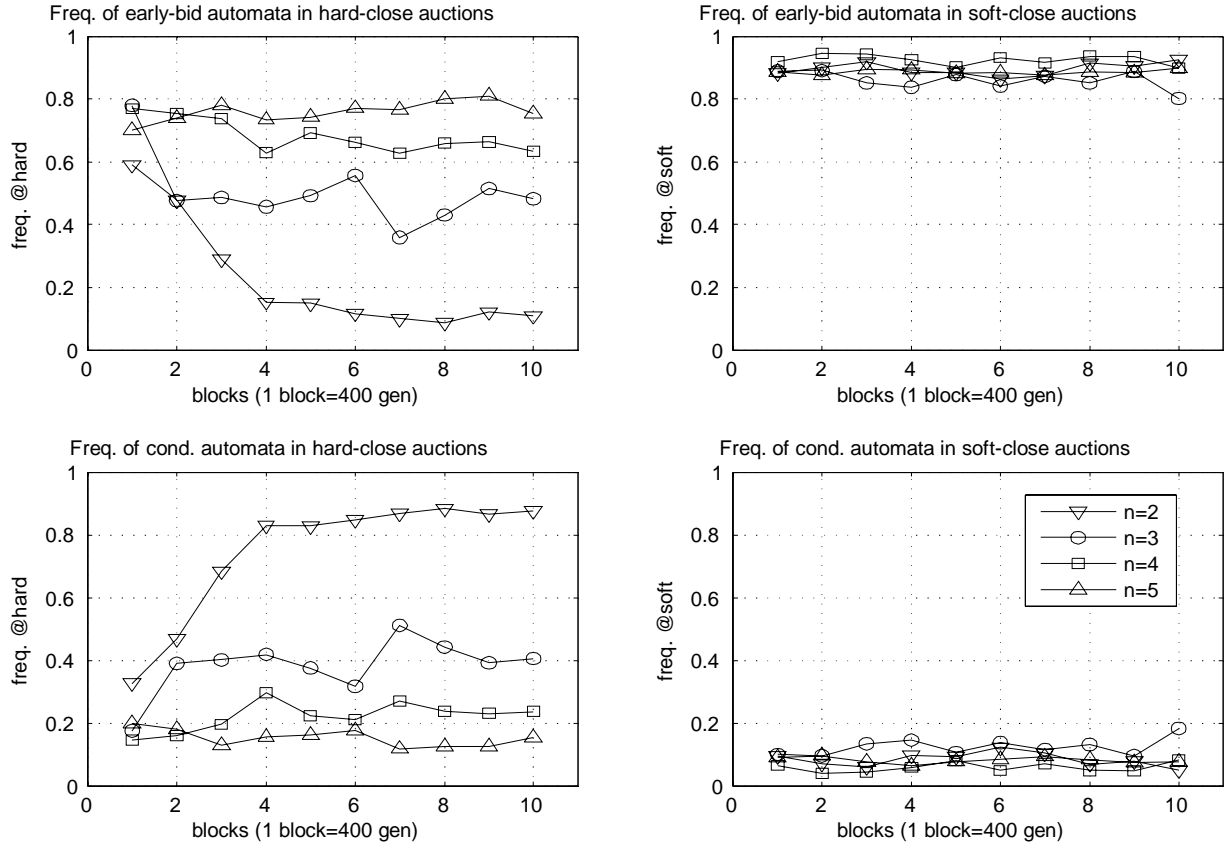


Figure 3: The average frequency of observing “unconditional early-bidding” and “conditional bidding” phenotypes of automata in hard-close and soft-close auctions with 2,3,4 and 5 bidders.

HARD-CLOSE	<i>Number of Bidders</i>			
<i>Phenotypes</i>	n=2	n=3	n=4	n=5
E	0.0893	0.428967	0.556583	0.7163
E 1 2 L 2 1	0.5727	0.0881	0.068667	0.01565
E 2 2 L 2 2	0.000267	0.064717	0.014217	0.003133
L	0.026967	0.157233	0.10435	0.102533
L 1 2 E 1 1	0.007067	0.046483	0.09425	0.003733
L 1 2 E 1 2	0.0009	0.040333	0.037317	0.02185
L 1 2 E 2 1	0.275317	0.047517	0.022467	0.004383
remaining:	0.027483	0.12665	0.10215	0.132417
SOFT-CLOSE				
<i>Phenotypes</i>	n=2	n=3	n=4	n=5
E	0.972017	0.8054	0.887133	0.8684
L 1 2 E 1 1	0.000183	0.047117	0.00085	0.002917
L 1 2 E 2 2	0.0004	0.034017	0.000817	0.007083
Remaining:	0.0274	0.113467	0.1112	0.1216

Table 2: Frequency of surviving automata phenotypes in benchmark simulations (Averages over the last 100 generations surpassing a frequency of 1/30)

7 Simulations With Adaptive Agents and Naive Incremental Bidders

Roth and Ockenfels' (2002) empirical findings suggest that there is a significant amount of "inexperienced" naive bidders participating in internet auctions. These are often first-time bidders. A typical first-time bidder behavior uses a naive "incremental bidding stage-game strategy." Following this strategy, a bidder increases the current price by bidding incrementally higher than the current second bid whenever he is not the current high bidder in the auction – in effect bidding until he discovers the current high bid, and bidding the necessary increment above that bid so as to achieve high bidder status. This incremental bidding strategy proceeds so long as the current price is lower than the 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 subsection, we introduce one naive bidder to each internet auction. These bidders should be viewed as "one-time bidders" and, indeed, they are replaced by a different naive bidder in each auction. These naive bidders use a simple incremental bidding strategy: they only bid whenever they are not the current high bidder, and 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 (second highest bid plus the increment Δ). We re-run all of our benchmark simulations by replacing one of the n bidders with a naive bidder. Figure 4, which is comparable to Figure 1, shows the frequency over time with which $n - 1$ adaptive bidders submit late bids in hard- and soft- close auctions when there is 1 naive incremental bidder present. Notice that, by comparison with our earlier baseline simulations without naive bidders (as shown in Figure 1) in these new simulation results with naive bidders, the frequency of late bidding is significantly greater in hard-close auctions and significantly lower in soft-close auctions. Indeed, in hard-close auctions with 1-3 adaptive bidders (we exclude the naive bidder), this percentage reaches and stays above 70 percent.

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 (2002) 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.

We note that in our earlier setup without naive bidders we can replicate Ariely et al.'s finding that late bidding increases as ρ is varied from a value less than 1 to being equal to 1. Indeed, we conducted such

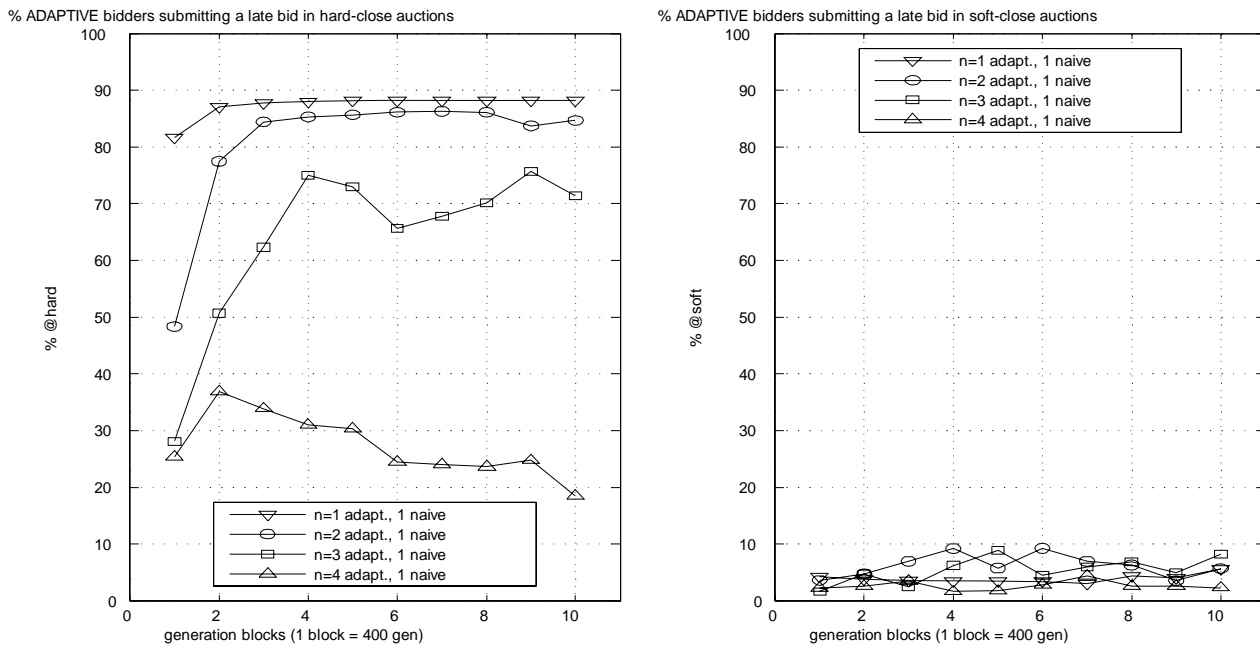


Figure 4: Simulation results with a naive incremental bidder. The percentage of “adaptive” bidders placing a successful late bid in hard- and soft-close auctions with 2,3,4, and 5 bidders. To find the percentage of all bidders placing a successful late bid, multiply the above numbers roughly by $\frac{n-1}{n}$, since naive bidders rarely place a late bid in the late stages of the evolution.

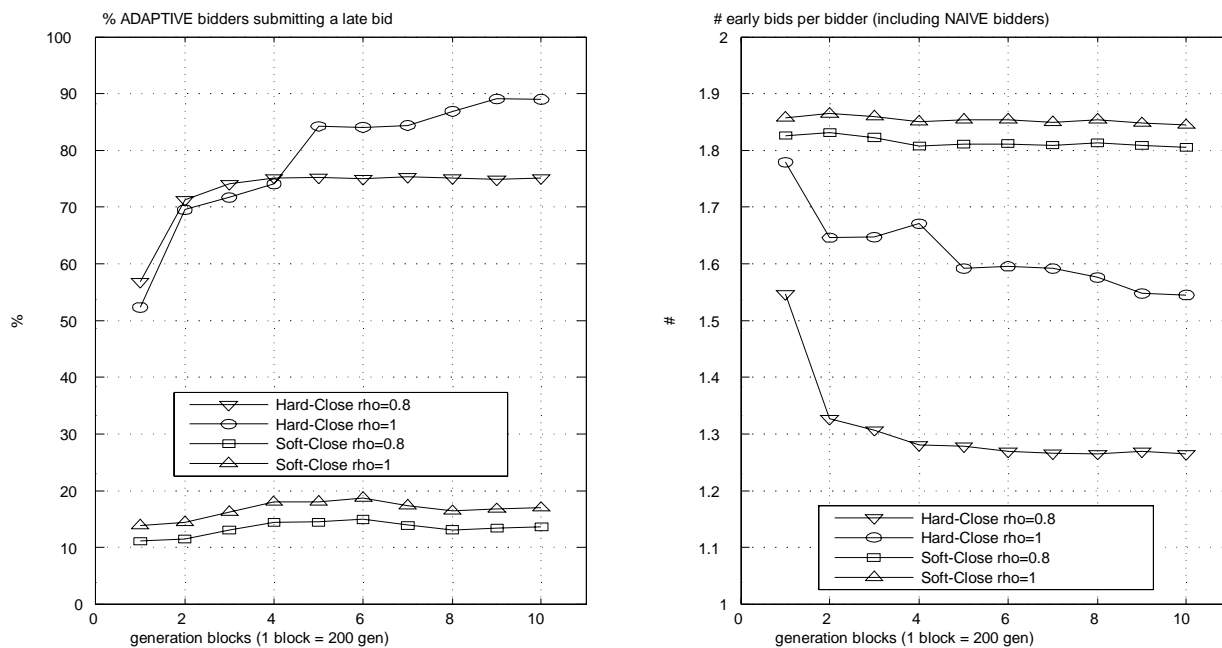


Figure 5: Simulations with 1 naive and 1 adaptive bidder. The frequency of the adaptive bidder submitting a successful late bid.

an exercise as part of our sensitivity analysis as reported on in Appendix B; this sensitivity analysis was limited to the 2-bidder case and involved changing only 1 parameter at a time from our baseline simulation values (as reported in Table 1). Table 3 in Appendix B reveals that in 2-bidder hard close auctions, the fraction of late bids is 33.5 percent when $\rho = .8$ and increases to 60.1 percent when $\rho = 1$.

To perform a similar comparison with naive incremental bidders, we modify our simulation setup somewhat so that it is more closely aligned with that of Ariely, Ockenfels and Roth (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 values for ρ - 0.8 and 1, the values adopted by Ariely et al. (2005). The simulation findings from this exercise are reported in Figure 5. The frequency with which adaptive bidders submit successful late bids is given in Figure 5a. Like Ariely et al., we observe a higher frequency of late bidding by adaptive bidders in hard-close auctions when $\rho = 1$ (nearly 100 percent) than when $\rho = 0.8$ (nearly 80 percent). By contrast, in soft-close auctions the frequencies of late bidding in the presence of a naive bidder are substantially lower, at around just 20 percent. Also, bidders submit significantly more early bids per bidder (including the naive bidder) in the soft-close auctions as compared with hard-close auctions (see Figure 5b).

8 Conclusions and Implications for Market Design

Our results show that late bidding is an evolutionarily sustainable phenomenon in hard-close auctions and is far less common or sustainable in soft-close auctions. Our findings are in accordance with empirical evidence of late bidding in hard-close 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.

Indeed, our simulation findings suggest that hard-close auctions raise *lower* revenue for sellers than soft-close auctions, and that bidders fare better in terms of their payoff in hard-close rather than soft-close auctions. This result bears additional emphasis from a market design perspective.²² Since internet auction web-sites view themselves as clearinghouses or intermediaries for the market transactions, their interests are not clearly aligned with sellers or buyers. However, information on which auction format favors sellers or buyers is of obvious use to these individuals, 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.²³ The points raised by our study set the stage for further investigation on the evolution of different market designs for internet auctions.

It should be noted that in our baseline simulations all bidders are “adaptive” learners, and eventually they learn to use “good” strategies, i.e., ones that have them bid their full valuations by the final period, T . 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. 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. 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.

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.

²²It 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.

²³Yahoo! auctions allow the seller to choose whether to have a hard- or soft-close to the auction.

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,” forthcoming, *RAND Journal of Economics*.
- [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] Borle, S., P. Boatwright and J. B. Kadane (2005) “The Timing of Bid Placement and Extent of Multiple Bidding: An Empirical Investigation Using eBay Online Auctions,” mimeo.
- [9] Dawid, H. (1999), *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models*, Second Edition. Springer, Berlin.
- [10] Engle-Warnick, J. and R. Slonim (2005) “Inferring Repeated Game Strategies from Actions: Evidence from Trust Game Experiments,” forthcoming, *Economic Theory*.
- [11] Goldberg, D. E. (1989), *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA.
- [12] Hasker, K., R. Gonzalez and R. C. Sickles (2004), “An Analysis of Strategic Behavior and Consumer Surplus in eBay Auctions,” mimeo.
- [13] Holland J. H. (1975), *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann-Arbor, MI.

- [14] Holland J. H. and J. H. Miller (1991), "Artificial Adaptive Agents in Economic Theory," *American Economic Review: Papers and Proceedings of the 103rd Annual Meeting of the American Economic Association*, 335-350.
- [15] Lucking-Reiley, D. (1999), "Using Field Experiments to Test Equivalence between Auction Formats: Magic on the Internet," *American Economic Review*, 89: 1063-1080.
- [16] Lucking-Reiley, D. (2000), "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics*, 48: 227-252.
- [17] Melnik, M. and J. Alm (2001), "Does a Seller's eCommerce Reputation Matter," mimeo.
- [18] McAfee, R. P. and J. McMillan (1987), "Auctions and Bidding," *Journal of Economic Literature*, 25: 699-738.
- [19] Michalewicz, Z. (1994), *Genetic Algorithms + Data Structures = Evolution Programs*, Second Edition. Springer-Verlag, Berlin.
- [20] Milgrom, P. R. and R. J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50: 1089-1122.
- [21] Miller, J. H. (1996), "The Coevolution of Automata in the Repeated Prisoner's Dilemma," *Journal of Economic Behavior and Organization*, 29: 87-112.
- [22] Moore, E. (1956), "Gedanken Experiments on Sequential Machines," *Automata Studies*. Princeton University Press, Princeton, NJ.
- [23] Ockenfels, A and A. E. Roth (2002), "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction," mimeo.
- [24] 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.
- [25] Vickrey, W. (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16: 8-37.

Appendix A: Proof of the Theorem

Proof of the Theorem: We prove the Theorem separately for hard-close and for 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 ex-post dominates any other strategy which does not involve bidding fraction 1 early in a stage auction. This will imply that the prior strategy ex-ante dominates any other strategy.

1. First, we consider a stage hard-close auction. Let σ^1 be a stage game strategy of bidder 1 with the highest fraction $\alpha \in \{0, \frac{1}{3}, \frac{2}{3}\}$ submitted in one of the first $T - 1$ periods and fraction $\alpha' \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$ 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 $\beta \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$ submitted in one of the first $T - 1$ periods and fraction $\beta' \in \{0, \frac{1}{3}, \frac{2}{3}, 1\}$ 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 .

HARD-CLOSE with 2 bidders	<i>Varying Parameters in the Benchmark Simulations, Ceteris Paribus</i>							
<i>Statistics</i>	$\epsilon = 20$	$\epsilon = 4000$	$m = 1000$	$m = 10^9$	$\rho = 0.8$	$\rho = 1$	$R = 1$	$R = 40$
number of early bids per bidder	1.9602	1.8830	1.9150	1.9854	1.8669	1.9322	1.7946	2.0210
fraction of bidders submitting a late bid	0.6188	0.6409	0.6068	0.6196	0.3355	0.6015	0.0459	0.6804
avr. revenue of seller / mean value	0.9535	0.9507	0.9448	0.9545	0.9429	0.9995	0.9950	0.9488
avr. payoff of bidders / mean value	0.0233	0.0247	0.0287	0.0228	0.0285	0.0003	0.0025	0.0256
freq. of early bidding automata	0.1110	0.0757	0.1237	0.1073	0.3812	0.2689	0.8896	0.0167
freq. of late bidding automata	0.0124	0.0192	0.0132	0.0505	0.0429	0.3711	0.0224	0.0057
freq. of cond. bidding automata	0.8766	0.9051	0.8632	0.8423	0.5760	0.3600	0.0880	0.9777
SOFT-CLOSE with 2 bidders								
<i>Varying Parameters in the Benchmark Simulations, Ceteris Paribus</i>								
<i>Statistics</i>	$\epsilon = 20$	$\epsilon = 4000$	$m = 1000$	$m = 10^9$	$\rho = 0.8$	$\rho = 1$	$R = 1$	$R = 40$
number of early bids per bidder	1.7842	1.4563	1.8564	1.7946	1.6401	1.8988	1.7347	1.6857
fraction of bidders submitting a late bid	0.0635	0.1708	0.0413	0.0800	0.1278	0.5335	0.0084	0.1414
avr. revenue of seller / mean value	0.9930	0.8053	0.9858	0.9941	0.9696	0.9996	0.9987	0.9812
avr. payoff of bidders / mean value	0.0035	0.0968	0.0086	0.0029	0.0151	0.0002	0.0007	0.0093
freq. of early bidding automata	0.9046	0.7328	0.9246	0.8713	0.7516	0.2876	0.9254	0.7917
freq. of late bidding automata	0.0181	0.0607	0.0187	0.0367	0.0305	0.2556	0.0028	0.0281
freq. of cond. bidding automata	0.0773	0.2065	0.0567	0.0920	0.2178	0.4568	0.0719	0.1802

Table 3: Sensitivity Analysis (Averages in the Last 100 Generations)

Appendix B: Sensitivity Analyses

We change one model parameter at a time in the comparative static exercises in auctions with 2 bidders only. In Table 3, the results of these exercises are reported. Specifically, we change 1) the spread of the value distribution from $\epsilon = 40$ to 20 and then to 4000, 2) the mean of the value distribution from $m = 10^6$ to 1000 and then to 10^9 , 3) the probability of last minute registration from $\rho = 0.9$ to 0.8 and then to 1, and finally, 4) the number of auction stages in a repeated game block from $R = 20$ to 1 and then to 40. The main findings from our sensitivity analysis may be summarized as follows (see Table 3 below for details).

- The probability of late bid registration should be less than 1 in order to observe more prominent early bidding in soft-close auctions. Otherwise, we observe a 50 percent frequency of late bidding even in soft-close auctions since the expected payoffs are the same for bidding in any period.
- An increase in the number of stages R in a repeated-auction block has a positive effect on late bidding in hard-close auctions. As R increases, the late bidding frequency increases in the hard-close auctions. There is no late bidding when $R = 1$.