Forward-looking bidding in online auctions

Robert Zeithammer

Assistant Professor of Marketing
The University of Chicago
Graduate School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637

e-mail: rzeitham@gsb.uchicago.edu
phone: (773)-834-3769
fax: (773)-702-045

October 16, 2005

I would like to thank John Little, Birger Wernerfelt, Elie Ofek, Dan Ariely, Sergei Izmalkov, Jeroen Swinkels, Jean-Pierre Dube, Uri Simonsohn, two anonymous reviewers, and seminar participants at several marketing departments around the country, where I presented an early version of this paper in 2002.

Forward-looking bidding in online auctions

Brief abstract

At Internet auction sites like eBay, nearly identical goods are often sold in a sequence of auctions, separated by small amounts of time. Upcoming auctions are announced several days in advance, so buyers can benefit from forward-looking strategies that take that information into account. This paper develops a model of such bidding, provides empirical evidence of the model's relevance to actual behavior on eBay, and discusses the general implications of forward-looking bidding for sequential auction-driven marketplaces.

Keywords: Auctions, Forward-looking consumers, Bidding, Auction econometrics, Game theory.

Internet auction sites like eBay are increasingly being used to sell mass-produced consumer durables: the largest eBay categories in dollar terms are: cars, consumer electronics, computers, clothing/accessories, and books/movies/music (2004 company report). Since the ending times of the individual auctions are not synchronized, each of these markets evolves as a sequence, allowing bidders to focus on the auction that will end first, while accounting for the fact that there will be other auctions later. Because online auctions are usually listed for several days before concluding, detailed information about what and when will be sold in the near future is available to bidders. Two important questions thus arise, namely how should rational consumers interested in buying just one unit of the good use such information in forming their bids, and whether eBay bidders actually use the information accordingly. To answer these questions, this paper develops a new model of equilibrium bidding in a very long sequence of auctions, and provides empirical evidence of the model's relevance to actual behavior of eBay bidders.

The model assumes that each product-category is horizontally differentiated into several types of goods, with each bidder having a unit demand for only one type of the good. For example, a consumer may be shopping for one DVD of her favorite movie, or for one unit of a specific brand and model of an MP3-player. Each bidder therefore faces a tradeoff between winning now and winning later, arising from the fact that the individual desired units are perfect substitutes: the winner of each auction exits the marketplace and hence foregoes the expected surplus from participating in future auctions that also sell her desired good, possibly for a lower price. The future expected surplus foregone by the current winner is an opportunity cost of winning now, and rational forward-looking bidders should reduce their bids relative to the myopic bidding strategy that would be optimal in the absence of future auctions selling the desired type of good.

The model departs from previous models of bidding in a sequence of heterogeneous substitutes (Engelbrecht-Wiggans 1994, Jofre-Bonet and Pesendorfer 2003) by assuming that bidders know not only the type of the current product they are bidding on, but also what type will be sold next and when. In other words, the bidders are not only forward-looking in that they anticipate a future auction, but also *forward-seeing* in that they know detailed information about several future auctions. The expected future surplus, and hence the opportunity cost of winning now, is a function of the available information about what and when will be sold in the near future – both the timing of the upcoming auctions as well as the types of products sold in those auctions. The equilibrium analysis of the game with forward-seeing bidders is complicated: the expected surplus function depends on the bidding strategy while the bidding strategy in turn depends on the expected surplus function. While the equilibrium bidding strategy is thus intractable in closed form, this paper shows that there exists a well-behaved symmetric pure-strategy Markov Perfect equilibrium bidding function whose comparative statics can be characterized qualitatively without relying on specific assumptions about the distribution of personal valuations in the bidder population.

The properties of the equilibrium bidding strategy depend on how much detail of the available information about near-future auctions do the bidders actually use, i.e. how sophisticated they are in taking the information into account. Three nested levels of such information-usage sophistication are considered and empirically tested: First, when bidders ignore the information completely, the model reduces to a special case of the model of Engelbrecht-Wiggans (1994), in which bids do not depend on short-term variation in the near-future frequency of auctions or on variation in the near-future incidence of specific product-types. Second, auctions ending within the next hour are highlighted in red on eBay, so frequency of near-future auctions is easier to discern than the specific attributes of the objects sold. An intermediate level of sophistication therefore involves bidders, who see only the general frequency of auctions in the near-future, but not the types of the individual future objects. Such bidders should reduce their bids more whenever there are more auctions ending soon, because that decreases the expected waiting time until another unit of their desired types comes up for sale, and hence increases the expected future surplus. Finally, each bidder can actually examine the near-future auctions closely and base her bidding strategy not only on timing, but also on the types of objects actually coming up for sale. Then, the opportunity cost of winning today becomes a function of personal preferences for the future items - the more personally desirable products coming up and the sooner they are coming up, the higher the opportunity cost and the lower the bids.

How much detail of the available information about near-future auctions do actual eBay bidders use is an empirical question, and this paper proposes an empirical strategy to answer this question using standard eBay data. The empirical strategy relies on measuring the relationship between current bids and both object-types and ending-times of near-future auctions, a relationship for which the three levels of information-use sophistication generate the above-described nested restrictions of the "most sophisticated" bidding function. Two different datasets are used in the test, one from the MP3-player category with each player brand-model combination considered to be a different type, and one from the DVD movie category with different product-types assigned to different movie-titles. The empirical test on both product categories rejects the two nested simpler models in favor of the "most sophisticated" model, in which bidders take their personal preferences for specific future products into account.

Therefore, the model of forward-looking behavior proposed here is relevant to understanding the demand-side of auction-driven marketplaces like eBay.

The paper is structured as follows: Section II provides a brief review of the relevant literature. Section III then presents the model that constitutes the main theoretical contribution of this paper, and Section IV discusses the robustness of the model predictions to perturbations in the assumptions. Section V presents the results of an empirical test, and Section VI concludes by discussing the implications of the present findings for both researchers and participants of online auction marketplaces.

II. Literature review

The theoretical model proposed here is not confined to online auctions, and contributes to the general auction theory literature. Unlike most previous work on sequential auctions that focuses on price-trends in finite sequences of auctions (motivated by the "price-decline anomaly" documented in real-world auction-sequences by Ashenfelter (1989) and others), this paper investigates the influence of public information about near-future auctions on infinite-horizon steady-state bidding. The model extends the finite-horizon identical-goods model of Milgrom and Weber (1999) to an infinite horizon and horizontally differentiated goods. A very simple differentiation into a finite number of mutually exclusive types is assumed, so the extension amounts to assuming several randomly-interlaced sequences of identical-goods sequences. The proposed model is thus the simplest model that involves unit-demand bidders and non-trivial information about the near-future auctions. Because the model considers a sequence of auctions for non-identical objects with knowledge of future objects, it also extends the model of Gale and Hausch (1994), who examine the case of continuously heterogeneous objects by focusing on the special case of two bidders and two auctions. The extension beyond two auctions is accomplished here thanks to the simplifying assumption on the product-heterogeneity being captured by a finite number of types. The closest simpler benchmark is provided by the model of heterogeneous but unseen future objects (Engelbrecht-Wiggans 1994), a special case of which is nested in the proposed model: when only some *fixed* common distribution of future products is known, bidders can still engage in forward-looking strategies, but they are unable to use the forward-seeing strategies investigated here. Within the online-auction literature, the issue of multi-auction bidding has not been addressed except for work by Bajari and Hortacsu (2003), who study bidder entry in common-value auctions, and Dholakia and Soltysinski (2001), who find a "herding bias" – consumers flocking to popular auctions despite the existence of other auctions for substitute items. The "herding bias" is especially relevant to this work because it provides another layer of behavioral complexity on top of the rational behavior described here.

On the empirical front, the most related work is a recent paper by Jofre-Bonet and Pesendorfer (2003). They investigate sequential auctions for highway construction procurement contracts in California, and they find evidence of forward-looking behavior using a structural model. Forward-seeing behavior is not part of their model because bidders are assumed to not take public information about upcoming auctions into account. Another important difference is that this paper conducts an empirical test non-parametrically without functional-form assumptions that would be necessary for a structural econometric model.

III. Theory of forward-seeing bidding

Several simplifying assumptions are needed to obtain a tractable model. Online auctions usually remain open for several days, potentially leading to strategically rich within-auction behavioral dynamics (see Ariely and Simonson 2001 for a discussion). I abstract from these within-auction dynamics, and model each individual auction as an instantaneous sealed-bid auction occurring at the time of the actual auction's end. Validity of this abstraction is supported by the fact that bidding on eBay both should and does tend to happen at the very end of each auction, not giving the competitive bidders time to react to each other's bids (Roth & Ockenfels 2002). To approximate the price-determination in the eBay's ascending auction within the sealed-bid abstraction, the models examine second-price auctions, in which the highest bidder wins the object, but only pays the second highest bid as the price.

Since ending times of online auctions are not synchronized, the sealed-bid abstraction results in a model of sequential auctions, in which one auction ends before the next one starts and several upcoming auctions are already known. The eBay webpage design reinforces this conceptualization by listing auctions in a sequence ordered by ending time, and by allowing bidders to place known future auctions on their private watch-lists. Anectodal evidence from an eBay community newsgroup suggests that the conceptualization resonates with at least some eBay bidders. When I posted a question to an eBay newsgroup asking how to bid in several different auctions for a particular model of a digital camera, one user replied: *"Place bids on only one item at a time and put all the rest on your watch-list. If you are outbid on the first item, move to the next ending time on your watching page"*.

Consumer valuations of a single unit of the good are assumed private and independent across bidders. This is a reasonable model of private-consumer utility in the economically largest eBay categories involving mass-produced consumer durables that are usually purchased for private use, and depreciate quickly due to obsolescence. In particular, the assumption is reasonable for MP3 players and DVDs considered in the empirical section. The independence assumption resonates with the sealed-bid abstraction outlined above because bidders with independent preferences do not try to learn about their own preferences from other people's bids.

One additional simplifying assumption is necessary for a tractable model, namely that bidders do not consider past prices. Suppose they did. Then, since past prices are upper bounds on the past bids of competing bidders who did not win the past auction and hence may have survived until today, the outcomes of past auctions could be informative about the level of future competition. To make matters even more complicated, the past price-determining bidder would have slightly different information about the remaining competition than the bidders whose past bids were less than the past price. This asymmetry would escalate over time as pointed out by Milgrom and Weber (1999). Assuming the effects of past prices away does not have a big impact on the model because the effects are likely to be second-order, especially in an eBay-like environment characterized by a fluctuating and unobservable bidder pool.

Basic 1-period look-ahead Model: Assumptions

There is an infinite sequence of instantaneous second-price sealed-bid auctions occurring at distinct and countable points of continuous time. The waiting time ω between auctions varies stochastically and independently, according to a known distribution. Bidders discount future utility exponentially by factor δ per unit of time.

Each auction sells one object. The objects offered for sale are horizontally differentiated into $K+1$ types, with probability of type k captured by rate ρ_k. Each bidder desires one of the first K product-types in the sense that all non-desired types give the bidder zero utility while the desired type gives the bidder a positive utility. For example, each bidder is interested in buying only one particular movie-title or only one particular model of MP3 player.[1] No bidders desire the last $K+1^{st}$ product type, which captures various suspect "free" offers as well as poorly described and misplaced goods that are bound to clutter any marketplace.

[1] The category of MP3 players is differentiated by brand-model combinations like "Rio 500", and the category of movies on DVD is differentiated by title of the movie.

Select an arbitrary "current" auction as the origin of indexing, so the current auction has index 0, the immediately following auction index 1, the auction after that index 2, etc. To capture the type information, let $\varphi_{j,k} \in \{0,1\}$ be the indicator function equal to one when auction j is of type k, and equal to zero otherwise. To capture the waiting-time information, let ω_j be the waiting time between auction $j-1$ and j. The key innovation of this model is that bidders of type k know not only the desirability of the current product $\varphi_{0,k}$, but also the information $(\varphi_{1,k}, \omega_1)$ that arises from seeing forward, namely whether they desire the next product and how far in the future will the next sale occur. The fundamental difference between the two kinds of forward-seen information is that $\varphi_{1,k}$ is inherently type-specific whereas ω_1 is the same for all types.

N_k bidders participate in each auction of type k. Bidder i of type k considers all her desired objects to be identical, and can derive a private value of $v_{i,k}$ dollars from any one of them. The individual private values $v_{i,k}$ are drawn independently across bidders from a known probability distribution F_k with full support on [0,1] and a continuous density f. Therefore, the private valuation to bidder i of type k of the current object is $\varphi_{0,k} v_{i,k}$, and the private valuation to the same bidder of the next object is $\varphi_{1,k} v_{i,k}$, where $v_{i,k} \sim F_k$. All bidders can only derive utility from a single unit of their desired good, so once a bidder owns one unit of her desired type, all subsequent units are worth zero to her. The bidders have no memory, so they cannot base their actions on outcomes of past auctions.

Assume resale is too costly for a private consumer to warrant speculative purchases of multiple objects for future resale, so each auction's winner exits the game and is replaced by another randomly-drawn bidder. Bidders also exit at random with an exogenously given probability $(1-\lambda)$ per time-period. Some bidder-replacement beyond the replacement of the winner is an essential feature of a realistic steady-state model, because when only the winner is replaced and bidders stay until they win an auction, the distribution of the steady-state survivors degenerates to a group of bidders with zero valuations.

Basic 1-period look-ahead Model: Equilibrium bidding strategy

In a symmetric pure-strategy Markov Perfect Equilibrium, the strategy can depend only on the publicly-known state $(\varphi_{0,k}, \varphi_{1,k}, \omega_1)$ and on each bidder's private information $v_{i,k}$. Since the product-types have their own bidder populations that evolve without interacting across type boundaries, the optimal bidding problem is symmetric across types: for each type k, the remaining types $\{1,2,\ldots,k\text{-}1,k\text{+}1,\ldots K,K\text{+}1\}$ can all be lumped into "other" undesired type. Without loss of generality, I can therefore solve for the equilibrium bidding strategy in the case K=1, suppressing all k subscripts for clarity.

Let $K=1$, and let $S(\varphi_0, \varphi_1, \omega_1, v_i \mid c_0)$ be the bidder i's continuation value of the game in case of a loss today, i.e. steady-state expected future surplus of bidder i conditional on him losing today's auction. All bidders use the same S in a symmetric equilibrium. It will become clear that the continuation value relevant at the margin depends on the current competition, so let c_0 be the highest competing bid that would arise if $\varphi_0 = 1$. Let G be the distribution of c_0. Then, each bidder with valuation v solves the following utility-maximization problem to find the optimal steady-state bid $b(\varphi_0, \varphi_1, \omega_1, v)$:

$$b(\varphi_0, \varphi_1, \omega_1, v) = \arg\max_b \int^{b} (\varphi_0 v - c_0)\, dG(c_0) + (\delta\lambda)^{\omega_1} \int_{b} S(\varphi_0, \varphi_1, \omega_1, v \mid c_0)\, dG(c_0) \qquad (1)$$

In equilibrium, the expected surplus function must account for the fact that other bidders are also reducing their current bids, so the current competition is weaker than if the competitors were not strategically forward-looking, and the future competition depends on the current competition. Furthermore, in an infinite-horizon setting employed here to capture a mature ongoing market, future bidders will again be reducing their bids as a function of the future's future and at least some of those future bidders will be current competitors who lost the present auction. Therefore, the expected surplus function *S* depends on the bidding strategy *b,* which in turn depends on the expected surplus function *S*. These dependencies make a closed-form solution of the model unavailable, but a well-behaved equilibrium exists as shown in Proposition 1:

Proposition 1 (proof in Appendix): There is a symmetric pure-strategy Markov Perfect equilibrium characterized by a bidding function $b(\varphi_0, \varphi_1, \omega_1, v)$ that satisfies:

$$\begin{aligned} b(1, \varphi_1, \omega_1, v) &= v - (\delta\lambda)^{\omega_1} S(1, \varphi_1, \omega_1, v \mid c_0 = b(1, \varphi_1, \omega_1, v)) \\ b(0, \varphi_1, \omega_1, v) &= 0 \end{aligned} \tag{2}$$

Where the function *S* satisfies a set of Bellman equations:

$$S(\varphi_{0,1}, \omega_1, v \mid c_0) = \tag{3}$$

$$= E_{\varphi_2, \omega_2}\left[\int^{b(\varphi_{1,2}, \omega_2, v)} (v - c_1)\, dG(c_1 \mid c_0, \varphi_{0,1,2}, \omega_{1,2}) + (\delta\lambda)^{\omega_2} \int_{b(\varphi_{1,2}, \omega_2, v)} S(\varphi_{1,2}, \omega_2, v \mid c_1)\, dG(c_1 \mid c_0, \varphi_{0,1,2}, \omega_{1,2}) \right]$$

The bidding strategy *b* has several striking properties: First, it does not directly depend on *G* – a consequence of the general truth-revealing property of the second-price sealed-bid auction. However, $b(\varphi_0, \varphi_1, \omega_1, v)$ does depend on the current competition inasmuch as the current competition is informative about the future competition: When evaluating the option value of the future, the bidder assumes that she will lose the current period to a competitive bid that exactly matches her current bid. This "tie" is the only situation in which raising the current bid slightly changes the outcome of the game, and *S* given $c_0 = b(1, \varphi_1, \omega_1, v)$ is therefore the opportunity cost relevant at the margin. In other words, each bidder assumes she is pivotal to the outcome of the game. Finally, it is notable that bidders only submit positive bids on products of their desired type – a result that will lead to identification of personal preferences in the empirical test.

The bidding function is fully characterized by the expected surplus function *S*, which is in turn characterized by the steady-state distribution of the future competition c_1 conditional on the current competition c_0 and all the state variables involved in the relationship between them: $G\left(c_1 \mid c_0, \varphi_0, \varphi_1, \varphi_2, \omega_1, \omega_2\right)$. In equilibrium, the surplus function must reflect the actual expected surplus given that everyone uses the optimal bidding strategy (2). Therefore, the equilibrium expected surplus function must satisfy the Bellman equation (3). Such an *S* function exists because of the continuity of *f*, compactness of its support, and the fact (shown in the proof of Proposition 1) that the slope of *S* in any of its arguments is uniformly bounded. However, equilibrium *S* cannot be expressed in closed form even for a simple distribution *F* and small *N*. Despite the lack of a closed-form solution, some general comparative statics of the bidding function can be derived from an analysis of the Bellman equation (see Appendix for a proof):

Proposition 2: For all *F*, the equilibrium $b\left(\varphi_0, \varphi_1, \omega_1, v\right)$ has the following properties:

1) $b\left(1, \varphi_1, \omega_1, v\right)$ increases in ω_1
2) $b\left(0, \varphi_1, \omega_1, v\right) = 0 < b\left(1, 1, \omega_1, v\right) < b\left(1, 0, \omega_1, v\right) < v$ for all $v>0$.
3) $b\left(1, \varphi_1, \omega_1, v\right)$ decreases in ρ

The first property shows that bids decrease when the future gets closer in the sense that ω_1 decreases. In the empirical section, a generalization of this result will be tested, namely the prediction that bids decrease as the number of auctions in the next hour increases. The second property contains several important results. The first inequality shows that all bidders with positive valuations submit positive bids on objects. This both guarantees trade and allows an analyst to identify individual type-preferences in the data by interpreting a positive bid on a type as an indication of that type's desirability to the given bidder. The second inequality in 2), also tested in the empirical section, is the main result of this paper because it shows that all forward-seeing bidders of all types bid less when they see their desired object in the next period compared to when they see an object they don't desire. Finally, the third inequality provides a comparison of forward-seeing behavior to the myopic benchmark: forward-seeing bidders always bid less than they would if they were myopic because myopic bidders have a dominant strategy to bid their valuation in a second-price sealed-bid auction (Vickrey 1961). The fundamental reason for the third inequality is the positive opportunity cost of winning, so it will hold for any forward-looking bidding strategy, even without forward-seeing.

The third property shows that as the long-term rate ρ of desired products increases, the bids decrease. The reason is that forward-seeing bidders are also forward-looking beyond the near future they can see. The result would hold even if the bidders did not know their (φ_1, ω_1) and hence could not be forward-seeing. The resulting model would be analogous to the model of Engelbrecht-Wiggans (1994), so this result shows how that important benchmark model is nested within the model proposed here. The empirical section will not be able to identify this effect from a type-specific effect because each type is, by definition, only observed with one long-term rate. A generalization of the basic model that will inform empirical testing is discussed next.

Multi-period look-ahead model

The above basic model contains all the intuition of a more general model, in which bidders see more than one period forward - a realistic extension most relevant to the data at hand. When bidders are able to see $A > 1$ periods ahead, the forward-seeing information states are of the form (Φ, Ω) where $\Phi = (\varphi_1, ..., \varphi_A)$ and $\Omega = (\omega_1, \omega_2, ..., \omega_A)$. Two empirically relevant summary statistics of (Φ,Ω) will be considered: $H(\Omega)$ = the number of auctions in the next hour implied by Ω, and $w(\Phi,\Omega)$ = waiting time until the first future auction that sells a product of the same type as the current product. $H(\Omega)$ is relevant because eBay auctions ending in the next hour are highlighted in red, and $H(\Omega)$ is thus very easy to discern at a glance. $w(\Phi,\Omega)$ is relevant because it is invariant to the way consumers actually use the eBay website, i.e. whether they search for the all listings in the category or for listings of their specific product-type only. Given these definitions, the same arguments as in the proofs of Propositions 1 and 2 can be used to show that there is a Markov Perfect symmetric-equilibrium pure strategy $b(\varphi_0,\Phi,\Omega,v)$, with the following properties:

Corollary 1: When bidders see *A* periods ahead, the equilibrium bidding function $b(\varphi_0,\Phi,\Omega,v)$ has the following properties:

1) $b(1,\Phi^1,\Omega^1,v)\le b(1,\Phi^0,\Omega^0,v)$ for all Φ^i,Ω^i such that (Φ^1,Ω^1) has one additional listing of any type in the next hour and is otherwise the same as (Φ^0,Ω^0), so $H(\Omega^1)>H(\Omega^0)$
2) $b(0,\Phi,\Omega,v)=0<b(1,\Phi^1,\Omega,v)<b(1,\Phi^0,\Omega,v)<v$ for all v>0 and for all $\Phi^1\succ\Phi^0$, where $\Phi^1\succ\Phi^0$ is defined by $\varphi_a^1\ge\varphi_a^0$ for all a, $\varphi_b^1>\varphi_b^0$ for some b.
   In particular, the inequality holds for all $\Phi^1\succ\Phi^0$ such that $\Phi^0=0$ and $\Phi^1\ge 0$.
3) $b(1,\Phi^1,\Omega^1,v)<b(1,\Phi^0,\Omega^0,v)$ for all Φ^i,Ω^i such that $w(\Phi^1,\Omega^1)<w(\Phi^0,\Omega^0)$ and the continuation of the (Φ^0,Ω^0) sequence is the same as the continuation of the (Φ^1,Ω^1) sequence after $w(\Phi^0,\Omega^0)$.
4) Δb_a decreases in with a,

   where $\Delta b_a=b(1,\Phi_a^0,\Omega,v)-b(1,\Phi_a^1,\Omega,v)$ and $\varphi_{a,j}^i=\begin{cases}0 \text{ when } j<a\\ i \text{ when } j=a\\ \varphi_{a,j}^{-i} \text{ when } j>a\end{cases}$

Claim 4) examines how the magnitude of the bid-decrement shown to be positive in 1) changes as the first occurrence of the desired product-type gets more distant in the sequence, keeping timing the same. The decrement becomes smaller as the future recedes into the distance - an immediate consequence of discounting and the chance of attrition.

To construct statements about average differences in bids for empirical testing, it is necessary to average over the parts of (Φ,Ω) kept constant in each claim of the Corollary 1. Let $\overline{b}(x,v)=E\left[b(1,\Phi,\Omega,v)\,|\,x\right]$ stand for the expected bid of a bidder with valuation *v* conditional on *x*, with the expectation over all of the components of the states that vary as *x* remains constant. Then, as long as future timing Ω is independent of future types Φ and the continuation of the sequence independent of its beginning, the claims in Corollary 1 average out to testable predictions:

Model predictions: If bidders act consistently with the proposed model, the following four relationships will hold for all valuations *v>0* and for all desired types *k* =1…K:

1) $\overline{b}(H,v)$ decreases in number of auctions in the next hour *H(*Ω*)*
2) $\overline{b}(\Phi^1,v)<\overline{b}(\Phi^0,v)$ for all $\Phi^1 \succ \Phi^0$ such that $\Phi^0=0$ and $\Phi^1 \geq 0$, so $\overline{b}(\Phi,v)$ decreases in the indicator function $\mathrm{I}\Phi^1$ of Φ^1.
3) $\overline{b}(w,v)$ increases in waiting time until the same type *w*
4) $\Delta\overline{b}(a,v)$ decreases in a, where $\Delta\overline{b}_a=\overline{b}(\Phi_a^0,v)-\overline{b}(\Phi_a^1,v)$ and $\varphi_{a,j}^i=\begin{cases}0 \text{ when } j<a\\ i \text{ when } j=a\end{cases}$

   So $\overline{b}(\Phi,v)$ "decreases" in the indicator function $\mathrm{I}\Phi_a^1$ of Φ_a^1, as shown in 2), and the decrease is attenuated by *a*.

The above predictions nest the predictions of models with less sophisticated bidders as follows: If the bidders see only the H summary of the near future auctions, then prediction 1) will hold but the other predictions will not because they all rely on the bidders considering specific types of the future products. When the bidders do not see forward, none of the predictions will hold.

IV. Robustness of the model predictions

Several assumptions of the basic model can be relaxed without altering the key predictions of Proposition 2, this section discusses these relaxations in turn.

Stochastic and unknown number of competing bidders

Within each product-type, number of bidders N_k is assumed to be the same in each period. Specifying the model with a fixed N_k simplifies the exposition without sacrificing generality because the model's qualitative conclusions will not be sensitive to variations in the assumption about the bidder pool as long as some bidders stay for multiple periods and there is a well-defined steady-state distribution of the number of bidders present. Since current competition generically does not enter the bidding strategy, the only difference a stochastic N_k would make to the results is that the entire RHS of equation (3) would have to be integrated over the steady-state distribution of future N_k, adding another argument to the expectation.

Bidders desiring more than one type of product

Another variation of the model that can be accommodated is allowing each bidder to idiosyncratically desire more than one type, but still have only unit demand in the category and still consider all desired types identical in terms of utility. Thus, all private single-unit valuations v_i would be drawn from some common distribution F, and φ would have to be specified for every bidder as $\varphi_{j,i} = 1$ when bidder *i* desires the type of product sold in auction *j*. Then, the private value to bidder *i* of product sold *j* auctions from now would be $\varphi_{j,i} v_i$. This structure of preferences may be relevant to the MP3 player category, where each bidder may only have use for one player but be indifferent among several models. In contrast, each movie-category bidder may have use for multiple DVDs as long as they all contain different movies. Allowing each bidder to have idiosyncratic preferences over multiple types would terminate the symmetry and independence across types that allowed the analysis of K=1 to be without loss of generality, and the model would have to be specified in terms of a set of type-specific equilibrium surplus functions S_k. Then, it could again be shown that there is a set of type-specific equilibrium bidding functions $b_k\left(\varphi_{0,i}, \varphi_{1,i}, \omega_1, v_i\right)$ that all satisfy an analogue of Proposition 2, with $\varphi_{1,i}$ assuming the role of φ_1. This model is investigated briefly in the empirical section, and as predicted above, the data lends no support to this model in the movie data while finding at least some weak evidence for it in the MP3 player data.

Presence of speculators

The model assumes that the bidders are private consumers who buy for private use and not for resale. This assumption can be justified by the fact that effective selling on eBay requires a much greater set of capabilities than buying on eBay: while buyers can treat eBay as any other online store, sellers need to have at least minimal web-publishing skills along with ability to ship goods and accept payments by various electronic methods. Speculation for resale on eBay does not seem to be empirically prevalent in either of the two datasets considered in Section V: over 99 percent of the bidders are not observed selling anything in the same category within a month, and of the several thousand sellers observed, only about 2.5 percent submitted at least one bid per month within the entire category, all together affecting less than three percent of the auctions.

It remains to be shown that even in the presence of a small fraction of speculators, Proposition 1 and 2 remain valid descriptions of private-consumer behavior in eBay-like sequential auctions. The main reason why speculative considerations do not affect the qualitative predictions of the theory in online-auction settings is that the predictions concern the impact of practically immediate future auctions (ω_1 is usually less than an hour), whereas any resale of the current item is delayed by at least several days (median auction-duration on eBay is a week). This necessarily delayed resale would manifest itself in the model as follows: Suppose that in addition to the N_k consumer bidders, there is one speculator intent on reselling any purchased units in the future. Because auctions take time to conclude and auction 1 is already listed, the speculator cannot resell the current unit (should he win it) until *after* period 1. In other words, the expected stream of upcoming auctions is unaffected until after period 1 even in the presence of the speculator. Therefore, equation (1) holds with c_0 reinterpreted as the maximum of the competitive consumer bid and the speculator's bid, and the $S\left(\varphi_0,\varphi_1,\omega_1,v_i \mid c_0\right)$ implicitly combining the possibility that the current winner is the speculator (making the distant future more desirable by virtue of the current unit returning to the marketplace) with the possibility that the current winner is a private consumer (who removes the current unit from the marketplace). Therefore, it is clear that the presence of the speculator would lower the overall level of consumer bids, but the qualitative characterization of the equilibrium relationship between forward-seen information and current bids shown in Propositions 1 and 2 would remain.

V. Empirical Evidence: Analysis of eBay data

This section uses the predictions highlighted in the end of Section III to construct a statistical test that uses actual eBay bidding data, and attempts to reject the proposed model in favor of one of two simpler models: the model with forward-seeing of only the type-independent summary statistic $H(\Omega)$, and a model without any forward-seeing. Before describing the test and its results, it is necessary to describe some key features of the available data.

Data

EBay kindly provided two datasets, each corresponding to a different product-category: MP3 players and movies on DVD. Both categories involve differentiated mass-produced consumer goods, so consumer preferences are likely to be well-approximated by assumptions of the model as discussed in the beginning of Section III. In particular, private consumers are likely to have unit demand for a specific model of an MP3 player or for a specific title of a movie, and their purchases are likely to be motivated by independent private utilities.

Data from any online-auction site that only facilitates the communication between sellers is bound to come without definitive identification of each item sold. To match each listing to a product type (movie title or player model), researchers have to rely on the item description written by the sellers, and some classification error is inevitable. In both datasets used here, a word-matching algorithm classified about 80 percent of the listings as likely selling a known product-type,[2] but the resulting classification is still only approximate. To refine the classification, a few outlier auctions of each type were removed from the data because their final prices were too far out of line with the bulk of their type, suggesting that they probably actually

[2] In the MP3 player category, 21 percent of all auction listings remained unclassified, either because their description was insufficient for identification ("new cool mp3 player for sale") or because they do not belong to the product category at all ("napster t-shirt" or "128mb sandisk memory card").

sell something else than a single unit of the type.[3] While both datasets are similar to each other in many respects, they differ slightly; the similarities are discussed first. Each dataset contains all submitted bids in each recorded auction, as well as information about each listing including its timing and a text description of the item sold. The bidding data captures all the proxy bids made, including the winning bid which remains undisclosed in reality.[4] Individual bidders and sellers are tracked over time with unique identification numbers. All auctions that involve reserve prices or bid-cancellations are eliminated from the analysis because their modeling is beyond the scope of this paper.

The MP3-player dataset is constructed to capture all the auctions for the top thirty types (models) of MP3 players held during a 4-month period in the beginning of 2001.The top 30 types account for about 91% of the identified products and 70% of all listings. Both Buy-It-Now (BIN) auctions and simple auctions are recorded in the MP3-player dataset. The minority (23 percent) of the BIN auctions that ended early at the BIN price level were excluded from the analysis because modeling the use of the BIN option is beyond the scope of this paper, and because their early termination makes them not useful as forward-seen future options[5]. The remaining BIN auctions that either reverted to simple auctions or remained unsold were retained in the analysis.[6] Of the resulting 6,967 auctions used in the analysis, about 50 percent were originally started as

---

[3] Removing top and bottom three percent of all bids on each type is sufficient to eliminate all bids that are either multiples of the median price on the type, probably indicating an undetected bundle, or that are less than 10 percent of the type's median price, indicating a listing that is just an accessory or that is problematic for reasons unobserved by the analyst but obviously observable to the buyers.

[4] Please see www.ebay.com for a definition of proxy bids and a thorough description of the bidding rules.

[5] Excluding BIN auctions from the analysis makes the remaining auctions seem closer to each other in the sequence, confounding the effect of the next auction on the current bid with the effects of more distant future auctions for some observations. Since effects of more distant future auctions are smaller than the effects of immediately following offerings, excluding BIN auctions biases the estimates of the focal next-auction effect down, making the resulting estimates of all effects' magnitudes conservative.

[6] On eBay, the BIN option disappears when a bid lower than the BIN price is made, and the auction reverts to a simple auction. Therefore, BIN auctions that reverted are indistinguishable from simple auctions to the bidders. BIN auctions that remained without any bids have had at least partial option-value to the bidders, so they are also retained as parts of the auction-stream.

BIN auctions. Listing all of the auctions were 2663 unique sellers. Participating in the 4852 (70 percent) auctions that received bids were 15,574 unique bidders, 3.2 per auction on average. Almost half of the bidders participated in multiple auctions, raising the average number of unique bidders in an auction to 7.5, median of 7.

The movie-dataset is constructed to capture all simple auctions for thirty popular titles in October 2002, where popularity was judged using bestseller lists.[7] BIN auctions were not recorded in the movie-data. The dataset contains 4864 auctions listed by 1607 unique sellers. Participating in the 3384 (69 percent) of auctions that received bids were 7,445 unique bidders, 2.2 per auction on average. About a third of the bidders participated in multiple auctions, resulting in an average number of unique bidders per auction of 3.7, median 3. The movie-market thus involves much lower bidder-competition than the MP3-player market.

Some preliminary evidence that the model is consistent with actual bidding behavior can be gleaned from simple summary statistics of the data: Most eventual winners won only one unit within the data-period (93% in MP3-players and 87% in movies). Yet, a substantial number of bidders participated in more than one auction (43% in MP3-players and 33% in movies). One alternative explanation of multi-auction bidding can be ruled out right away: It does not seem that the multi-auction bidders simply submitted a very low bid initially to learn about the auction process or their true valuation, and only later raised their bid to their "full" willingness to pay. Instead, of the 2276 bidders who bid on the same movie-title at least twice in a row, only 49 percent submitted a higher second bid. The corresponding figure among the 4543 MP3-player bidders who bid on the same player-model at least twice in a row is 59 percent. A more precise test of the proposed model based on the empirical relationship between bids and properties of near-future auctions is described next.

[7] I would like to thank Uri Simonsohn for selecting the popular titles and processing the movie dataset.

Econometric test

On eBay, a bid can only be submitted if it exceeds the highest bid at the moment, so the dataset contains relatively more high bids and relatively fewer low bids than a random sample of willingness-to-bid modeled by the sealed-bid abstraction $\overline{b}(x,v)$. While many latent bids may be truncated by the eBay ascending-auction procedure, two bids in every auction are always recorded – the highest and the second-highest bid in each auction. Therefore, the first- and second-order statistics of the population distribution of bids $\overline{b}(x,v)$ conditional on x, $\overline{b}_{(1)}(x)$ and $\overline{b}_{(2)}(x)$ respectively, are observed in the data without any bias. Since all the model predictions concerning $\overline{b}(x,v)$ are true for all v and valuations are by definition independent of the near-future details, the *qualitative* predictions will be true for the order-statistics of $\overline{b}(x,v)$ as well, and $\overline{b}(x,v)$ can thus be replaced with $\overline{b}_{(1)}(x)$ or $\overline{b}_{(2)}(x)$ in all the model predictions of Section IV.[8] The following linear regression can then be used to test the qualitative predictions about the relationship between bids and both object-types Φ and ending-times Ω of near-future auctions:

$$b_{(m),i} = \alpha_{m,k(i)} + \beta_m H_i + \gamma_m x_{i,k(i)} + \theta_m z_i + \varepsilon_{m,i} \tag{4}$$

where *i* indexes auctions, *k* indexes types, and *m* is the order of the bid-statistic, $\alpha_{m,k(i)}$ is the type-specific fixed-effect, β_m is the effect of the type-independent forward-seeing variable H_i, γ_m is the effect of forward-seeing variables specific to auction *i*'s type *k(i)*:

[8] Note that this approach does not require the knowledge of private valuations because the predictions concern the impact of commonly-known forward-seen types rather than the impact of privately known valuations.

$x_{i,k(i)} \in \left\{ \mathrm{I}\Phi^1_{i,k}, w_{i,k}, \left(\mathrm{I}\Phi^1_{1,i,k}, \mathrm{I}\Phi^1_{2,i,k} ..., \mathrm{I}\Phi^1_{A,i,k} \right) \right\}$, z_i is a vector of control-variables specific to the auction i, and $\varepsilon_{m,i}$ is mean-zero error.

Consistent estimates of all parameters can be obtained by Ordinary Least Squares. Since the three different specifications of $x_{i,k(i)}$ are at least partially correlated, three separate specifications of (4) were estimated for both levels of m and both datasets. To improve the theoretical quality of the linear approximation implicit in (4), the analysis of the MP3-player dataset was further split into two separate analyses because of the high variance in the price across the types. The median type sold for a median price of \$105, so the players were split into 15 “low-priced” players with median prices of less than \$100, and 15 “high-priced” players with median prices above \$100. In each level-specific analysis, the players corresponding to the other price-level are retained as part of the auction-stream, lumped together into the 16-th “other” type. The control variables z_i are discussed next.

Since order-statistics of the bidding distribution are used as dependent variables, the most important control variable is the number of current competing bidders, which varies from auction to auction, and which clearly increases any order-statistic ceteris paribus. The number of bidders in an auction is not perfectly observed because of the truncation issue, but the number of observed unique bidders is likely highly correlated with the true number, and is used throughout. In specifying z_i in equation (4), a fully non-parametric specification using a separate dummy for each number of observed bidders was considered, but it did not contribute much beyond the simple linear effect used in the final analysis.

Besides the influence implied by the proposed model, including current competition also controls for the following alternative explanation of a potential negative correlation between the near-future desirability and order-statistics of current bids: If the bidders, instead of acting sequentially as proposed, randomly chose only one auction of their product-type to bid in and subsequently acted myopically by bidding their valuation, more auctions in a short time-period would imply *both* a more desirable near future, and fewer bidders per auction - hence a lower order-statistic of the bids. Therefore, there would arise a mechanical negative correlation between near-future desirability and the order-statistic of the current bids, and including the current number of bidders as a control is critical to rule out this explanation.

Other control-variables included in z_i were a measure of seller reputation[9] shown by Resnick et al (2002), Wilcox(2000) and many others to have a positive impact on bids, a dummy variable for the description of the unit containing words like "new" or "mint" as a coarse measure of within-type vertical differentiation of the products, and (not included by eBay in the movie dataset) seller-controlled differentiation indicators of the listing itself like "photo included", "bold-type listing", or "gallery listing".

Since only the highest and the second highest bid in each auction are used in the analysis and eBay does not allow a bidder to outbid herself, the two order-statistics correspond to bids submitted by different people, and each is the highest bid in the auction submitted by its respective bidder. The analysis thus resolves the issue of "multiple bidding", i.e. the fact that some eBay bidders submit multiple bids in the same auction, by retaining the highest bid for each bidder as "the" bid of that bidder in the auction.[10]

---

[9] Seller feedback score used in the MP3-player analysis was not included in the movie-dataset, so a dummy indicating whether the seller was an "eBay Top Seller" was used instead to capture the effect of reputation.

[10] Multiple bidding remains an unresolved curiosity of eBay because it cannot constitute equilibrium behavior even if the bidders have a common value (as shown by Bajari & Hortacsu 2003). Instead, multiple bidding has been linked to bidder inexperience by both Roth & Ockenfels (2002) and by Wilcox (2000), and explained at least partially as a "naïve English" bidding scheme that ignores the eBay proxy system and instead bids as if the eBay auctions were open-outcry English auctions and bidders paid their bids. In the datasets used here, multiple bidding is not very prevalent (71 percent of movie-bidders and 46 percent of MP3-player bidders never engage in multi-bidding) and it is also negatively correlated with bidder experience as suggested by previous findings.

In both datasets, special care was taken to exclude bids obviously not made by private bidders modeled by the theory. In the MP3-player data, bids made by sellers (about 2 percent) as well as bids made by bidders who won more that one unit within the data-period (about 12 percent of highest and 7 percent of second-highest) were eliminated, resulting in 4068 highest bids and 4099 second-highest bids. Among the movie-auctions, around 3 percent of both highest and second-highest bids were eliminated because of being made by a seller or by a bidder who won multiple units of the same title, and another approximately 0.4 percent of bids were eliminated because they were made by bidders who bid on too many types, resulting in 3114 highest and 2433 second-highest bids. Please see Table 1 for summary statistics of all the variables in the final datasets used in estimation and additional summaries of the data.

Results

In both datasets and according to both order-statistics of bids, bidders seem to engage in at least one form of forward-seeing bidding. The two datasets are discussed in turn, Table 2 presents the parameter estimates for movies, Table 3 for Low-priced MP3 players, and Table 4 for High-priced MP3 players. In the movie-dataset, all type-specific forward-seeing variables have coefficients γ_m consistent with the theory: waiting times until the next auction of the same type increase bids, the same type offered in the next 5 auctions decreases bids, and the impact of another offering in the near future decreases with the number of intervening auctions. The coefficient β_m on the type-independent variable (number of auctions in the next hour) is not significant but generally negative as predicted. Interestingly, the effects are smaller for highest bids than for second-highest bids. The measured effects on second-highest bids seem quite large, and they have the added relevance of essentially capturing the effects on price[11]. With the average price in the category around $10, the same movie offered in the immediately following auction leads to an average price-reduction of 72 cents, and the same movie offered at least once within the next five auctions reduces price by about 31 cents. All control-variable parameters are significant and have the anticipated signs.

Bids on high-priced MP3-players exhibit large and significant β_m and γ_m consistent with the theory. In the subsample of bids on low-priced players, β_m is still as predicted by theory, but γ_m are not significant. One explanation for this difference is that on lower-priced players, detailed examination of the near future is not worth the effort, and bidders find it sufficient to just glance at the red ink and account for the number of auctions ending in the next hour. This potential explanation is not ruled out by the above case of movies (even cheaper than low-priced players) because there tend to be many more bidders in the MP3-player auctions than in the movie-auctions (median 7 versus 3), and such increased competition exponentially reduces the expected future surplus.

Another interesting property of bidding on the low-priced players comes from one of the model-extensions outlined in Section IV, namely from the model involving bidders desiring multiple types. When the analysis is focused on multi-type bidders and each bidder's desired types are identified as all of the 30 types on which that bidder ever submitted a valid bid, a regression analogous to (4) reveals that bidders submitting $\overline{b}_{(1)}$ and $\overline{b}_{(2)}$ on low-priced MP3 players significantly decrease their bids by $2-$3 whenever a high-priced type they also desire is available within the next five auctions (details not reported). The converse is not true for high

[11] Second-highest bids are just a constant increment different from prices, so their analysis is the same as the analysis of price conditional on there being at least two bidders in the auction.

bidders on high-priced players and, as predicted, this model-extension is also not empirically supported in the movie dataset. These results illustrate the richness of information contained in forward-seeing behavior, but their further development is beyond the scope of this paper.

The β_m effects of the number of auctions in the next hour are small in both MP3-player sub-samples: doubling of the number of auctions ending in the next hour is associated with approximately 2 percent reduction in prices[12]. On the other hand, the type-specific effects γ_m on high-priced players are substantial: On roughly $180 items, the same product being available within the next 5 auctions reduces prices by $8 (4.4 percent) on average, $10 (5.6 percent) when the same product is available in the immediately subsequent auction. Analogously, delaying the next offering of the same product by a mere hour from the average of 53 minutes is correlated with an increase in bids of over $2. Note that all these estimates are actually conservative because they suffer from the errors-in-variables problem and are hence biased towards zero. All control-variable parameters have the anticipated signs, but unlike in the movie dataset, some variables are insignificant, notably most of the listing-variables under seller control like photo, bold and gallery. The insignificance does not imply that these instruments are of no value because their usage is endogenous.

[12] log(2) β_m /mean(price) is -0.025 in low-priced and -0.019 in high-priced players.

V. Discussion

This paper documents what happens when the role of an auction changes from selling unique objects at Sotheby's to driving large sequential markets for consumer durables on eBay and other online auction sites. In such markets, seemingly independent auctions become linked through the demand-side strategies. When participating in a sequence of auctions for substitutes, rational forward-looking bidders reduce their bids in anticipation of future auctions offering the same products. When details of some future offerings are already common knowledge as near-future offerings are on eBay, the bidders base their bids on the available information. This paper proposed a new model of such *forward-seeing* bidding, and found support for the predicted behavior in eBay data. The model departs from previous models of sequential auctions by assuming that bidders know not only the type of the current product they are bidding on, but also what and when will be sold next. Then, all rational bidders reduce their current bids when there are more near-future auctions in general, as well as when more of those auctions offer units of the good specifically desired by the bidders and when those units are offered sooner. These findings contribute to the auction-theory literature, and they have obvious relevance to bidders in sequential auctions on eBay or elsewhere.

In both eBay product-categories studied (MP3 players and DVD movies), a test of the model predictions rejects the alternative simpler model without forward-seeing. Moreover, in both categories, at least some bidders seem to take detailed information about the near future into account, leading to price-reductions between three and seven percent whenever the same type of good is available in the next five auctions. The empirical evidence therefore suggests that the bidders behave consistently with the proposed model, and that such behavior does have a sizeable economic impact on the seller revenues.

It is interesting to contrast the demand-side of a sequential auction market with a demand-side of a traditional posted-price market. In both marketplaces, consumers engage in sequential search. The difference is not just in the increased rapidity and built-in non-reversibility of the implied "search" in the auction marketplace. This paper demonstrates that in the online-auction marketplace, useful information about the other (future) purchase opportunities is available, and this information enters current observed demand, effectively blending elements of simultaneous choice among several purchase opportunities in the the underlying sequential search. This paper provides the first step, but more research is clearly needed to see exactly how consumers should and do cope with this new shopping environment.

The empirical findings give focus to future modelers of online auction marketplaces by providing a fairly high lower bound on the sophistication of eBay bidders: eBay bidders seem to look beyond a single auction as they should, and they seem to take what they see into account consistently with a theory of rational bidding. The eBay markets for MP3 players and DVD movies are therefore examples of large internet-auction markets, in which stand-alone analysis of individual auctions would be inappropriate. Instead, this paper demonstrates that individual auctions need to be interpreted within their context of other auctions selling similar objects, and provides a model that can be used to achieve such analysis. The model implies that observed bids are actually always negatively biased measures of true valuations because winning involves an additional opportunity cost arising from not participating in future auctions for the same good.

Because the present results only provide a lower bound on bidder sophistication, there is a lot of room for further empirical modeling of buyer behavior in sequential auction marketplaces. For example, it may be possible to extend structural estimation methodologies like Jofre-Bonet and Pesendorfer (2003) to study exact properties of auction-market demand. The findings reported here may also have impact on seller strategies, raising the question of the scope of auction-driven markets, i.e. whether the inter-auction competition found here limits the potential of sequential auctions as trading institutions. Throughout this paper, the seller was assumed to be exogenous, but allowing for strategic selling may both qualitatively and quantitatively change the bidder's strategy. A companion paper (Zeithammer 2004) provides a model of such strategic sellers facing forward-looking buyers, showing that the scope of auctions may not necessarily be limited by competition arising from forward-looking bidders.

Table 1: Summary statistics

	Movies				Low-priced MP3 players				High-priced MP3 players			
	Highest bid (N=3113)		Second highest bid (N=2431)		Highest bid (N=1693)		Second highest bid (N=1646)		Highest bid (N=2372)		Second highest bid (N=2451)	
	Mean	*Std*	Mean	*Std*	Mean	*Std*	Mean	*Std*	Mean	*Std*	Mean	*Std*
log (Seller Reputation + 6)	n/a	*n/a*	n/a	*n/a*	4.803	*1.742*	4.709	*1.715*	4.431	*1.417*	4.463	*1.402*
top-seller dummy	0.466	*0.499*	0.466	*0.499*	n/a	*n/a*	n/a	*n/a*	n/a	*n/a*	n/a	*n/a*
photo-listing dummy	n/a	*n/a*	n/a	*n/a*	0.745	*0.436*	0.765	*0.424*	0.699	*0.459*	0.699	*0.459*
bold-listing dummy	n/a	*n/a*	n/a	*n/a*	0.034	*0.182*	0.037	*0.189*	0.045	*0.208*	0.047	*0.211*
gallery-listing dummy	n/a	*n/a*	n/a	*n/a*	0.094	*0.292*	0.095	*0.294*	0.069	*0.254*	0.068	*0.252*
new dummy	0.441	*0.497*	0.429	*0.495*	0.569	*0.495*	0.554	*0.497*	0.667	*0.471*	0.668	*0.471*
current competition (# unique bidders)	3.683	*2.455*	4.450	*2.238*	7.080	*3.969*	7.845	*3.493*	9.341	*4.084*	9.659	*3.862*
log(# auctions next hour+1)	2.403	*0.671*	2.366	*0.669*	1.469	*0.599*	1.478	*0.590*	1.580	*0.586*	1.576	*0.586*
log(time until next+1)	3.179	*0.885*	3.270	*0.807*	4.047	*0.880*	4.043	*0.882*	3.714	*1.040*	3.698	*1.043*
dummy (same type next 5)	0.178	*0.383*	0.172	*0.377*	0.299	*0.458*	0.291	*0.454*	0.497	*0.500*	0.505	*0.500*
dummy (same type 1 auction from now)	0.065	*0.246*	0.056	*0.231*	0.129	*0.335*	0.123	*0.329*	0.223	*0.416*	0.236	*0.425*
dummy (same type 2 auctions from now)	0.027	*0.161*	0.027	*0.163*	0.047	*0.212*	0.048	*0.214*	0.116	*0.320*	0.116	*0.320*
dummy (same type 3 auctions from now)	0.028	*0.165*	0.026	*0.159*	0.048	*0.215*	0.050	*0.218*	0.073	*0.261*	0.068	*0.252*
dummy (same type 4 auctions from now)	0.037	*0.188*	0.037	*0.189*	0.039	*0.194*	0.037	*0.188*	0.051	*0.219*	0.047	*0.212*
dummy (same type 5 auctions from now)	0.023	*0.148*	0.025	*0.156*	0.035	*0.185*	0.033	*0.180*	0.035	*0.184*	0.038	*0.190*

Notes to Table 1:
In the movie-data, the shares of each type range fairly continuously from 1.5% to 14.1% (Black Hawk Down).
In the MP3-player data, each both low- and high-priced players have dominant products, KB Gear Jamp3 (27%) in the low-priced category and Diamond Rio 500 (48%) in the high-priced category. The shares of the remaining products are all below 10% and decline fairly continuously to 1% for the 30-th product.

In the raw MP3 player dataset, the proportions of listings by selling-format (BIN auction vs. simple auction) are as follows: of 100 listings, 18 involve a reserve price and 12 end up being sold by BIN price (57 start with a BIN price). The listings with a reserve price and those that end with a BIN price are eliminated since their future option-value is more complicated than that of a simple auction modeled by the theory. The remaining 70 enter the empirical analysis. The raw movie dataset does not include BIN auctions or auctions with a reserve price, making an analogous breakdown by selling format unavailable.

Table 2: Estimation results: movies

	Specification 1		*Specification 2*		*Specification 3*	
	Highest bid	**2nd highest bid**	**Highest bid**	**2nd highest bid**	**Highest bid**	**2nd highest bid**
Variable	Estimate *(t-stat)*	Estimate *(t-stat)*	Estimate *(t-stat)*	Estimate *(t-stat)*	Estimate *(t-stat)*	Estimate *(t-stat)*
α (30 type-specific dummies)	suppressed for parsimony (mean 7.92, standard deviation 1.39, minimum 5.40, maximum 10.9)					
θ (top-seller dummy)	0.645 *(9.30)*	0.575 *(8.51)*	0.637 *(9.33)*	0.561 *(8.41)*	0.628 *(9.18)*	0.553 *(8.31)*
θ (new dummy)	0.756 *(10.63)*	0.826 *(11.87)*	0.747 *(10.70)*	0.817 *(11.93)*	0.748 *(10.71)*	0.825 *(12.05)*
θ (current competition)	0.128 *(8.81)*	0.094 *(6.20)*	0.127 *(8.89)*	0.096 *(6.38)*	0.125 *(8.71)*	0.094 *(6.26)*
β (log (# next hour+1))	**-0.045** *-(0.83)*	**0.033** *(0.63)*	**-0.087** *-(1.74)*	**-0.039** *-(0.80)*	**-0.084** *-(1.68)*	**-0.039** *-(0.81)*
γ (log time until next)	**0.06** *(2.35)*	**0.108** *(4.07)*				
γ (same type next 5 auctions)			**-0.17** *-(1.92)*	**-0.313** *-(3.51)*		
γ (same type 1 a. from now)					**-0.348** *-(2.52)*	**-0.722** *-(5.00)*
γ (same type 2 a. from now)					**-0.433** *-(2.10)*	**-0.385** *-(1.93)*
γ (same type 3 a. from now)					**0.136** *(0.67)*	**0.098** *(0.48)*
γ (same type 4 a. from now)					**-0.051** *-(0.29)*	**-0.182** *-(1.05)*
γ (same type 5 a. from now)					**0.07** *(0.31)*	**0.025** *(0.12)*
	N=3017 $R^2=0.42$	N=2356 $R^2=0.53$	N=3113 $R^2=0.42$	N=2431 $R^2=0.53$	N=3113 $R^2=0.42$	N=2431 $R^2=0.53$

Note to Table 2: Parameters of primary interest are shown in bold. 30 movie-title fixed-effects are suppressed for parsimony. The first model has a slightly smaller sample-size because calculating time until next auction of the same type for every type requires a longer forward-seeing horizon and hence there is more truncation of forward-seeing information in the end of the data-period.

Table 3: Estimation results: Low-priced MP3 players

	Specification 1				*Specification 2*				*Specification 3*			
	Highest bid		**2nd highest bid**		**Highest bid**		**2nd highest bid**		**Highest bid**		**2nd highest bid**	
Variable	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*
α (15 type-specific dummies)	suppressed for parsimony (mean 65-70, standard deviation 19-20, minimum 40, max 103-108)											
θ (log (Seller Reputation+6))	0.869	*(2.92)*	0.038	*(0.16)*	0.803	*(2.79)*	0.041	*(0.17)*	0.799	*(2.77)*	0.037	*(0.16)*
θ (photo-listing dummy)	1.855	*(1.63)*	0.207	*(0.23)*	2.04	*(1.83)*	0.428	*(0.48)*	2.065	*(1.85)*	0.483	*(0.53)*
θ (bold-listing dummy)	-1.318	*-(0.53)*	-1.673	*-(0.88)*	-0.996	*-(0.41)*	-1.664	*-(0.89)*	-0.996	*-(0.40)*	-1.605	*-(0.86)*
θ (gallery-listing dummy)	4.633	*(3.03)*	4.322	*(3.56)*	4.339	*(2.90)*	4.092	*(3.43)*	4.331	*(2.89)*	4.088	*(3.42)*
θ (new dummy)	2.951	*(3.03)*	4.264	*(5.53)*	3.215	*(3.36)*	4.378	*(5.75)*	3.217	*(3.36)*	4.369	*(5.73)*
θ (current competition)	0.28	*(2.46)*	0.537	*(5.37)*	0.287	*(2.57)*	0.544	*(5.51)*	0.286	*(2.56)*	0.544	*(5.51)*
β (# next hour)	**-2.578**	*-(3.48)*	**-2.392**	*-(4.00)*	**-2.587**	*-(3.69)*	**-2.576**	*-(4.55)*	**-2.596**	*-(3.69)*	**-2.596**	*-(4.57)*
γ (log time until next)	**0.048**	*(0.15)*	**0.176**	*(0.70)*								
γ (same type next 5 auctions)					**-0.965**	*-(1.00)*	**-0.358**	*-(0.46)*				
γ (same type 1 a. from now)									**-0.454**	*-(0.34)*	**0.296**	*(0.27)*
γ (same type 2 a. from now)									**-1.376**	*-(0.68)*	**-0.359**	*-(0.23)*
γ (same type 3 a. from now)									**-1.436**	*-(0.73)*	**-0.221**	*-(0.14)*
γ (same type 4 a. from now)									**-1.586**	*-(0.73)*	**-1.914**	*-(1.07)*
γ (same type 5 a. from now)									**-0.88**	*-(0.39)*	**-1.046**	*-(0.56)*
	N=1645	R^2=0.63	N=1600	R^2=0.73	N=1693	R^2=0.63	N=1646	R^2=0.72	N=1693	R^2=0.63	N=1646	R^2=0.72

Note to Table 3: Parameters of primary interest are shown in bold.Please also see note to Table 2 for an explanation of different sample-sizes.

Table 4: Estimation results: High-priced MP3 players

	Specification 1				*Specification 2*				*Specification 3*			
	Highest bid		**2nd highest bid**		**Highest bid**		**2nd highest bid**		**Highest bid**		**2nd highest bid**	
Variable	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*	Estimate	*(t-stat)*
α (15 type-specific dummies)	suppressed for parsimony (mean 186-171, standard deviation 57-58, minimum 99-114, max 316-334)											
θ (top-seller dummy)	0.729	*(1.67)*	1.332	*(3.41)*	0.816	*(1.89)*	1.405	*(3.60)*	0.787	*(1.82)*	1.392	*(3.57)*
θ (photo-listing dummy)	0.009	*(0.01)*	0.977	*(0.85)*	-0.162	*-(0.13)*	0.746	*(0.65)*	-0.102	*-(0.08)*	0.852	*(0.74)*
θ (bold-listing dummy)	4.427	*(1.50)*	2.597	*(1.00)*	4.369	*(1.49)*	2.898	*(1.12)*	4.176	*(1.42)*	2.693	*(1.04)*
θ (gallery-listing dummy)	1.088	*(0.45)*	-1.11	*-(0.51)*	0.966	*(0.40)*	-1.228	*-(0.57)*	1.146	*(0.48)*	-1.171	*-(0.54)*
θ (new dummy)	7.112	*(5.07)*	7.131	*(5.65)*	7.259	*(5.21)*	7.112	*(5.62)*	7.117	*(5.10)*	6.918	*(5.47)*
θ (current competition)	0.582	*(3.90)*	0.646	*(4.64)*	0.536	*(3.62)*	0.634	*(4.55)*	0.538	*(3.64)*	0.64	*(4.60)*
β (# next hour)	**-5.057**	*-(4.58)*	**-3.542**	*-(3.61)*	**-6.755**	*-(6.61)*	**-5.783**	*-(6.35)*	**-6.577**	*-(6.41)*	**-5.522**	*-(6.05)*
γ (log time until next)	**2.617**	*(5.51)*	**3.232**	*(7.73)*								
γ (same type next 5 auctions)					**-7.441**	*-(4.89)*	**-8.172**	*-(6.07)*				
γ (same type 1 a. from now)									**-8.83**	*-(4.95)*	**-10.248**	*-(6.57)*
γ (same type 2 a. from now)									**-7.624**	*-(3.55)*	**-7.696**	*-(4.04)*
γ (same type 3 a. from now)									**-8.364**	*-(3.37)*	**-8.959**	*-(3.97)*
γ (same type 4 a. from now)									**-3.214**	*-(1.13)*	**-4.832**	*-(1.86)*
γ (same type 5 a. from now)									**-4.023**	*-(1.22)*	**-1.389**	*-(0.49)*
	N=2317	R^2=0.86	N=2393	R^2=0.88	N=2372	R^2=0.85	N=2451	R^2=0.87	N=2372	R^2=0.85	N=2451	R^2=0.87

Note to Table 4: Parameters of primary interest are shown in bold.Please also see note to Table 2 for an explanation of different sample-sizes.

Appendix: Proofs of propositions

Proof of Proposition 1: First, consider $b\left(0,\varphi_1,\omega_1,v\right)$: Since the expected surplus function is obviously positive, and bidding any positive amount on a personally worthless object yields a negative current-period payoff, any positive bid is dominated by a zero bid. Second, consider $b\left(1,\varphi_1,\omega_1,v\right)$: As long as $\frac{\partial S\left(1,\varphi_1,\omega_1,v\mid c_0\right)}{\partial c_0} > -\frac{1}{\left(\lambda\delta\right)^{\omega_1}}$, a solution to the optimal bidding problem is characterized by first-order condition in (2) because the problem is concave at the solution to (2). Moreover, the solution to (2) is unique for every φ_1,ω_1 and v because for all c_0: $\frac{\partial}{\partial c_0}\left(v-\left(\lambda\delta\right)^{\omega_1} S\left(1,\varphi_1,\omega_1,v\mid c_0\right)\right)<1$, and since $v-\left(\lambda\delta\right)^{\omega_1} S\left(1,\varphi_1,\omega_1,v\mid c_0=0\right)>0$ and $v-\left(\lambda\delta\right)^{\omega_1} S\left(1,\varphi_1,\omega_1,v\mid c_0=1\right)<1$, it follows by continuity of S and the Intermediate Value Theorem, there is exactly one $b\left(1,\varphi_1,\omega_1,v\right)$ that satisfies (2). To show that $\left|\frac{\partial S\left(1,\varphi_1,\omega_1,v\mid c_0\right)}{\partial c_0}\right|$ is sufficiently uniformly bounded to ensure concavity (and also existence of S using Schauder Fixed point theorem) let the steady-state order-statistics of the current competitive bids be $Y_{(1)}>Y_{(2)}>\ldots>Y_{(N-1)}$, and note that $\left|\frac{\partial S\left(1,\varphi_1,\omega_1,v\mid c_0\right)}{\partial c_0}\right|<\left|\frac{\partial \tilde{S}\left(1,1,\omega_1,1\mid Y_{(1)}=c_0\right)}{\partial c_0}\right|$, where $\tilde{S}\left(1,\varphi_1,\omega_1,v\mid c_0\right)$ is the future surplus conditional on $Y_{(2)}$ surviving until the next auction for sure (as opposed to with probability λ^{ω_1}). The impact of today's winning bid on tomorrow's competition is clearly the highest when today's highest loser survives for sure, and the impact increases in the chance of winning tomorrow's auction because of attrition and discounting, and hence is bounded by the impact on the highest possible bidder v=1. Let z be the expected maximum bid of new entrant(s). Then, $\tilde{S}\left(1,1,\omega_1,1\mid c_0\right)$ can be evaluated as $\tilde{S}\left(1,1,\omega_1,1\mid c_0\right)=E_{\varphi_2,\omega_2}\left[1-E\left(c_1\mid c_0\right)\right]=E_{\varphi_2,\omega_2}\left[1-E\left(\max\left(Y_{(2)},z\right)\mid Y_{(1)}=c_0\right)\right]$, and it is clear that the impact of c_0 on $\tilde{S}\left(1,1,\omega_1,1\mid c_0\right)$ is limited by the slope of the expectation of the second order-statistic in the first-order statistic, which is less than unity for any distribution by elementary statistics: $\left|\frac{\partial \tilde{S}\left(1,1,\omega_1,1 \mid Y_{(1)}=c_0\right)}{\partial c_0}\right| < \left|\frac{\partial E\left(Y_{(2)} \mid Y_{(1)}=c_0\right)}{\partial c_0}\right| < 1$. This concludes the proof that there is a well-defined pure-strategy characterized by (2). Equation (3) is the Bellman equation that a steady-state strategy must satisfy to be perfect. The assumption of no memory allows future surpluses to depend only on future information. *QED*

Proof of Proposition 2:

The claim that $b(1,\varphi_1,\omega_1,v)$ increases in ω_1, follows immediately from the optimal bidding condition (2). The claim that $b(1,\varphi_1,\omega_1,v) < v$ follows from the obvious fact that S is positive. The claim that $0 < b(1,\varphi_1,\omega_1,v)$ follows from the fact that $S(\varphi_0,\varphi_1,\omega_1,v \mid c_0) < v$, which is always true because the bidder cannot get more utility than that from getting a unit of the desired type for free, for sure, and immediately, and hence receiving a surplus of v. The central claim of the first part of Proposition 2, $b(1,1,\omega_1,v) < b(1,0,\omega_1,v)$, hinges on the fact that bidding on a desired type always gives the bidder at least as much surplus, ceteris paribus, as bidding on an undesired type, and will be shown in two steps: 1) for every c_0 $S(1,0,\omega_1,v \mid c_0) < S(1,1,\omega_1,v \mid c_0)$ and 2) to show 1) implies $b(1,0,\omega_1,v) > b(1,1,\omega_1,v)$.

To show 1), it is instructive to write down the four Bellman equations characterizing the steady-state expected-surplus functions in all possible combinations of current and future desirability states, keeping timing ω_1=ω_2=1 constant and suppressing it from all equations:

$$S(1,1,v \mid c_0) = E_{\varphi_2}\left[\int^{b(1,\varphi_2,v)} (v-c_1)\,dG(c_1 \mid c_0,1,1,\varphi_2) + \lambda\delta \int_{b(1,\varphi_2,v)} S(1,\varphi_2,v \mid c_1)\,dG(c_1 \mid c_0,1,1,\varphi_2)\right]$$

$$S(0,1,v \mid c_0) = E_{\varphi_2}\left[\int^{b(1,\varphi_2,v)} (v-c_1)\,dG(c_1 \mid c_0,0,1,\varphi_2) + \lambda\delta \int_{b(1,\varphi_2,v)} S(1,\varphi_2,v \mid c_1)\,dG(c_1 \mid c_0,0,1,\varphi_2)\right]$$

$$S(1,0,v \mid c_0) = E_{\varphi_2}\left[\lambda\delta \int S(0,\varphi_2,v \mid c_1)\,dG(c_1 \mid c_0,1,0,\varphi_2)\right]$$

$$S(0,0,v \mid c_0) = E_{\varphi_2}\left[\lambda\delta \int S(0,\varphi_2,v \mid c_1)\,dG(c_1 \mid c_0,0,0,\varphi_2)\right]$$

The G distribution arises from all the surviving losers of the current auctions as well as from all the new entrants, and since the number of each is random, G is not simple to evaluate. However, it will always be true that the expected competition after a desired type is sold today is always slightly weaker than when today's type is not desired, because a trade means that the highest competing bidder certainly exited the bidder pool while no trade means that the highest competing bidder only exited the bidder pool with probability *(1-λ)*. Therefore, $G(c_1 \mid c_0, 0, \varphi_1, \varphi_2) < G(c_1 \mid c_0, 1, \varphi_1, \varphi_2)$, and since S decreases in c as shown in the proof to Proposition 1, $S(0, \varphi_1, v \mid c_0) < S(1, \varphi_1, v \mid c_0)$. On the other hand, $G(c_1 \mid c_0, 1, 1, \varphi_2) = G(c_1 \mid c_0, 1, 0, \varphi_2)$ because the bidding function is increasing in v, and so φ_1 has no differential impact on the kind of bidders likely to survive from the past period 0. Therefore, I can write the key difference between surpluses as:

$$S(1,1,v \mid c_0) - S(1,0,v \mid c_0) = E_{\varphi_2}\left[\int^{b(1,\varphi_2,v)} \left(v - c_1 - S(0,\varphi_2,v \mid c_1)\right) dG(c_1 \mid c_0, \varphi_2)\right] +$$

$$+\lambda\delta \int_{b(1,\varphi_2,v)} \left[S(1,\varphi_2,v \mid c_1) - S(0,\varphi_2,v \mid c_1)\right] dG(c_1 \mid c_0, \varphi_2)$$

, and this difference is positive because $S(0,\varphi_1,v \mid c_0) < S(1,\varphi_1,v \mid c_0)$ and because $\left[v - c_1 - S(1,\varphi_2,v \mid c_1)\right] > 0$ for all $c_1 < b(1,\varphi_2,v)$ which follows from the single-crossing property discussed in the proof of Proposition 1. Step 2) also follows from the single-crossing property of the first-order condition: Since $\left[v - c_1 - S(1,\varphi_2,v \mid c_1) = 0\right]$ has a unique solution and $\left[v - 0 - S(1,\varphi_2,v \mid 0) > 0\right]$, $b_0 = b(1,0,v)$ implies $\left[v - b_0 - S(1,1,v \mid b_0) < 0\right]$ and hence the point b_1 such that $\left[v - b_1 - S(1,1,v \mid b_1) = 0\right]$ must lie to the left of b_0.

The fact that $b(1,\varphi_1,\omega_1,v)$ decreases in ρ follows from differentiation of the Bellman equation (3) after writing the expectation $E_{\varphi_2}[.]$ as $\rho(\varphi_2=1)+(1-\rho)(\varphi_2=0)$ and noting by an argument analogous to the one above that $S(\varphi_0, 0, \omega_1, v \mid c_0) < S(\varphi_0, 1, \omega_1, v \mid c_0)$, i.e. whatever today's type, it is always better to face a desired type tomorrow than not. Since higher ρ increases the chance of $\varphi_2=1$, the result follows. *QED*

References

Ariely, Dan and Simonson, Itamar (2003), "Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions," *Journal of Consumer Psychology*, vol. 13, 113-123.

Ashenfelter, Orley (1989), "How Auctions Work for Wine and Art," *Journal of Economic Perspectives*, vol. 3, No. 3, 23-36.

Bajari, Patrick and Hortacsu, Ali (2003), "Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, vol. 34, No. 2.

Dholakia, Utpal (Paul) and Soltysinski, Kerry (2001), "Coveted or Overlooked? The Psychology of Bidding for Comparable Listings in Digital Auctions," *Marketing Letters*, vol. 12:3, 225-237.

Engelbrecht-Wiggans, Richard (1994), "Sequential Auctions of Stochastically Equivalent Objects," *Economic Letters*, vol. 44, 87-90.

Engelbrecht-Wiggans, Richard (1993), "Optimal Auctions Revisited," *Games and Economic Behavior*, vol. 5, 227-239.

Gale, Ian L. and Hausch, Donald B. (1994), "Bottom-fishing and declining prices in sequential auctions," *Games and Economic Behavior*, vol. 7, 318-331.

Jofre-Bonet, Mireia and Pesendorfer, Martin (2003), "Estimation of a Dynamic Auction Game," *Econometrica*, vol. 71, 1443-1489.

Milgrom, Paul R. and Weber, Robert J. (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, vol. 50, 1089-1122.

Milgrom, Paul R. and Weber, Robert J. (1999), "A Theory of Auctions and Competitive Bidding II," in *The Economic Theory of Auctions*, Paul Klemperer (ed.), Edward Elgar Publishing.

Resnick, Paul, Zeckhauser, Richard, Swanson, John, and Lockwood, Kate (2002), "The Value of Reputation on eBay: A Controlled Experiment," working paper, University of Michigan.

Roth, Al and Ockenfels, Axel (2002), "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, Vol. 92, No. 4, pp. 1093-1103.

Weber, Robert J. (1983), "Multi-object auctions," in *Auctions, Bidding and Contracting: Uses and Theory*, eds. Engelbrecht-Wiggans, Shubik and Stark, New York University Press.

Wilcox, Ronald T. (2000), "Experts and Amateurs: The Role of Experience in Internet Auctions", *Marketing Letters*, Vol. 11(4), 363-74.

Zeithammer, Robert (2004), "An Equilibrium Model of a Dynamic Auction Marketplace," working paper, Graduate School of Business, University of Chicago.