

ROBERT ZEITHAMMER*

At Internet auction sites, such as 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 this information into account. This article 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.

Forward-Looking Bidding in Online Auctions

Internet auction sites, such as eBay, are increasingly being used to sell mass-produced consumer durables. The largest eBay categories in terms of dollars are cars, consumer electronics, computers, clothing/accessories, and books/movies/music (according to eBay's 2004 company report). Because 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 other, subsequent auctions. Because online auctions are usually listed for several days before concluding, detailed information about what will be sold and when it will be sold in the near future is available to bidders. Two important questions arise: How should rational consumers interested in buying just one unit of the good use such information in forming their bids? and Do eBay bidders actually use the information accordingly? To answer these questions, this article 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 and that each bidder has a unit demand for only one type of good. For example, a consumer may be shopping for one DVD of his or her favorite movie or for one unit of a specific brand and model of an MP3 player. Therefore, each bidder faces a trade-off between winning now and winning later. This trade-off arises from individual desired units being perfect substitutes; that is, the winner of each auction exits the marketplace and thus forgoes the expected surplus from participating in future auctions that also sell the desired good, pos-

sibly for a lower price. The current winner's forgone future surplus is an opportunity cost of winning now. Thus, 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 their 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, thus, the opportunity cost of winning now is a function of the available information about what will be sold in the near future and when. The equilibrium analysis of the game with forward-seeing bidders is complicated; the expected surplus function depends on the bidding strategy, and the bidding strategy depends on the expected surplus function. Although the equilibrium bidding strategy is intractable in closed form, this article shows that there is 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 the bidders actually use (i.e., how sophisticated they are in accounting for the information). 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 Engelbrecht-Wiggans's (1994) model, 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,

*Robert Zeithammer is Assistant Professor of Marketing, Graduate School of Business, University of Chicago (e-mail: rzeitham@chicagoGSB.edu). The author thanks John Little, Birger Wernerfelt, Elie Ofek, Dan Ariely, Sergei Izmalkov, Jeroen Swinkels, Jean-Pierre Dube, Uri Simonsohn, the two anonymous *JMR* reviewers, and seminar participants at several marketing departments around the country, where he presented a version of this article in 2002.

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. Therefore, an intermediate level of sophistication involves bidders who use 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 this decreases the expected waiting time until another unit of their desired types comes up for sale, thus increasing the expected future surplus. Third, each bidder can actually examine the near-future auctions closely and base his or her bidding strategy on both the timing and the types of objects that actually come up for sale. The opportunity cost of winning today becomes a function of personal preferences for the future items; the sooner personally desirable products come up, the higher are the opportunity costs and the lower are the bids.

How much detail of the available information about near-future auctions eBay bidders actually use is an empirical question, and this article 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 previously described nested restrictions of the most-sophisticated bidding function. Two different data sets are used, one from the MP3-player category, in which each player brand-model combination is considered a different type, and one from the DVD movie category, in which different product types are 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 account for their personal preferences for specific future products. Therefore, the model of forward-looking behavior proposed here is relevant to understanding the demand side of auction-driven marketplaces, such as eBay.

The article is structured as follows: After a brief review of the relevant literature, I present the model that constitutes the main theoretical contribution of this article. Next, I discuss the robustness of the model predictions relative to perturbations in the assumptions. Then, I present the empirical test and conclude with a discussion of the implications of the findings for both researchers and participants of online auction marketplaces.

LITERATURE REVIEW

The theoretical model proposed here is not confined to online auctions; it 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 article investigates the influence of public information about near-future auctions on infinite-horizon steady-state bidding. The model extends Milgrom and Weber’s (2000) finite-horizon and identical-goods model to an infinite horizon and horizontally differentiated goods model. A 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. Thus, the proposed model is the simplest model that involves unit-demand bidders and nontrivial information about the near-future auctions. Because the model considers a sequence of auctions for nonidentical 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 as a result of the simplifying assumption that the product heterogeneity is 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 multiauction bidding has not been addressed, except in the work of Bajari and Hortacsu (2003), who study bidder entry in common-value auctions, and Dholakia and Soltysinski (2001), who find a “herding bias” (i.e., consumers flocking to popular auctions despite the existence of other auctions for substitute items). The herding bias is especially relevant to the current work because it provides another layer of behavioral complexity beyond the rational behavior described here.

On the empirical front, the most related work is a recent article 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. However, forward-seeing behavior is not part of their model, because they assume that bidders do not account for public information about upcoming auctions. Another important difference is that the current article conducts an empirical test nonparametrically without functional-form assumptions that would be necessary for a structural econometric model.

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 (for a discussion, see Ariely and Simonson 2003). 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 because bidding on eBay both should and does tend to happen at the end of each auction, not giving the competitive bidders time to react to one another’s bids (Roth and Ockenfels 2002). To approximate the price determination in 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.

Because 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 begins, and several upcoming auctions are already known. The eBay Web page design reinforces this conceptualization by listing auctions in a sequence ordered by end-

ing time and by allowing bidders to place known future auctions on their private “watch lists.” Anecdotal 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 to be private and independent across bidders. This is a reasonable model of private-consumer utility in the economically largest eBay categories that involve mass-produced consumer durables that are usually purchased for private use and depreciate quickly because of obsolescence. In particular, the assumption is reasonable for MP3 players and DVDs, which I consider in the empirical section. The independence assumption resonates with the sealed-bid abstraction outlined previously because bidders with independent preferences do not try to learn about their own preferences from other people’s bids.

An additional simplifying assumption is necessary for a tractable model, namely, that bidders do not consider past prices. If they did, because past prices are upper bounds on the past bids of competing bidders who did not win the past auction and thus 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 from the bidders whose bids were less than the past price. This asymmetry would escalate over time, as Milgrom and Weber (2000) point out. Assuming the effects of past prices away does not have a large 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 One-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, and the probability of type k is captured by rate ρ_k . Each bidder desires one of the first K product types in the sense that all nondesired types give the bidder zero utility, whereas 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 (the category of MP3 players is differentiated by brand-model combinations, such as “Rio 500,” and the category of movies on DVD is differentiated by movie title). No bidders desire the last $K + 1$ product type, which captures various suspect “free” offers and poorly described, misplaced goods that are bound to clutter any marketplace.

An arbitrary “current” auction is selected as the origin of indexing, so the current auction has an index of 0, the

immediately following auction has an index of 1, the auction after that has an index of 2, and so forth. To capture the type information, let $\phi_{j,k} \in \{0, 1\}$ be the indicator function equal to one when auction j is of type k ; otherwise, let it equal zero. 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 $\phi_{0,k}$ but also the information $(\phi_{1,k}, \omega_1)$ that arises from seeing forward, namely, whether they desire the next product and how far in the future the next sale will occur. The fundamental difference between the two kinds of forward-seeing information is that $\phi_{1,k}$ is inherently type-specific, whereas ω_1 is the same for all types.

In each auction of type k , N_k bidders participate. Bidder i of type k considers all the desired (i.e., type- k) objects identical and can derive a private value of $v_{i,k}$ dollars from any 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_k . Therefore, the private valuation to bidder i of type k of the current object is $\phi_{0,k}v_{i,k}$, and the private valuation to the same bidder of the next object is $\phi_{1,k}v_{i,k}$, where $v_{i,k} \sim F_k$. All bidders can derive utility only from a single unit of their desired good, so when a bidder owns one unit of the desired type, all subsequent units are worth zero to that bidder. The bidders have no memory, so they cannot base their actions on outcomes of previous auctions.

Assume that 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 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 One-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 $(\phi_{0,k}, \phi_{1,k}, \omega_1)$ and on each bidder’s private information $v_{i,k}$. Because 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 $\{1, 2, \dots, k - 1, k + 1, \dots, K, K + 1\}$ can all be lumped into “other” undesired type. Without loss of generality, I can then solve for the equilibrium bidding strategy in the case $K = 1$, suppressing all k subscripts for clarity.

Let $K = 1$, and let $S(\phi_0, \phi_1, \omega_1, v_i | c_0)$ be 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 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 $\phi_0 = 1$. 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 $b(\varphi_0, \varphi_1, \omega_1, v)$:

$$(1) \quad b(\varphi_0, \varphi_1, \omega_1, v) = \arg \max_{\beta \geq 0} \int_0^\beta (\varphi_0 v - c_0) dG(c_0) \\ + (\delta\lambda)^{\omega_1} \int_\beta^1 S(\varphi_0, \varphi_1, \omega_1, v | c_0) dG(c_0).$$

In equilibrium, the expected surplus function must account for other bidders 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 the infinite-horizon setting employed here to capture a mature, ongoing market, future bidders again reduce their bids as a function of their respective future, and at least some of these future bidders will be current competitors who lost the current 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 P_1 :

P_1 (proof in the 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

$$(2) \quad b(1, \varphi_1, \omega_1, v) = v - (\delta\lambda)^{\omega_1} S[1, \varphi_1, \omega_1, v | c_0 = b(1, \varphi_1, \omega_1, v)] \\ b(0, \varphi_1, \omega_1, v) = 0,$$

where the function S satisfies a set of Bellman equations:

$$(3) \quad S(\varphi_{0,1}, \omega_1, v | c_0) \\ = E_{\varphi_{2,2}} \left[\int_{b(\varphi_{1,2}, \omega_2, v)}^{b(\varphi_{1,2}, \omega_2, v)} (v - c_1) dG(c_1 | c_0, \varphi_{0,1,2}, \omega_{1,2}) + (\delta\lambda)^{\omega_2} \int_{b(\varphi_{1,2}, \omega_2, v)} S(\varphi_{1,2}, \omega_2, v | c_1) dG(c_1 | c_0, \varphi_{0,1,2}, \omega_{1,2}) \right].$$

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, $b(\varphi_0, \varphi_1, \omega_1, v)$ depends on the current competition insofar as the current competition is informative of the future competition: When evaluating the option value of the future, the bidder assumes that he or she will lose the current auction to a competitive bid that exactly matches his or her current bid. This “tie” is the only situation in which raising the current bid slightly changes the outcome of the game, and therefore S , given $c_0 = b(\varphi_0, \varphi_1, \omega_1, v)$, is the opportunity cost relevant at the margin. In other words, each bidder assumes that he or she is pivotal to the outcome of the first period. Finally, it is notable that bidders submit only positive bids on products of their desired type, a result that leads to the identification of personal preferences in the empirical test.

The bidding function is fully characterized by the expected surplus function S , which in turn is 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: $G(c_1 | c_0, \varphi_0, \varphi_1, \varphi_2, \omega_1, \omega_2)$. In equilibrium, the surplus function must reflect the actual expected surplus, given that everyone uses the optimal bidding strategy (Equation 2). Therefore, the equilibrium expected-surplus function must satisfy the Bellman equation (Equation 3). Such an S function exists because of the continuity of f , the compactness of its support, and the slope of S in any of its arguments being uniformly bounded (shown in the proof of P_1 ; see the Appendix). 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 (for a proof, see the Appendix):

P_2 : For all F , the equilibrium $b(\varphi_0, \varphi_1, \omega_1, v)$ has the following properties:

- $b(1, \varphi_1, \omega_1, v)$ increases in ω_1 ,
- $b(0, \varphi_1, \omega_1, v) = 0 < b(1, 1, \omega_1, v) < b(1, 0, \omega_1, v) < v$ for all $v > 0$, and
- $b(1, \varphi_1, \omega_1, v)$ decreases in ρ .

The first property shows that bids decrease when the future draws nearer in the sense that ω_1 decreases. In the empirical section, a generalization of this result is 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 enables 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 Equation 2, which is also tested in the empirical section, is the main result of this article because it shows that all forward-seeing bidders of all types bid less when they see their desired object in the next period than when they see an object they do not 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 holds 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 for this 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 thus could not be forward seeing. The resultant model would be analogous to Engelbrecht-Wiggans’s (1994) model, so this result shows that an important benchmark model is nested within the model proposed here. The empirical section is not able to identify this effect from a type-specific effect, because by definition, each type is observed with only one

long-term rate. A generalization of the basic model that informs empirical testing is discussed next.

Multiperiod Look-Ahead Model

The basic model contains all the intuition of a more general model, in which bidders see more than one period forward, a realistic extension that is 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 $\Phi = (\phi_1, \dots, \phi_A)$ and $\Omega = (\omega_1, \dots, \omega_A)$. Two empirically relevant summary statistics of (Φ, Ω) are considered: $H(\Omega)$ is the number of auctions in the next hour implied by Ω , and $w(\Phi, \Omega)$ is the waiting time until the first future auction that sells a product of the same type as the current product. The statistic $H(\Omega)$ is relevant because eBay auctions ending in the next hour are highlighted in red, and thus $H(\Omega)$ is easy to discern at a glance. The statistic $w(\Phi, \Omega)$ is relevant because it is invariant to the way consumers actually use the eBay Web site (i.e., whether they search for 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 P_1 and P_2 can be used to show that there is a Markov perfect symmetric-equilibrium pure strategy $b(\phi_0, \Phi, \Omega, v)$ with the following properties:

Corollary 1: When bidders see A periods ahead, the equilibrium bidding function $b(\phi_0, \Phi, \Omega, v)$ has the following properties:

- $b(1, \Phi^1, \Omega^1, v) \leq b(1, \Phi^0, \Omega^0, v)$ for all Φ^1, Ω^1 , such that (Φ^1, Ω^1) has one additional listing of any type in the next hour and otherwise is the same as (Φ^0, Ω^0) , so $H(\Omega^1) > H(\Omega^0)$;
- $b(0, \Phi^1, \Omega, v) = 0 < b(1, \Phi^1, \Omega, v) < b(1, \Phi^0, \Omega, v) < v$ for all $v > 0$ and for all $\Phi^1 \succ \Phi^0$, where $\Phi^1 \succ \Phi^0$ is defined by $\phi_a^1 \geq \phi_a^0$ for all a and by $\phi_b^1 > \phi_b^0$ for some b . In particular, the inequality holds for all $\Phi^1 \succ \Phi^0$, such that $\Phi^0 = 0$ and $\Phi^1 \geq 0$;
- $b(1, \Phi^1, \Omega^1, v) \leq b(1, \Phi^0, \Omega^0, v)$ for all Φ^1, Ω^1 , 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 $w(\Phi^0, \Omega^0)$; and
- Δb_a decreases in a , where

$$\Delta b_a = b(1, \Phi_a^0, \Omega, v) - b(1, \Phi_a^1, \Omega, v) \text{ and}$$

$$\phi_{a,j}^i = \begin{cases} 0 & \text{when } j < a \\ i & \text{when } j = a \\ \phi_{a,j}^{-i} & \text{when } j > a \end{cases}$$

The fourth claim examines how the magnitude of the bid decrement shown to be positive in the second claim changes as the first occurrence of the desired product type becomes more distant in the sequence, if timing is kept the same. The decrement becomes smaller as the future recedes into the distance, an immediate consequence of discounting, chance of attrition, and unit demand.

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 Corollary 1. Let $\bar{b}(x, v) = E[b(1, \Phi, \Omega, v) | x]$ stand for the expected bid of a bidder with valuation v conditional on x , with the expectation over the remaining state-components. Then, as long

as future timing Ω is independent of future types Φ and the continuation of the sequence is independent of its beginning, the claims in Corollary 1 average to testable predictions.

Model Predictions

If bidders act consistently with the proposed model, the following four relationships hold for all valuations $v > 0$ and for all desired types $k = 1, \dots, K$:

1. $\bar{b}(H, v)$ decreases in number of auctions in the next hour $H(\Omega)$;
2. $\bar{b}(\Phi^1, v) < \bar{b}(\Phi^0, v)$ for all $\Phi^1 \succ \Phi^0$, such that $\Phi^0 = 0$ and $\Phi^1 \geq 0$, so $\bar{b}(\Phi, v)$ decreases in the indicator function $I\Phi^1$ of Φ^1 ;
3. $\bar{b}(w, v)$ increases in waiting time until the same type w ; and
4. $\Delta \bar{b}(a, v)$ decreases in a , where

$$\Delta \bar{b}_a = \bar{b}(\Phi_a^0, v) - \bar{b}(\Phi_a^1, v) \text{ and } \phi_{a,j}^i = \begin{cases} 0 & \text{when } j < a \\ i & \text{when } j = a \end{cases}$$

Thus, $\bar{b}(\Phi, v)$ “decreases” in the indicator function $I\Phi_a^1$ of Φ_a^1 , as shown in the second claim, and the decrease is attenuated by a .

These 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, the first claim 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.

ROBUSTNESS OF THE MODEL PREDICTIONS

Several assumptions of the basic model can be relaxed without altering the key predictions of P_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. Because current competition does not enter the bidding strategy, the only difference a stochastic N_k would make to the results is that the entire right-hand side of Equation 3 would need 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 desire more than one type idiosyncratically 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 need to be specified for every bidder as $\phi_{j,i} = 1$ 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 φ_{i,v_i} . This structure of preferences may be relevant to the MP3-player category, in which each bidder may have use for only one player but be indifferent among several models. In contrast, each movie-category bidder may have use for multiple DVDs as long as they are all 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 need 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 $b_k(\varphi_{0,i}, \varphi_{1,i}, \omega_1, v_i)$ that all satisfy an analogue of P_2 , with $\varphi_{1,i}$ assuming the role of φ_1 . This model is investigated briefly in the empirical section, and as predicted, the data lend no support to this model in the movie data, but there is 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 because effective selling on eBay requires a much greater set of capabilities than buying on eBay; whereas buyers can treat eBay as any other online store, sellers need to have at least minimal Web-publishing skills along with the 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 data sets considered in the empirical section. More than 99% of the bidders are not observed selling anything in the same category within a month, and of the several thousand sellers observed, only approximately 2.5% submitted at least one bid per month within the entire category, all together affecting less than 3% of the auctions.

It remains to be shown that even in the presence of a small fraction of speculators, P_1 and P_2 remain valid descriptions of private-consumer behavior in eBay-like sequential auctions. The main reason that speculative considerations do not affect the qualitative predictions of the theory in online auction settings is that the predictions pertain to the impact of almost 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 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 or she 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 $S(\varphi_0, \varphi_1, \omega_1, v_i|c_0)$ 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 char-

acterization of the equilibrium relationship between forward-seen information and current bids shown in P_1 and P_2 would remain.

EMPIRICAL EVIDENCE: ANALYSIS OF eBAY DATA

This section uses the predictions noted in the end of the theory section to construct a statistical test that uses actual eBay bidding data, and it attempts to reject the proposed model in favor of one of two simpler models: (1) the model with forward seeing of only the type-independent summary statistic $H(\Omega)$ and (2) a model without any forward seeing.

Data

Two data sets were provided by eBay, each corresponding to a different product category: MP3 players and DVD movies. 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 the theory section. 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 third-party auction site that only facilitates the communication between buyers and sellers are bound to come without definitive identification of each item sold. To match each listing to a product type (movie title or player model), researchers must rely on the item description written by the seller, and some classification error is inevitable. In both data sets used here, a word-matching algorithm classified approximately 80% of the listings as likely selling a known product type, but the resulting classification is still only approximate. In the MP3-player category, 21% of all auction listings remained unclassified, either because their description was insufficient for identification ("new cool MP3 player for sale") or because they did not belong to the product category ("Napster T-shirt" or "128 MB SanDisk memory card"). 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 sell something other than a single unit of the type. Removing the top and bottom 3% of all bids on each type is sufficient to eliminate all listings that either sell for multiples of the median price on the type, probably indicating an undetected bundle, or are less than 10% 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.

Although both data sets are similar to each other in many respects, they differ slightly; the similarities are discussed first. Each data set contains all submitted bids in each recorded auction and information about each listing, including its timing and a text description of the item sold. The bidding data capture all the proxy bids made, including the winning bid, which remains undisclosed to eBay participants.¹ Individual bidders and sellers are tracked over time with unique identification numbers. All auctions that involve reserve prices or bid cancellations are eliminated

¹For a definition of proxy bids and a thorough description of the bidding rules, see www.ebay.com.

from the analysis because their modeling is beyond the scope of this article.

The MP3-player data set is constructed to capture all the auctions for the top 30 types (models) of MP3 players held during a four-month period in the beginning of 2001. The top 30 types account for approximately 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 data set. The minority (23%) 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 article and because their early termination makes them not useful as forward-seen future options. Excluding BIN auctions from the analysis makes the remaining auctions seem closer to one another 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. Because 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 downward, making the resulting estimates of all effects' magnitudes conservative. The remaining BIN auctions that either reverted to simple auctions or remained unsold were retained in the analysis. 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, and 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. Of the resulting 6967 auctions used in the analysis, approximately 50% were originally started as BIN auctions. There were 2663 unique sellers listing all the auctions, and there were 15,574 unique bidders (on average, 3.2 per auction) participating in the 4852 (70%) auctions that received bids. Almost half of the bidders participated in multiple auctions, raising the average number of unique bidders in an auction to 7.5, with a median of 7.

The movie data set is constructed to capture all simple auctions for 30 popular titles in October 2002, and popularity was judged using best-seller lists.² Buy-it-now auctions were not recorded in the movie data. The data set contains 4864 auctions listed by 1607 unique sellers, and there were 7445 unique bidders (on average, 2.2 per auction) participating in the 3384 (69%) auctions that received bids. About one-third of the bidders participated in multiple auctions, resulting in an average number of unique bidders per auction of 3.7, with a median of 3. Thus, the movie market 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 the MP3-player category and 87% in the movie category). Yet a substantial number of bidders participated in more than one auction (43% in the MP3-player category and 33% in the

movie category). An alternative explanation of multiauction bidding can be ruled out right away: It does not seem that the multiauction bidders simply submitted a very low bid initially to learn about the auction process or their true valuation, only later raising 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% submitted a higher second bid. The corresponding figure among the 4543 MP3-player bidders who bid on the same model at least twice in a row is 59%. Next, a more precise test of the proposed model based on the empirical relationship between bids and properties of near-future auctions is described.

Econometric Test

On eBay, a bid can be submitted only if it exceeds the highest bid at the moment, so the data set contains relatively more high bids and relatively fewer low bids than a random sample of willingness to bid modeled by the sealed-bid abstraction $\bar{b}(x, v)$. Although 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 $\bar{b}(x, v)$ conditional on x , $\bar{b}_{(1)}(x)$ and $\bar{b}_{(2)}(x)$, respectively, are observed in the data without any bias. Because all the model predictions pertaining to $\bar{b}(x, v)$ 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 $\bar{b}(x, v)$ as well, and thus $\bar{b}(x, v)$ can be replaced with $\bar{b}_{(1)}(x)$ or $\bar{b}_{(2)}(x)$ in all the model predictions of the theory section. Note that this approach does not require the knowledge of private valuations, because the predictions pertain to the impact of commonly known forward-seen types rather than the impact of privately known valuations. 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:

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

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 $k(i)$: $x_{i,k(i)} \in \{I\Phi_{i,k}^1, w_{i,k}, (I\Phi_{1,i,k}^1, I\Phi_{2,i,k}^1, \dots, I\Phi_{A,i,k}^1)\}$; 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. Because the three different specifications of $x_{i,k(i)}$ are at least partially correlated, three separate specifications of Equation 4 were estimated for both levels of m and both data sets. To improve the theoretical quality of the linear approximation implicit in Equation 4, the analysis of the MP3-player data set 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 greater than \$100. In each level-specific analysis, the players corresponding to the other price level are retained as part of the auction

²I thank Uri Simonsohn for selecting the popular titles and processing the movie data set.

stream, lumped together into the 16th “other” type. Next, the control variables z_i are discussed.

Because 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, all else being equal. 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 nonparametric 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.

In addition to the influence implied by the proposed model, including current competition 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 randomly chose only one auction of their product type to bid on and subsequently acted myopically by bidding their valuation rather than acting sequentially as proposed, more auctions in a short period would imply both a more desirable near future and fewer bidders per auction; thus, there would be a lower order statistic of the bids. Therefore, a mechanical negative correlation would arise 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, which has been shown by Resnick and colleagues (in press) and Wilcox(2000), among others, to have a positive impact on bids; a dummy variable for the description of the unit containing words such as “new” or “mint” as a coarse measure of within-type vertical differentiation of the products; and (not included by eBay in the movie data set) seller-controlled differentiation indicators of the listing itself (e.g., “photo included,” “bold-type listing,” “gallery listing”). The seller feedback score used in the MP3-player analysis was not included in the movie data set, so a dummy indicating whether the seller was an “eBay Top Seller” was used instead to capture the effect of reputation.

Because only the highest and the second-highest bid in each auction are used in the analysis and because eBay does not allow a bidder to outbid him- or 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. Thus, the analysis resolves the issue of “multiple bidding” (i.e., 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.

In both data sets, special care was taken to exclude bids that were obviously not made by private bidders modeled by the theory. In the MP3-player data, bids made by sellers (approximately 2%) and bids made by bidders who won more than one unit within the data period (approximately 12% of highest and 7% of second highest) were eliminated, resulting in 4075 highest bids and 4107 second-highest bids. Among the movie auctions, approximately 3% of both

highest and second-highest bids were eliminated because they were made by a seller or by a bidder who won multiple units of the same title, and approximately .4% of bids were eliminated because they were made by bidders who bid on too many types, resulting in 3118 highest and 2436 second-highest bids. For summary statistics of all the variables in the final data sets used in estimation and additional summaries of the data, see Table 1.

Results

In both data sets and according to both order statistics of bids, bidders seem to engage in at least one form of forward-seeing bidding. The two data sets are discussed in turn. Table 2 presents the parameter estimates for movies, Table 3 presents the parameter estimates for low-priced MP3 players, and Table 4 presents the parameter estimates for high-priced MP3 players. In the movie data set, all type-specific forward-seeing variables have γ_m coefficients consistent with the theory: Waiting times until the next auction of the same type increase bids, the same type offered in the next five 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, in general, is negative, as predicted. Notably, the effects are smaller for highest bids than for second-highest bids. This is consistent with the notion that bidders have idiosyncratic preferences for the particular versions of the good that are not captured by type alone. The measured effects on second-highest bids seem quite large, and they have the added relevance of essentially capturing the effects on price (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). With the average price in the category being approximately \$10, the same movie offered in the immediately following auction leads to an average price reduction of \$.72, and the same movie offered at least once within the next five auctions reduces the price by approximately \$.31. 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 is not significant. An explanation for this difference is that on lower-priced players, a 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 aforementioned 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 the theory-robustness section, namely, from the model involving bidders desiring multiple types. When the analysis is focused on multitype bidders and when each bidder’s desired types are identified as all the 30 types on which that

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)			
	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD		
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	.466	.499	.466	.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.	.745	.436	.765	.424	.699	.459	.699	.459		
Bold-listing dummy	n.a.	n.a.	n.a.	n.a.	.034	.182	.037	.189	.045	.208	.047	.211		
Gallery-listing dummy	n.a.	n.a.	n.a.	n.a.	.094	.292	.095	.294	.069	.254	.068	.252		
New dummy	.441	.497	.429	.495	.569	.495	.554	.497	.667	.471	.668	.471		
Current competition (number of 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(number of auctions next hour + 1)	2.403	.671	2.366	.669	1.469	.599	1.478	.590	1.580	.586	1.576	.586		
Log(time until next same-type auction + 1)	3.179	.885	3.270	.807	4.047	.880	4.043	.882	3.714	1.040	3.698	1.043		
Dummy (same type, next five auctions)	.178	.383	.172	.377	.299	.458	.291	.454	.497	.500	.505	.500		
Dummy (same type, one auction from now)	.065	.246	.056	.231	.129	.335	.123	.329	.223	.416	.236	.425		
Dummy (same type, two auctions from now)	.027	.161	.027	.163	.047	.212	.048	.214	.116	.320	.116	.320		
Dummy (same type, three auctions from now)	.028	.165	.026	.159	.048	.215	.050	.218	.073	.261	.068	.252		
Dummy (same type, four auctions from now)	.037	.188	.037	.189	.039	.194	.037	.188	.051	.219	.047	.212		
Dummy (same type, five auctions from now)	.023	.148	.025	.156	.035	.185	.033	.180	.035	.184	.038	.190		

Notes: 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, 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 30th product. In the raw MP3-player data set, the proportions of listings by selling format (BIN auction versus 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 because 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 data set does not include BIN auctions or auctions with a reserve price, making an analogous breakdown by selling format unavailable. "Same type, A auctions from now" means that the same type of good is available A auctions from current auction and that auctions {1, 2, ..., A-1} sell goods of different types. n.a. = not available.

Table 2
ESTIMATION RESULTS: MOVIES

Variable	Specification 1			Specification 2			Specification 3			
	Highest Bid		Second-Highest Bid	Highest Bid		Second-Highest Bid	Highest Bid		Second-Highest Bid	
	Estimate	(t-Statistic)	Estimate	(t-Statistic)	Estimate	(t-Statistic)	Estimate	(t-Statistic)	Estimate	(t-Statistic)
α (30 type-specific dummies)										
θ (top-seller dummy)	.645	9.298	8.509	9.325	.637	9.408	.628	9.180	.553	8.308
θ (new dummy)	.756	10.634	.826	11.869	.747	10.698	.748	10.712	.825	12.051
θ (current competition)	.128	8.807	.094	6.202	.127	8.894	.125	8.710	.094	6.259
β (log(auctions next hour + 1))	-.045	-8.33	.033	.631	-.087	-1.740	-.039	-1.682	-.039	-8.06
γ (log time until next same-type auction + 1)	.060	2.350	.108	4.071	-.170	-1.920	-.313	-3.510	-.722	-5.004
γ (same type, next five auctions)									-.385	-1.928
γ (same type, one auction from now)									.136	.674
γ (same type, two auctions from now)									-.051	-2.88
γ (same type, three auctions from now)									.070	.314
γ (same type, four auctions from now)										
γ (same type, five auctions from now)										
	N = 3017	R ² = .42	N = 2356	R ² = .53	N = 3113	R ² = .42	N = 2431	R ² = .53	N = 3113	R ² = .42

Notes: Parameters of primary interest appear in bold. Thirty movie-title fixed-effects are suppressed for parsimony. Specification 1 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 thus there is more truncation of forward-seeing information in the end of the data period. Specification 2 truncates only five observations in the end of the data period; thus, there is the reduced number of observations by five relative to the data. For variable definitions, see the notes to Table 1.

Table 4
ESTIMATION RESULTS: HIGH-PRICED MP3 PLAYERS

Variable	Specification 1		Specification 2		Specification 3	
	Highest Bid Estimate	Second-Highest Bid Estimate (t-Statistic)	Highest Bid Estimate	Second-Highest Bid Estimate (t-Statistic)	Highest Bid Estimate	Second-Highest Bid Estimate (t-Statistic)
α (15 type-specific dummies)						
θ (top-seller dummy)	.729	1.671	.816	1.890	.787	1.819
θ (photo-listing dummy)	.009	.007	-.162	-.126	-.102	-.080
θ (bold-listing dummy)	4.427	1.504	4.369	1.489	4.176	1.422
θ (gallery-listing dummy)	1.088	.445	.966	.403	1.146	.478
θ (new dummy)	7.112	5.070	7.259	5.206	7.117	5.098
θ (current competition)	.582	3.903	.536	3.621	.538	3.635
β (auctions next hour)	-5.057	-4.579	-6.755	-6.610	-6.577	-6.406
γ (log time until next same-type auction + 1)	2.617	5.511	-7.441	-4.886	-8.172	-6.073
γ (same type, next five auctions)						
γ (same type, one auction from now)					-8.830	-4.946
γ (same type, two auctions from now)					-7.624	-3.546
γ (same type, three auctions from now)					-8.364	-3.369
γ (same type, four auctions from now)					-3.214	-1.132
γ (same type, five auctions from now)					-4.023	-1.223
	N = 2317	R ² = .86	N = 2372	R ² = .85	N = 2372	R ² = .85
			N = 2451	R ² = .87	N = 2451	R ² = .87

Notes: Parameters of primary interest appear in bold. For variable definitions, see the notes to Table 1. For an explanation of different sample sizes, see the notes to Table 2.

bidder ever submitted a valid bid, a regression analogous to Equation 4 reveals that bidders submitting $\bar{b}_{(1)}$ and $\bar{b}_{(2)}$ 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 are not reported here). The converse is not true for high bidders and low-priced players, and as predicted, this model extension is not empirically supported in the movie data set. These results illustrate the richness of information contained in forward-seeing behavior, but their further development is beyond the scope of this article.

The β_m effects of the number of auctions in the next hour are small in both MP3-player subsamples. Doubling of the number of auctions ending in the next hour is associated with approximately 2% reduction in prices.³ Conversely, the type-specific effects γ_m on high-priced players are substantial: For items that cost approximately \$180, the same product being available within the next five auctions reduces prices by \$8 (4.4%) on average and by \$10 (5.6%) when the same product is available in the next 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 more than \$2. Note that all these estimates are conservative because they suffer from the errors-in-variables problem and thus are biased toward zero. All control variable parameters have the anticipated signs, but unlike in the movie data set, some variables are insignificant—notably, most of the listing variables under seller control (e.g., photo, bold, gallery). However, the insignificance does not imply that these instruments are of no value, because their usage is endogenous.

DISCUSSION

This article 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 rational, forward-looking bidders participate in a sequence of auctions for substitutes, they 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 article 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 will be sold next and when. Then, all rational bidders reduce their current bids when there are more near-future auctions in general, as well as when more of the auctions offer units of the good specifically desired by the bidders and when the units are offered sooner. These findings contribute to the auction theory literature and are relevant to bidders in sequential auctions on eBay and 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 account for detailed information about the near future, leading to price reductions between 3% and 7% whenever the same type of good is available in the next five auctions. Therefore, the empirical evidence suggests that the bidders behave consistently with the proposed model and that such behavior has a sizable 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 nonreversibility of the implied "search" in the auction marketplace. This article 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 into the underlying sequential search. This article provides the first step, but more research is needed to determine 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; that is, 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 examples of large Internet auction markets, in which stand-alone analysis of individual auctions would be inappropriate. Instead, this article demonstrates that individual auctions need to be interpreted within their context of other auctions that sell similar objects, and it provides a model that can be used to achieve such analysis. The model implies that observed bids are always negatively biased measures of true valuations because winning involves an additional opportunity cost that arises from not participating in future auctions for the same good.

Because the current results provide only 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 the structural estimation methodologies of Jofre-Bonet and Pesendorfer (2003) to study exact properties of auction market demand. The findings reported here may also affect seller strategies, raising the question of the scope of auction-driven markets (i.e., whether the interauction competition found here limits the potential of sequential auctions as trading institutions). Throughout this article, the seller was assumed to be exogenous, but allowing for strategic selling may both qualitatively and quantitatively change the bidder's strategy. A companion article (Zeithammer 2005) 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.

³ $\text{Log}(2)\beta_m/\text{mean}(\text{price})$ is $-.025$ in low-priced and $-.019$ in high-priced players.

APPENDIX: PROOFS OF PROPOSITIONS

Proof of P_1

First, consider $b(0, \varphi_1, \omega_1, v)$. Because the expected surplus function is positive and because 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 $b(1, \varphi_1, \omega_1, v)$. As long as

$$\frac{\partial S(1, \varphi_1, \omega_1, v|c_0)}{\partial c_0} > -\frac{1}{(\lambda\delta)\omega_1},$$

a solution to the optimal bidding problem is characterized by the first-order condition in Equation 2, because the problem is concave at the solution to Equation 2. Moreover, the solution to Equation 2 is unique for every φ_1, ω_1 , and v : Because for all c_0 , $\partial/\partial c_0(v - [\lambda\delta]^{\omega_1}S[1, \varphi_1, \omega_1, v|c_0]) < 1$, and because $v - (\delta\lambda)^{\omega_1}S(1, \varphi_1, \omega_1, v|c_0 = 0) > 0$ and $v - (\delta\lambda)^{\omega_1}S(1, \varphi_1, \omega_1, v|c_0 = 1) < 0$, it follows by continuity of S and the intermediate value theorem that there is exactly one $b(1, \varphi_1, \omega_1, v)$ that satisfies Equation 2. To show that

$$\left| \frac{\partial S(1, \varphi_1, \omega_1, v|c_0)}{\partial c_0} \right|$$

is sufficiently uniformly bounded to ensure concavity (and also the existence of S using the Schauder fixed point theorem), let the steady-state order statistics of the current competitive bids be $Y_{(1)} > Y_{(2)} > \dots > Y_{(N-1)}$, and note that

$$\left| \frac{\partial S(1, \varphi_1, \omega_1, v|c_0)}{\partial c_0} \right| < \left| \frac{\partial \tilde{S}(1, 1, \omega_1, 1|Y_{(1)} = c_0)}{\partial c_0} \right|,$$

where $\tilde{S}(1, \varphi_1, \omega_1, v|c_0)$ is the future surplus conditional on $Y_{(2)}$ surviving until the next auction for sure (as opposed to with probability $\lambda\omega_1$). 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 thus is bounded by the impact on the highest possible bidder $v = 1$. Let z be the expected maximum bid of new entrants. Then, $\tilde{S}(1, 1, \omega_1, 1|c_0)$ can be evaluated as $\tilde{S}(1, 1, \omega_1, 1|c_0) = E_{\varphi_2, \omega_2} [1 - E(c_1|c_0)] = E_{\varphi_2, \omega_2} \{1 - E[\max(Y_{(2)}, z)|Y_{(1)} = c_0]\}$, and it is clear that the impact of c_0 on $\tilde{S}(1, 1, \omega_1, 1|c_0)$ 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:

$$\left| \frac{\partial \tilde{S}(1, 1, \omega_1, 1|Y_{(1)} = c_0)}{\partial c_0} \right| < \left| \frac{\partial E(Y_{(2)}|Y_{(1)} = c_0)}{\partial c_0} \right| < 1.$$

This concludes the proof that there is a well-defined pure strategy characterized by Equation 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. Q.E.D.

Proof of P_2

The claim that $b(1, \varphi_1, \omega_1, v)$ increases in ω_1 follows immediately from the optimal bidding condition (Equation

2). The claim that $b(1, \varphi_1, \omega_1, v) < v$ follows from S being positive. The claim that $0 < b(1, \varphi_1, \omega_1, v)$ follows $S(\varphi_0, \varphi_1, \omega_1, v|c_0) < v$, which is always true because the bidder cannot derive more utility than that from getting a unit of the desired type for free, for sure, and immediately, thus receiving a surplus of v . The central claim of the first part of P_2 , $b(1, 1, \omega_1, v) < b(1, 0, \omega_1, v)$, hinges on bidding on a desired type always giving the bidder at least as much surplus, all else being equal, as bidding on an undesired type; this is shown in two steps: (1) for every c_0 , $S(1, 0, \omega_1, v|c_0) < S(1, 1, \omega_1, v|c_0)$, and (2) to show Step 1 implies $b(1, 1, \omega_1, v) < b(1, 0, \omega_1, v)$.

To show Step 1, it is instructive to write down the four Bellman equations that characterize the steady-state expected-surplus functions in all possible combinations of current and future desirability states, keeping timing $\omega_1 = \omega_2 = 1$ constant and suppressing it from all equations:

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

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

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

and

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

The G distribution arises from all the surviving losers of the current auctions and from all the new entrants, and because 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 exited the bidder pool, whereas no trade means that the highest competing bidder exited the bidder pool only with probability $(1 - \lambda)$. Therefore, $G(c_1|c_0, 0, \varphi_1, \varphi_2) < G(c_1|c_0, 1, \varphi_1, \varphi_2)$, and because S decreases in c , as shown in the proof to P_1 , $S(0, \varphi_1, v|c_0) < S(1, \varphi_1, v|c_0)$. Conversely, $G(c_1|c_0, 1, 1, \varphi_2) = G(c_1|c_0, 1, 0, \varphi_2)$ because the bidding function is increasing in v , and thus φ_1 has no differential impact on the kind of bidders who are likely to survive from the past period 0. Therefore, the key difference between surpluses can be written as

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

and this difference is positive because $S(0, \varphi_1, v|c_0) < S(1, \varphi_1, v|c_0)$ and because $(v - c_1 - S[1, \varphi_2, v|c_1]) > 0$ for all $c_1 < b(1, \varphi_2, v)$, which follows from the single-crossing property discussed in the proof of P_1 . Step 2 also follows from the single-crossing property of the first-order condition. Because $(v - c_1 - S[1, \varphi_2, v|c_1]) = 0$ has a unique solution and $(v - 0 - S[1, \varphi_2, v|0]) > 0$, $b_0 = b(1, 0, v)$ implies that $(v - b_0 - S[1, 1, v|b_0]) < 0$, and thus the point b_1 , such that $(v - b_1 - S(1, 1, v|b_1)) = 0$, must lie to the left of b_0 .

That $b(1, \varphi_1, \omega_1, v)$ decreases in ρ follows from differentiation of the Bellman Equation 3 after writing the expectation $E_{\varphi_2}[\cdot]$ as $\rho(\varphi_2 = 1) + (1 - \rho)(\varphi_2 = 0)$ and noting, by an argument analogous to preceding one, that $S(\varphi_0, 0, \omega_1, v|c_0) < S(\varphi_0, 1, \omega_1, v|c_0)$; that is, whatever today's type, it is always better to face a desired type tomorrow than not. Because higher ρ increases the chance of $\varphi_2 = 1$, the result follows. Q.E.D.

REFERENCES

- Ariely, Dan and Itamar Simonson (2003), "Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions," *Journal of Consumer Psychology*, 13 (1), 113-23.
- Ashenfelter, Orley (1989), "How Auctions Work for Wine and Art," *Journal of Economic Perspectives*, 3 (3), 23-36.
- Bajari, Patrick and Ali Hortacsu (2003), "Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34 (2), 85-101.
- Dholakia, Utpal (Paul) and Kerry Soltysinski (2001), "Coveted or Overlooked? The Psychology of Bidding for Comparable Listings in Digital Auctions," *Marketing Letters*, 12 (3), 225-37.
- Engelbrecht-Wiggans, Richard (1994), "Sequential Auctions of Stochastically Equivalent Objects," *Economic Letters*, 44, 87-90.
- Gale, Ian L. and Donald B. Hausch (1994), "Bottom-Fishing and Declining Prices in Sequential Auctions," *Games and Economic Behavior*, 7, 318-31.
- Jofre-Bonet, Mireia and Martin Pesendorfer (2003), "Estimation of a Dynamic Auction Game," *Econometrica*, 71 (5), 1443-89.
- Milgrom, P. and R. Weber (2000), "A Theory of Auctions and Competitive Bidding II," in *The Economic Theory of Auctions*, Paul Klemperer, ed. Cheltenham, UK: Edward Elgar Publishing.
- Resnick, Paul, Richard Zeckhauser, John Swanson, and Kate Lockwood (in press), "The Value of Reputation on eBay: A Controlled Experiment," *Experimental Economics*, forthcoming.
- Roth, Al and Axel 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 (4), 1093-1103.
- Vickrey, William (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16 (1), 8-37.
- Wilcox, Ronald T. (2000), "Experts and Amateurs: The Role of Experience in Internet Auctions," *Marketing Letters*, 11 (4), 363-74.
- Zeithammer, Robert (2005), "An Equilibrium Model of a Dynamic Auction Marketplace," working paper, Graduate School of Business, University of Chicago.

Copyright of *Journal of Marketing Research* (JMR) is the property of American Marketing Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.