MARKETING SCIENCE
Vol. 29, No. 6, November–December 2010, pp. 964–987
ISSN 0732-2399 | EISSN 1526-548X | 10 | 2906 | 0964

DOI 10.1287/mksc.1100.0561


# The Sealed-Bid Abstraction in Online Auctions

Robert Zeithammer
Anderson School of Management, University of California, Los Angeles, Los Angeles, California 90095,
rzeitham@ucla.edu

Christopher Adams
Federal Trade Commission, Washington, DC 20580, cadams@ftc.gov

This paper presents five empirical tests of the popular modeling abstraction that assumes bids from online auctions with proxy bidding can be analyzed "as if" they were bids from a second-price sealed-bid auction. The tests rely on observations of the magnitudes and timings of the top two proxy bids, with the different tests stemming from different regularity assumptions about the underlying distribution of valuation signals. We apply the tests to data from three eBay markets—MP3 players, DVDs, and used cars—and we reject the sealed-bid abstraction in all three data sets. A closer examination of these rejections suggests that they are driven by less experienced bidders. This consistent rejection casts doubt on several existing theories of online auction behavior and suggests some demand estimates based on the abstraction can be biased. To assess the direction and magnitude of this bias, we propose and estimate a new model in which some bidders conform to the abstraction while other bidders bid in a reactive fashion. Because reactive bidding can be at least partially detected from the data, we are able to estimate the underlying distribution of demand and compare it to what the sealed-bid abstraction implies. We find that our proposed model fits the data better, and our demand estimates reveal a large potential downward bias were we to assume the second-price sealed-bid model instead.

*Key words*: auctions; online auctions; demand estimation; sealed-bid model; eBay; order statistics
*History*: Received: August 3, 2008; accepted: January 10, 2010; processed by Teck Ho. Published online in *Articles in Advance* August 11, 2010.

## 1. Introduction: The Sealed-Bid Abstraction

The rules of most online auctions, such as eBay, resemble an ascending English auction because the bidders can revise their bids upward whenever they are outbid. However, all major auction sites also provide proxy-bidding software designed to save the bidder effort by automatically bidding on his behalf up to a secret maximum. Because of this proxy system, the economically relevant endgame of each auction resembles a second-price sealed-bid auction, and so it seems reasonable to model bids in online auctions as bids arising from a second-price sealed-bid auction. We propose and apply a series of empirical tests of this *sealed-bid abstraction*, and our tests reject it consistently in three diverse data sets. Testing this particular abstraction is important for two reasons. First, our current theoretical understanding of online auctions is dominated by models that involve the abstraction as their prediction or assumption. A rejection thus focuses future theorists on more-complicated theories that do not abstract from within-auction dynamics. Second, the state-of-the-art method for nonparametric identification of demand in online auctions (Song 2004) depends crucially on the abstraction. A rejection focuses our search for a more realistic empirical model. We propose a new empirical model in which only some bidders conform to the abstraction while others bid in a reactive fashion. Assumptions about the nature of this reactive bidding allow us to estimate the underlying demand. Our demand estimates reveal a large potential downward bias were we to assume the second-price sealed-bid model instead. Before summarizing our results, we discuss in more detail the two reasons for testing the sealed-bid abstraction.

At least five existing theories of online bidding involve the sealed-bid abstraction. Most of the theories that explicitly model bid timing predict bidding at the last moment of the auction. Placing a proxy bid (equal to an equilibrium bid in the corresponding second-price sealed-bid auction) at the last moment can be a symmetric equilibrium strategy for at least three reasons. First, this strategy allows tacit collusion against the seller (Roth and Ockenfels 2002). Second, it protects private information in a common value setting (Bajari and Hortaçsu 2003). Third, it avoids bidding wars with an irrational fringe of incremental bidders (Ariely et al. 2005). Instead of predicting last-minute bidding, several existing theories simply abstract away from the timing of bids under the assumption of independent private values (IPV) because of the strategic equivalence between the English auction and the second-price sealed-bid auction (Zeithammer 2006, Yao and Mela 2008). Also under IPV, Song (2004) assumes bidders do not have the luxury of waiting until the last moment, but instead each has a different last opportunity to bid. Placing a proxy bid equal to one's valuation[1] at one's last opportunity to bid is an equilibrium strategy of Song's online auction game. See the survey by Bajari and Hortaçsu (2004) for an in-depth exposition of most of the aforementioned theories. Our rejection of the sealed-bid abstraction leaves several models that do not predict or assume it. For example, Peters and Severinov (2006) obtain nonsealed bidding in an IPV model with sequential arrivals of bidders to many simultaneous auctions. Another model that does not predict the abstraction is the model of a bidder who does not know his valuation but can tell whether his valuation exceeds any given price once he sees the price. Hossain (2008) shows that such a bidder would bid low and often in an eBay auction to learn his valuation.

Note that bidding that conforms to the sealed-bid abstraction is not the same concept as bidding just once per auction. Although some models that conform to the abstraction involve bidding just once per auction (e.g., Bajari and Hortaçsu 2003), others do not (e.g., Song 2004 does not rule out multiple bidding prior to the last opportunity to bid). Therefore, testing the abstraction is not as simple as checking whether all bidders submit more than one bid in an auction. Even if many bidders engage in multiple bidding, their final bids in the auctions may still look "as if" they arose in a sealed-bid auction. For example, the winner in Peters and Severinov (2006) might bid less than his valuation, i.e., less than what he would bid in a second-price sealed-bid auction under IPV.

The sealed-bid abstraction is not only consistent with much extant theory but is also critical for nonparametric identification of demand from online auction data. In standard auction models, demand is nonparametrically identified by a single order statistic (e.g., the second order statistic, often equal to the closing price) and the number of bidders (Athey and Haile 2002, 2005). When the number of bidders is unobservable, as in most online auctions, nonparametric identification is still possible in the IPV context if any *two* order statistics of the bid distribution are observed (Song 2004). The identification strategy relies heavily on the sealed-bid abstraction, so empirical researchers need a test to determine whether it holds in their particular data sets.

[1] Throughout this paper, "consumer's valuation" is the consumer's maximum willingness to pay at the moment, i.e., the dollar utility of the good *net* of all other opportunities to buy a substitute good elsewhere. Specifically, valuation is not the consumer's intrinsic value of the product (see also Bajari and Hortaçsu 2003, Chan et al. 2007).

Our rejection of the sealed-bid abstraction leaves several empirical modeling approaches that do not invoke it. One way to proceed is to model the bids directly using a reduced-form stochastic model. For example, Park and Bradlow (2005) model bid increments and bidder arrivals based on an evolving latent willingness to bid. Bradlow and Park (2007) explicitly model how the latent valuations of the bidders evolve over time and model the observed bids through a record-breaking model. Another way to proceed is to somehow infer the latent number of bidders and then use the standard identification while assuming the second-largest bid equals the second-highest valuation. For example, Chan et al. (2007) extend the bounds approach of Haile and Tamer (2003) by inferring latent bidders from observed bidders in concurrent auctions. Alternatively, Adams (2007) relies on an exogenous proxy for the number of bidders. Finally, Yao and Mela (2008) simply take the observed number of bidders to be the number of bidders who actually participated in the auction. We take a third approach and propose an empirical model that preserves the conditional order statistic approach of Song (2004) and thus does not rely on inference about the number of bidders. Besides not relying on any inference procedure about the number of latent bidders, our model does not assume the second-highest valuation is equal to the second-highest bid: when some bidders snipe and some bid in a reactive fashion, raising their bid only gradually whenever they are outbid, a sniper might win the auction without giving the second-highest reactive bidder a chance to counter.

We now briefly summarize the proposed tests and our new model, starting with the former. The tests use data on magnitudes and timing of the top two bids in each auction, where "bid" stands for the secret maximum proxy bid. A model of bidding is defined to conform to the sealed-bid abstraction if the final bid of each bidder in an auction depends only on that bidder's private signal about the valuation of the object sold. Under the null hypothesis that the abstraction holds, the following five things should be true, the first four are as follows: (T1) the top bid should be equally likely to arrive before or after the second-highest bid, (T2) the top two bids should not be exactly one bid increment apart too often, (T3) the difference between the top two bids should not be a function of which was placed first, and (T4) the difference between the top two bids should not be a function of time remaining in the auction.

Our last test (T5) is a nontrivial contribution to the empirical auction literature. Instead of relying on timing or increment data, it considers only the joint distribution of the top two bids and asks what should be true about the conditional distribution of the top bid given the second-highest bid. The sealed-bid abstraction implies that the conditional distribution of the top bid given the second-highest bid should have the same right tail for any particular value of the second-highest bid. When bids depend on auction-level observables as well as private signals, the T5 test can be applied to the residuals of the appropriate truncated regression of bids on observables. In this paper, we specify the different assumptions under which each of these five tests works, and we propose a way to operationalize the novel test T5. Our operationalization of T5 generalizes the nonparametric Wilcoxon test to deal with draws from tails of distributions.

Theoretically, there is nothing special about the top two bids in an auction—most of the above tests would be valid with any other pair of order statistics of the bidding distribution. Practically speaking, however, the top two order statistics are the only reliably observable ones because of entry truncation in online auctions: on eBay, one can submit a bid only if it exceeds the highest bid at the moment; thus eBay data contain relatively more high bidders and relatively fewer low bidders than the underlying population. Although many latent bidders may thus be truncated, the highest and the second-highest bidder in each auction are always recorded. Conversely, the lower order statistics of the observed bids do not necessarily correspond to the order statistics of the latent bidders. Another reason to focus on the top two bids is the obvious alternative "incremental" model of bidding in which each bidder manually raises his proxy bid by the minimum increment up to his maximum (i.e., up to his valuation in an IPV context). The top two bids are special in that incremental bidding by all participants predicts sharply different outcomes of T1–T5 from sealed bidding. Specifically, when bidders bid in an incremental fashion, the top bid always arrives after the second-highest bid, the top two bids are always exactly one bid increment apart, and the conditional distribution of the top bid given the second bid degenerates to a mass point one increment above the second-highest bid.

We apply the tests to three different data sets from eBay (MP3 players, movies on DVD, and used cars). Our data are provided directly by eBay, so we observe the proxy bid of the winner not available from the eBay website. Therefore we report on empirical regularities of uniquely complete data sets.[2] Taken together, our application of all five tests rejects the sealed-bid abstraction as a general property of bidding in eBay auctions. Three tests reject consistently, with surprising empirical regularity of the test statistics across the three diverse data sets. First (T1), the top bid is placed after second-highest bid in about two thirds of the auctions. Second (T2), about 15% of the auctions end with the two bids exactly one increment apart (compared to only 4% that end in an exact tie). Finally (T3), the one-increment-apart outcome is about three times more likely when the top bid comes after the second-highest bid compared with the reverse order. The easiest way to explain these findings is that a significant proportion (more than 15%) of the auctions within each data set is better described by incremental bidding. Although sealed bidding is thus not a general property of all auctions, it may still describe a subset of auctions large enough for demand estimation based on pairs of order statistics. We explore this possibility with two "plausibly sealed" subsets: the "OverInc" auctions in which the top bid exceeds the second-highest bid by more than an increment, and the "HighFirst" auctions in which the top bid is placed before the second-highest bid. Interestingly, tests T1–T4 still consistently reject the sealed-bid abstraction in these plausibly sealed subsets. We also applied the novel T5 test to the OverInc data after conditioning on auction-level observables, and it rejected the abstraction in two of our three data sets (MP3 players and cars).

To assess the direction and magnitude of bias one would encounter if one were to rely on the abstraction in demand estimation following Song (2004), we develop an alternative empirical model. When we examine the root causes of the T1–T5 rejections, we find the bidding style of the auction winner is an important correlate. Specifically, auctions won by multibid bidders (bidders observed submitting multiple bids per auction) tend to conform to the abstraction less than auctions won by single-bid bidders (bidders observed bidding only once per auction). Although multiple bidding is not clear evidence against the sealed-bid abstraction (as discussed above), our data indicate that multibid bidders bid systematically differently in that they tend to submit smaller bids than single-bid bidders. Therefore we propose that bidders have different personal bidding styles, and only the "sealed style" conforms to the abstraction, whereas the "reactive-style" bidders initially bid only a fraction of their valuation and subsequently raise their bid gradually toward their valuation whenever they are outbid. Reactive bidders can be at least partially detected by bidding multiple times in a single auction, and we can rely on bidder characteristics, such as experience, to estimate the probability that any given bidder is reactive. The link between experience and multiple bidding replicates previous findings by Wilcox (2000) and Borle et al. (2006). Consistent with the results presented in

[2] These data are not completely unique anymore because eBay now sells its data through third-party providers such as Advanced Ecommerce Research Systems. Obtaining data from eBay is not the only way to observe the proxy bid of the winner. For example, Bapna et al. (2008) observe it by running their own sniping agent.

List (2003) and Simonsohn and Ariely (2008), we find that experienced traders are more likely to behave "rationally" in the sense of conforming to the sealed-bid abstraction.

Given the bidding-style probability for each bidder and a standard IPV assumption, we can estimate the underlying distribution of valuations (a.k.a. "demand") by interpreting the observed bid distribution as a weighted combination of bids coming from sealed bidders and reactive bidders. The estimation uses the conditional order statistic approach, and so it does not need to infer the number of latent bidders. We estimate the model on the DVD data (in which the rejection of the sealed-bid abstraction was the weakest), and we compare the estimates to an alternative model that assumes all bidders are using the sealed-bidding style. The proposed model fits the data better than the all-sealed model and corrects a downward bias caused by the reactive bidders not bidding their true valuations. We find that this potential bias is large: specifically, the estimation results suggest that valuations of DVDs have a population mean and variance that are both more than double the all-sealed estimates. In addition, the two models imply substantially different public reserve prices. For example, a seller with a realistic marginal cost of \$3 per DVD should use a starting price of \$4.20 under the all-sealed model estimates versus \$6.00 under the proposed model. Therefore our rejection of the sealed-bid abstraction has substantial managerial implications.

The rest of this paper is organized as follows. The next section introduces the tests and the assumptions on which they are based. Section 3 then describes our data and applies the tests. Section 4 considers robustness of our results to relaxation of the assumptions. Section 5 discusses the new model with reactive bidders. Section 6 concludes by summarizing our results and outlining how the empirical regularities we document constrain theories of online auction behavior and econometric methods for estimating demand from eBay data.

## 2. Tests of the Sealed-Bid Abstraction in Online Auctions

### 2.1. Definitions

Our tests are geared toward understanding whether the economically important final bids in eBay auctions behave analogously to bids in a second-price sealed-bid auction. However, the tests' applicability extends to other auctions and other sealed-bid scenarios. To facilitate the widest possible scope of application, we define our primitives in maximum generality. Our concept of an *online auction* encompasses any auction that receives bids over time, with each bid associated with a unique time stamp. For the purposes of this paper, every current Internet auction is thus an online auction, but so is every other auction that receives mail-in or call-in bids. The feedback to bidders during the auction can range from none (as in a government auction with mail-in bids) to a full record of successful bids to date (as in an eBay auction).

A model of an online auction involves a *sealed-bid abstraction* whenever the bidders bid as if they were in a standard sealed-bid auction. Mathematically speaking, each bidder in a sealed-bid auction receives a private scalar signal $x$ and bids according to a strictly increasing function $\beta(x)$: $x \rightarrow$bid that depends only on $x$. For example, Bajari and Hortaçsu (2003) assume $x$ is a private signal about the common value of the auctioned good (a collectible coin). In such a common value environment, a symmetric equilibrium exists in which all bidders bid at the last moment as if they were in a second-price sealed-bid auction. The equilibrium bid function is increasing and mitigates winner's curse by $\beta(x) < x$. Similarly, bidders in Song's (2004) model receive independent signals about their private valuations, and each bidder has an exogenous last opportunity to bid. In such a private value model, each bidder has a dominant strategy to bid $x$ at his last opportunity to bid. The observable bidding behavior may not look sealed in the sense that bidders might be submitting multiple bids and following arbitrary dynamic bidding strategies before their last opportunity. However, the final bid each bidder submits should be the bid he would submit in a second-price sealed-bid auction. Note that in both examples, the second-price nature of the sealed-bid auction arises from eBay's proxy-bidding agent. Our tests do not depend on the equilibrium assumption—$\beta(x)$ could be any ad hoc behavioral regularity. For example, bidders may not actually be strategic but may follow eBay's instructions that direct them to submit their maximum willingness to pay as their proxy bids on arrival to the auction.

### 2.2. Model Properties

Having defined the two key primitives of our theory, we now turn to the properties of auction models from which our tests arise. (Please refer to Table A.5 in the appendix for notation used throughout this paper.) We will present a series of tests, and each test will rely on different properties of the auction model. Therefore, even if one or two of the following properties do not hold, some of the proposed tests will still work. We will be using the following properties.

- A1 (timing independent of signals): time $t_i$ is independent from signal $x_i$ for every bidder $i$.
- A2 (continuity): signals $x_{ij}$ of bidder $i$ in auction $j$ are drawn from a continuous distribution.
- A3 (conditional iid): conditional on auction-specific observables $Z_j$, signals $x_{ij}$ are independent

and identically distributed (iid) across auctions $j$ and bidders $i$. The function $\log\beta(x)$ is additively separable in auction-specific observables $Z_j$ and residual private valuation shocks.

- A4 (increment): the auction is actually an ascending auction with a minimum bid increment *inc*, which the econometrician observes.

The first property (A1) says no link exists between the magnitude of private signals and the timing of the bids associated with those private signals. As discussed in §1, this property does not hold in the models of Peters and Severinov (2006), Bradlow and Park (2007), and Hossain (2008). The second property (A2) says the signals are drawn from a continuous distribution that may vary from auction to auction and may involve arbitrary correlations across bidders, auctions, or both. A popular property that allows pooling of data across auctions in empirical research is that the signals are iid across auctions and bidders. This iid property is the essence of property A3, with A3 also conditioning on auction-level observables. The additional additive separability assumption is not needed if one can obtain a large number of observations with the same $Z_j$; then the implied test T5 can be validly carried out on the bids themselves. To pool across observations with different $Z_j$, we will approximate conditioning on $Z_j$ with a linear regression of log bids on observables and focus on the residuals. One setting that satisfies the additive separability assumption is an IPV setting with $\beta(x) = x$ and valuations that are multiplicatively separable into $Z_j$ and private shocks. Finally, property A4 is an assumption about the rules of the auction that generates the data.

### 2.3. Tests

Suppose we have the data on the top two bids in an auction $b_1 \geq b_2$, their timing $t_i = \text{time}(b_i)$, and the increment *inc*. In the eBay setting, $(b_1, b_2)$ are the top two proxy bids submitted in the auction and *inc* is eBay's minimum increment.[3] Given data on $(t_1, t_2, b_1, b_2, inc)$, the simplest test is based on timing alone: as long as timing of bids is independent of private signals (A1), the sealed-bid abstraction predicts neither of the bids should be more likely to appear first:

$$(\text{A1}) \;\Rightarrow\; (\text{T1}){:}\; \Pr(t_1 > t_2) = 1/2. \tag{1}$$

In the alternative incremental bidding model, $\Pr(t_1 > t_2) = 1$; thus the highest bid usually coming in after the second-highest suggests a departure toward incremental bidding. To operationalize this test, we compute the empirical probabilities with their associated standard errors.

In online ascending auctions with a minimum increment (A4), another simple test uses only data on bids and the continuity property (A2):

$$(\text{A2, A4}) \;\Rightarrow\; (\text{T2}){:}\; \Pr(\Delta b = inc) = 0, \tag{2}$$

where $\Delta b = b_1 - b_2$. In contrast, the incremental model predicts that $\Pr(\Delta b = inc) = 1$, so T2 failing should be a strong indication of incremental behavior. To operationalize T2, we again compute the empirical probability with its standard error. A2 is an unrealistic assumption about the auction environment when some bidders value goods in whole dollar amounts. Section 4 will generalize A2 to allow mass points in the distribution and discuss how such a generalization would impact testing. With mass points, $\Pr(\Delta b = inc) > 0$, and test T2 becomes weaker.

Combining the bid data with the timing data allows for more detailed testing based on the independence assumption (A1). By definition of independence between timing and bidding, A1 implies $\Delta b$ should be completely invariant to the timing of the bids. Two fruitful tests arise, again motivated by the incremental behavior alternative. First, the sealed-bid abstraction predicts that the difference between the top two bids should not depend on which bid came first:

$$(\text{A1}) \;\Rightarrow\; (\text{T3}){:}\; \text{cdf}(\Delta b \mid t_1 > t_2) = \text{cdf}(\Delta b \mid t_1 < t_2), \tag{3}$$

where cdf stands for cumulative distribution function. In contrast, if some auctions behave more like incremental auctions and others more like sealed-bid auctions, $(\Delta b \mid t_1 > t_2)$ will be smaller and related to the minimum bid increment, whereas $(\Delta b \mid t_1 < t_2)$ will be "more continuous" with a tail. To operationalize this test idea, we use the nonparametric Wilcoxon–Mann–Whitney (hereafter WMW) rank test and compare the subsample $(\Delta b \mid t_1 > t_2)$ to the subsample $(\Delta b \mid t_1 < t_2)$.[4] Analogous to conditioning on the relative timing of the top two bids, conditioning on time remaining in the auction should also leave $\Delta b$ unchanged. Define bidding as *late* versus *early* by doing a median split on the time left $\bar{t} - \min(t_1, t_2)$, where $\bar{t}$ is the ending time of the auction. The sealed-bid abstraction implies

$$(\text{A1}) \;\Rightarrow\; (\text{T4}){:}\; \text{cdf}(\Delta b \mid late) = \text{cdf}(\Delta b \mid early). \tag{4}$$

To operationalize this test, we again use the WMW test.

Bringing in the increment information allows us to zoom in on the value of $\Delta b = inc$ to conduct special cases of T3 and T4 with a sharp alternative hypothesis

[3] eBay's increment is a function of $b_2$ and varies from 5 cents to \$100 as $b_2$ increases from \$1 to \$5,000.

[4] In parallel, we also computed the standard $t$-test of the $E(\log\Delta b \mid t_1 > t_2) = E(\log\Delta b \mid t_1 < t_2)$ hypothesis, implicitly assuming lognormality of $\Delta b$. This parametric test almost perfectly agrees with the WMW test.

under the incremental model. Specifically, we can test whether

$$
\begin{aligned}
(\text{A1, A4}) \;\Rightarrow\; (\text{T3}')\text{:}\;& \Pr(\Delta b = inc \mid t_1 > t_2) \\
&= \Pr(\Delta b = inc \mid t_1 < t_2), \\
(\text{A1, A4}) \;\Rightarrow\; (\text{T4}')\text{:}\;& \Pr[\Delta b = inc \mid late] \\
&= \Pr[\Delta b = inc \mid early].
\end{aligned} \tag{5}
$$

Note that none of the above tests requires particular assumptions on the correlation of signals, either within or across auctions. When such correlations are only due to observables, another test is possible that does not rely on timing data. Consider a single auction $j$, and suppose (as in A3) the signals are only correlated within the auction because of some common shock $Z_j$ that is observable to both the bidders and the econometrician (on eBay, auctions differ in ending time, characteristics of the good sold, and characteristics of the seller). As long as $Z_j$ is the only source of dependence, the sealed-bid residuals $\varepsilon_{ij}$ after conditioning the bids on $Z_j$ are iid according to some distribution $F_j$. The second part of A3 then allows pooling of data across auctions by assuming the distribution of residuals is the same in all auctions in the sample:

$$
\begin{aligned}
&(\text{A3 within auction}) \\
&\quad\Rightarrow\; \varepsilon_{ij} \equiv \log\beta(x_{ij}) - E[\log\beta(x_{ij}) \mid Z_j] \overset{\text{iid}}{\sim} F_j, \\
&(\text{A3 across auctions}) \;\Rightarrow\; \forall k, j\text{: } F_k = F_j \equiv F.
\end{aligned} \tag{6}
$$

The residual $\varepsilon_{ij}$ is the component of the bid $\beta(x_{ij})$ arising from private information of bidder $i$ in auction $j$. $F$ is the distribution that demand analysts seek to recover from bidding data: when the auction sells private value goods, $F$ is the distribution of private valuations in the population. In a common value setting, $F$ is the distribution of bids that all correct for the winner's curse by bidding below the private signal of value, so $F$ can be used together with $\beta^{-1}(\cdot)$ to recover the population distribution of private signals. Given the above iid assumption, Song (2004) shows that knowing any two order statistics of the bidding distribution and their order is sufficient for identification of $F$ even when the analyst does not know the number of bidders in each auction. When the two observed order statistics of bids are $(b_1, b_2)$—as in our situation—the identification of $F$ is particularly simple: $\Pr(\varepsilon_1 < z \mid \varepsilon_2 = w) = (F(z) - F(w))/(1 - F(w))$, where $\varepsilon_1$ is the residual of $b_1$ and $\varepsilon_2$ is the residual of $b_2$. In particular, the conditional distribution of $\varepsilon_1$ given $\varepsilon_2$ is the right tail of $F$ truncated at $\varepsilon_2$ with a probability density function $\text{pdf}(\varepsilon_1 \mid \varepsilon_2 = w) = f(\varepsilon_1)/[1 - F(w)]$. Independence across bidders within the auction is a key assumption for this result to go through; independent and identical distribution across auctions is only useful to pool data across auctions.[5] This regularity implies an entire range of tests based on the fact that the shape of the conditional distribution of $(\varepsilon_1 \mid \varepsilon_2)$ does not depend on the particular value of $\varepsilon_2$:

**Figure 1** **Illustration of the Tail Comparison Test**

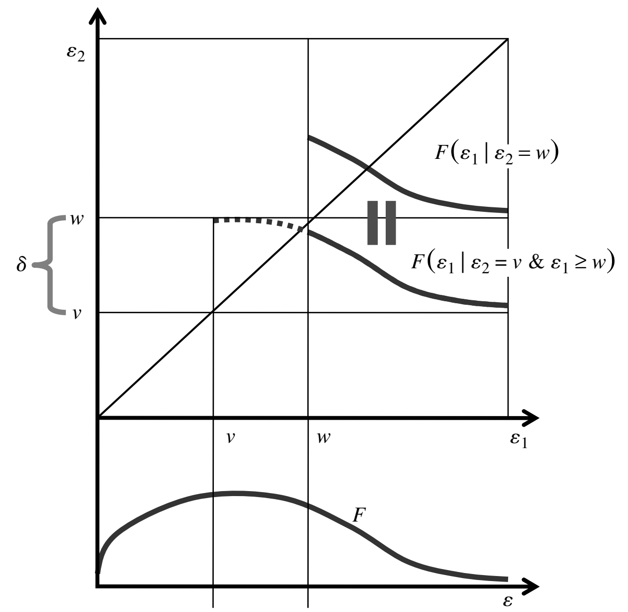


*Notes.* The lower part of the graph shows the population distribution $F$ of the private bid residual $\varepsilon$. The upper part of the graph shows the distribution of the $\varepsilon_1$ given $\varepsilon_2$ for two values of $\varepsilon_2$: $w = v + \delta > v$. Note how the nondashed tails of the two conditional distributions are the same as the tail of $F$ truncated below at $w$ and so are equal to each other.

$$
\begin{aligned}
&(\text{A3}) \;\Rightarrow\; \forall w \text{ and } \delta > 0\text{:} \\
&\quad \text{cdf}(\varepsilon_1 \mid \varepsilon_2 = w) = \text{cdf}(\varepsilon_1 \mid \varepsilon_2 = w - \delta \text{ and } \varepsilon_1 \geq w).
\end{aligned} \tag{7}
$$

Figure 1 illustrates the resulting "tail comparison test" for a particular value of $w$ and $\delta$, with the equal sign corresponding to the equal sign in the above equation. (Table A.1 in the appendix illustrates what the conditional distributions of $(b_1 \mid b_2)$ look like in practice for a specific value of $w$ (in raw bid data for the most popular movie title, i.e., without conditioning on auction-level observables $Z_j$).) Theoretically, every possible $w$ could produce a separate test like the one in Figure 1, but combining all the tests would be cumbersome. To combine all the data into a single test statistic, we develop a new nonparametric statistical test based on all pairs of auctions.

Proposition 1. *Let $(\varepsilon_{1,j}, \varepsilon_{2,j})$ be the top two order statistics from a sample of $N_j \geq 2$ iid draws from some*

[5] Independence implies the joint distribution $\text{pdf}(\varepsilon_1, \varepsilon_2) = N(N-1)f(\varepsilon_1)f(\varepsilon_2)F(\varepsilon_2)^{N-2}$, i.e., the probability of the event that one draw is $\varepsilon_1$, another draw is $\varepsilon_2$, and all other draws are below $\varepsilon_2$. The resulting multiplicative form of the density separates the two order statistics. Thus, when $\varepsilon_2$ is a constant $w$, $\text{pdf}(\varepsilon_1, \varepsilon_2 \mid \varepsilon_2 = w) \propto f(\varepsilon_1)$.

*distribution F. Call a pair of auctions j and k such that* $\min(\varepsilon_{1,j}, \varepsilon_{1,k}) > \max(\varepsilon_{2,j}, \varepsilon_{2,k})$ *a "feasible pair," and denote j to be the auction with the higher second bid*: $\varepsilon_{2,j} = \max(\varepsilon_{2,j}, \varepsilon_{2,k})$. *Then, in all the feasible pairs, the higher* $\varepsilon_2$ *is equally likely to occur in the same auction as the higher* $\varepsilon_1$:

$$\text{T5} \equiv \Pr[\varepsilon_{1,j} > \varepsilon_{1,k} \mid \varepsilon_{2,j} > \varepsilon_{2,k} \text{ and } \min(\varepsilon_{1,j}, \varepsilon_{1,k}) > \max(\varepsilon_{2,j}, \varepsilon_{2,k})] = 1/2.$$

*When the data set is selected such that* $\varepsilon_{1,j} > \varepsilon_{2,j} + inc_j$, *then the feasibility criterion needs to be stricter; namely,* $\min(\varepsilon_{1,j}, \varepsilon_{1,k}) > \max(\varepsilon_{2,j} + inc_j, \varepsilon_{2,k} + inc_k)$.

Please see the appendix for a formal proof.

The intuition for the feasibility requirement is that $\varepsilon_{2,j} = \max(\varepsilon_{2,j}, \varepsilon_{2,k})$ further truncates the already truncated $\Pr(\varepsilon_{1,k} < z \mid \varepsilon_{2,k})$, hence making the resulting conditional distribution of $\varepsilon_1$ above $\max(\varepsilon_{2,j}, \varepsilon_{2,k})$ completely independent of the realization of $\varepsilon_2$—that is, independent of $j$ and $k$. Since the distributions of the top bids in a feasible pair are identical, the prediction of Proposition 1 is that in all the feasible pairs, both auctions are equally likely to have a higher top bid. Specifically, the auction with a relatively higher price is *not* more likely to also have a higher top bid. By relying on the 50–50 chance of one draw from a distribution exceeding another draw, this test is a generalization of the WMW test. The intuition for the stricter feasibility criterion in the subset of auctions with $b_{1,j} > b_{2,j} + inc_j$ is that the subset's selection builds in additional truncation. To test the prediction outlined in Proposition 1, we simply compute an empirical estimate of the T5 probability across all feasible pairs and compare it to 1/2. To gauge statistical significance, we bootstrap the test statistic using 10,000 resamples with replacement from the original data set (Efron and Tibshirani 1994).

We computed the T5 test statistic under the null (iid) hypothesis while assuming different parametric data-generating distributions $F$, and we found the distribution of the test statistic under the null hypothesis seems to be invariant to the distribution that generated the bids. Therefore we conjecture that T5 is a nonparametric test in that its distribution under the null hypothesis does not depend on the distribution of bids in the population. Please see the appendix for details. We believe this test statistic is new and potentially useful outside auction research. Whenever an analyst has data that naturally group into pairs of order statistics, the test statistic can be used to test whether the data is iid in the population without making any specific distributional assumptions. For example, suppose we have data on the heights of the two tallest children in many families that all have at least two children (but can differ in the number of children). We can use T5 to check whether child heights are iid in the entire population (and we will probably reject this hypothesis because height is partly genetic).

**Table 1 Which Tests Rely on Which Model Properties**

| Tests | Model properties (assumptions of tests) | | | |
|---|---|---|---|---|
| | A1 (Timing indep.) | A2 (Signals continuous) | A3 (Signals cond. iid) | A4 (*inc* known) |
| T1: $\Pr(t_1 > t_2) = 1/2$ | ✓ | | | |
| T2: $\Pr(\Delta b = inc) = 0$ | | ✓ | | ✓ |
| T3: $\text{cdf}(\Delta b \mid t_1 > t_2) = \text{cdf}(\Delta b \mid t_1 < t_2)$ | ✓ | | | |
| T3′: $\Pr(\Delta b = inc \mid t_1 > t_2) = \Pr(\Delta b = inc \mid t_1 < t_2)$ | ✓ | | | ✓ |
| T4: $\text{cdf}[\Delta b \mid late] = \text{cdf}[\Delta b \mid early]$ | ✓ | | | |
| T4′: $\Pr(\Delta b = inc \mid late) = \Pr(\Delta b = inc \mid early)$ | ✓ | | | ✓ |
| T5: $\Pr[\varepsilon_{1,j} > \varepsilon_{1,k} \mid \varepsilon_{2,j} > \varepsilon_{2,k}$ and $(j,k)$ feasible$] = 1/2$ | | | ✓ | |

The null hypothesis prediction of T5 = 1/2 is completely nonparametric and does not depend on assumptions A1 or A2. However, large samples of exactly identical auctions are rare, so nonparametric conditioning on $Z_j$ is not practical. Operationalizing T5 as a semiparametric test with the conditioning on $Z_j$ achieved although a parametric model is therefore necessary. We assume bids are distributed lognormally, and the effect of $Z_j$ on bids is additively separable from the private bidder-specific residuals: $\log\beta(x_{i,j}) = \theta Z_j + \varepsilon_{i,j}$ and $\varepsilon_{i,j} \sim N(0, \sigma^2)$.

Under this specification, a Normal regression of $\log b_{1,j}$ on $Z_j$ truncated at $\log b_{2,j}$ recovers $\theta$, and the residuals $(\hat{\varepsilon}_{1,j}, \hat{\varepsilon}_{2,j})$ can be used to compute the test statistic T5. One practical problem with a truncated Normal regression is the lack of convergence in likelihood maximization when the model is severely enough misspecified. When we cannot achieve convergence, we resort to simply regressing $\log b_{2,j}$ on $Z_j$ and applying the test to the resulting ordinary least-squares (OLS) residuals. To further reduce the potential confounding impact of unobserved heterogeneity, we perform the test separately for every product—that is, separately for different movie titles and MP3 player models.

This discussion concludes the theoretical development of econometric tests. Table 1 illustrates how the different tests depend on different subsets of properties (assumptions).

## 3. Application of the Tests to eBay Data

### 3.1. Data

We have three data sets at our disposal that include observations of $(t_1, t_2, b_1, b_2, inc)$. The data sets capture bidding on popular MP3 players in 2001, popular

DVDs in 2002, and cars in 2003. Each data set was selected to not end by the buy-it-now option and to have at least two bidders who bid above the minimum bid set by the seller (and above the reserve price, wherever applicable). Ties are resolved as follows: $b_1$ is defined as the winning bid, so a tie at the top ($b_1 = b_2$) implies $t_1 < t_2$. A tie for second-highest bid is resolved in favor of the bid that would have won in the absence of $b_1$; that is, $t_2$ is the time of the earliest second-highest bid.

The MP3 player data set captures 6,316 auctions. For each auction, we observe the seller reputation, auction characteristics like "photo included," the brand and model of the product sold, and whether the seller advertised the player as "new." One player was particularly popular on eBay during the time of the data—the Diamond Rio 500 with 1,328 listings. The second-most popular player (KB Gear JamP3) only had 603 listings. The movie data set captures 3,512 auctions and excludes auctions with a secret reserve. Each auction includes an indicator of whether the seller was a "top seller," the title of the movie, and whether the seller advertised the DVD as "new." The most popular movie in our sample is *Black Hawk Down* with 375 listings.

The car data capture thousands of used cars sold in 2003 as well as a longer series on one of the most popular cars—the C5 Corvette. The Corvette data capture more than 700 auctions between 2001 and 2003, with exact timing available only for the 2003 data (543 auctions). The C5 is the fifth version of the Corvette—a relatively new and homogenous version (produced without major changes from 1997 to 2004). To get the observable attributes of each car, we selected the auctions with a valid vehicle identification number (VIN). A valid VIN gives information about the car, including make, model, year, engine type, and model style, corroborating the car information the seller provides. We eliminated several observations in which the VIN information did not agree with the information the seller provided. Finally, we eliminated about two dozen auctions that sold for less than \$10,000 or more than \$40,000 because these were outliers on the log scale (the median price was \$26,200, standard deviation was \$8,000).

These three data sets span a wide range of products sold on eBay, and they all contain the information on $(t_1, t_2, b_1, b_2, inc)$ because eBay provided them directly (in data obtained from the eBay website, $b_1$ would not be available). This information allows us to apply the tests developed in the previous section. We conduct these tests in two stages. First, we use tests T1–T4, which do not rely on the strong A3 assumption, and we can convincingly classify about 30% of the auctions as "not sealed" because of $\Delta b$ being "too small"—that is, one increment or less. Using test T5 on all the data would thus have a foregone negative conclusion, but the remaining 70% of plausibly sealed OverInc auctions could possibly still be used to identify demand. We therefore apply test T5 only to the OverInc subset of auctions in the second stage.

### 3.2. Results of Tests T1–T4 Based on Continuity of $F$ or on the Independence of Timing and Magnitude

Table 2 documents the results of tests T1–T4 by data set and within each data set by two popular products. Three tests reject the sealed-bid abstraction consistently across all data sets, with surprising empirical regularity of the test statistics. First (T1), $b_1$ comes after $b_2$ in about two thirds of the auctions. This result suggests some but not prevalent incremental bidding. Second, about 15% of the auctions end with the two bids exactly one increment apart (T2). This result again suggests incremental behavior. In §4, we will show the empirical $\Pr(\Delta b = inc)$ is "too high" to come from sealed bidding even when we relax the A2 assumption to allow mass points at whole dollars and arbitrary multiples of the increment. Further suggesting that some auctions are better captured by the incremental model, the exactly-one-increment-apart outcome is about three times more likely when $b_1$ comes after $b_2$ compared with the reverse order (T3′). Therefore the incremental-like timing tends to co-occur with the incremental-like $\Delta b$. Note that these uniformly rejecting tests are based on different assumptions, and at least one of them remains valid whenever either A1 or A2 holds.

The remaining tests (T4 and T4′) reject the sealed-bid abstraction in some but not all data sets. With only a few exceptions, test T4′ suggests auctions with late bidding are more likely to involve $\Delta b = inc$. With the exception of the car data, test T4 finds larger $\Delta b$ in auctions with early rather than late bidding. The late bid auctions include all auctions in which both of the top two bidders sniped. Therefore the T4 result suggests that even sniping bidders react to the early bids, and even the snipe bids do not conform to the sealed-bid abstraction. We did not conduct test T3 on the full data because of the way we resolved ties: because exact ties in the top two bids involve $t_1 < t_2$ by definition, the subsample of $\{\Delta b \mid t_1 < t_2\}$ involves an ad hoc substantial mass point at zero.

Table 2 reports the significant rejections at 5% of probability of Type 1 error within every test. Because there are 27 different slices of the data and between three and five tests per slice, it is important to rule out the possibility that we reject the null hypothesis because of a multiple-testing problem. Specifically, we need to demonstrate that we do not simply reject the null hypothesis by running many tests

**Table 2 Results of Tests T1–T4**

| | MP3 players | | | DVDs | | | Cars | | |
|---|---|---|---|---|---|---|---|---|---|
| Selected product | All | Rio 500 | KB Gear JamP3 | All | *Black Hawk Down* | *A Beautiful Mind* | All Fords | Ford F150 | C5 Corvette |
| All auctions | | | | | | | | | |
| Number of observations | 6,316 | 1,328 | 603 | 3,512 | 375 | 274 | 5,621 | 510 | 543 |
| T1: $\Pr(t_1 > t_2)$ | **64.3** | **61.1** | **64.2** | **72.9** | **72.5** | **75.6** | **63.2** | **68.9** | **67.8** |
| T2: $\Pr(\Delta b = inc)$ | **14.2** | **12.6** | **18.9** | **14.4** | **16.8** | **13.5** | **11.4** | **14.1** | **17.7** |
| T3′: $\Pr(\Delta b = inc \mid t_1 < t_2)$ baseline | 4.6 | 2.2 | 8.4 | 4.9 | 7.5 | 1.6 | 4.6 | 4.6 | 10.9 |
| $\Delta \Pr(\Delta b = inc \mid$ ">" versus "<") | **14.4** | **15.5** | **16.1** | **12.9** | **12.7** | **15.7** | **10.5** | **12.5** | **10.5** |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | **<1.0** | **<1.0** | **3.6** | **1.9** | 71.4 | 56.1 | 75.1 | 30.9 | 36.5 |
| T4′: $\Pr(\Delta b = inc \mid$ *early*) baseline | 12.7 | 10.4 | 17.0 | 12.9 | 16.0 | 9.5 | 11.9 | 14.1 | 17.3 |
| $\Delta \Pr(\Delta b = inc \mid$ *late* versus *early*) | **2.9** | **4.5** | **3.7** | **3.0** | 1.6 | **8.1** | −0.9 | 0.0 | 0.4 |
| OverInc ($\Delta b > inc$) | | | | | | | | | |
| Number of observations | 4,484 | 922 | 419 | 2,365 | 229 | 198 | 4,162 | 352 | 370 |
| T1: $\Pr(t_1 > t_2)$ | **64.6** | **62.0** | **64.2** | **73.4** | **73.8** | **73.7** | **60.4** | **67.6** | **66.7** |
| T3: ($\Delta b \mid t_1 > t_2$) versus ($\Delta b \mid t_1 < t_2$) | **<1.0** | **<1.0** | 16.7 | **<1.0** | 5.4 | 58.2 | **<1.0** | 4.8 | 3.8 |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | **<1.0** | **<1.0** | **<1.0** | 35.7 | 88.3 | 78.7 | 20.0 | 68.7 | 19.6 |
| HighFirst ($t_1 < t_2$ and $\Delta b > 0$) | | | | | | | | | |
| Number of observations | 1,620 | 335 | 154 | 780 | 81 | 57 | 1,559 | 125 | 128 |
| T2: $\Pr(\Delta b = inc)$ | **5.6** | **2.7** | **10.4** | **5.5** | **8.6** | 1.8 | **4.9** | **4.8** | **11.7** |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | **<1.0** | **1.3** | 58.1 | 17.6 | 98.5 | 85.0 | 7.4 | 89.2 | 50.6 |
| T4′: $\Pr(\Delta b = inc \mid$ *early*) baseline | 4.2 | 0.6 | 6.5 | 5.6 | 10.0 | 0.0 | 4.7 | 3.2 | 15.6 |
| $\Delta \Pr(\Delta b = inc \mid$ *late* versus *early*) | **2.7** | **4.2** | **7.8** | −0.3 | −2.5 | 3.6 | 0.2 | 3.2 | −7.8 |

*Notes.* Bold values are rejections of the sealed-bid abstraction significant at the 5% level. All numbers are probabilities or differences in probabilities scaled between 0 and 100. T1, T2, T3′, and T4′ entries show the test statistics (themselves probabilities), whereas T3 and T4 entries show the *p*-values of the WMW rank-sum test, scaled as probabilities from 0 to 100.

and finding at least one rejection. The most conservative way to rule out a multiple testing alternative explanation of our results is to assume all 99 tests in Table 2 are independent replicates of each other, and to use the Bonferroni correction for all the *p*-values (Games 1977). The Bonferroni correction suggests that we can reject the null hypothesis at a "familywise" 5% confidence level when we find at least one rejection at a *p*-value of $0.05/99 \approx 0.0005$. (Table A.4 in the appendix shows all the *p*-values associated with the test statistics in Table 2, and almost half (46%) of the *p*-values are below the conservative Bonferroni threshold.) Therefore, our rejection of the sealed-bid abstraction is not a spurious effect of multiple testing.

Tests T1–T4 cast serious doubt on the applicability of the sealed-bid abstraction to any of the three data sets because of mass points of $\Delta b$ at zero and *inc*. The incremental model obviously captures the auctions corresponding to those mass points better. A question still stands whether the remaining approximately 70% of the auctions with $b_1 > b_2 + inc$ fare better on the tests—that is, whether the OverInc auctions seem to conform to the sealed-bid abstraction more. We therefore apply the tests that do not rely on *inc* (T1, T3, and T4) to the OverInc subset of the data. The results are also shown in Table 2 and suggest OverInc auctions do not conform to the abstraction either. First, not only does $\Pr(t_1 > t_2)$ remain significantly above 1/2, restricting attention to OverInc subset does not seem to reduce $\Pr(t_1 > t_2)$ at all. Second, test T3 rejects the sealed-bid abstraction in five of the nine slices of the data. The emergent empirical regularity based on a Hodges-Lehmann estimate is that $\Delta b$ is smaller when the highest bid follows after the second-highest bid compared to the other order (not reported). Finally, test T4 rejects the sealed-bid abstraction in the MP3 player data. In summary, based on the joint distribution of timing and bids, the sealed-bid abstraction seems doubtful even within the OverInc subset.

Another subset of the data exists in which incremental bidding should not operate, namely, $t_1 < t_2$: if the higher proxy bid is placed before the lower proxy bid, the high bidder probably did not react to $b_2$ in submitting his bid $b_1$. We call this subset "HighFirst," and we exclude $b_1 = b_2$ observations from it as well even though these involve $t_1 < t_2$ by our definition. We can apply tests T2, T4, and T4′ to test the HighFirst data. Test T2 still rejects in that the mass point $\Delta b = inc$ is significantly bigger than zero. However, $\Pr(\Delta b = inc)$ is smaller than in full data (around 5% compared with 15% in MP3 players and DVDs, about 10% versus 20% in cars). Tests T4 and T4′ reject only in the MP3 player data. Therefore the sealed-bid abstraction is still rejected in the HighFirst subset of auctions but less strongly and less consistently than in the OverInc subset or in the full data.

Is the nonsealed behavior we observe associated with any observable differences across auctions, such as different days of the week, different sellers, or different buyers? To investigate the correlates of the T1's rejections, we compared the "more sealed" HighFirst auctions to the "less sealed" HighSecond auctions (auction with $t_1 > t_2$) in the DVD data set. We find no significant difference in the product sold (new versus used), the seller's feedback, or the timing of the auction ending (weekend versus weekday). In contrast, the HighFirst auctions are associated with more experienced winners (642 days registered on eBay versus 596 in HighSecond) and with winners who did not bid multiple times in the same auction (82% in HighFirst versus 66% in HighSecond). All these differences are significant at the 5% level. Analogously for T2, we compared the $\Delta b = inc$ auctions to the $\Delta b \neq inc$ subset and again found that the only significant differences are those between bidder experience and multibidding. Specifically, $\Delta b = inc$ auctions are associated with less experienced winners (507 versus 625 days on eBay, 74 versus 127 feedback points) and with more multibidding winners (46% versus 27%). Of course, multibidding and less experience are themselves mutually positively correlated (Wilcox 2000, Borle et al. 2006). Therefore, bidder heterogeneity in experience is likely driving our rejections, with more experienced bidders conforming more to the sealed-bid abstraction.

In the next section, we apply the timing-free test T5 to the OverInc data. Although tests T1–T4 reject the sealed-bid abstraction based on the joint distribution of $(t_1, t_2, b_1, b_2)$ even in OverInc and HighFirst subsets, a carefully selected subset of the $(b_1, b_2)$ observations could still be used in demand identification. From the test application perspective, seeing how one would test the sealed-bid assumption without the timing data is also interesting. Of the two subsets described so far, we focus on the OverInc subset because it contains a lot more data than the HighFirst one and because it avoids the $\Delta b = inc$ mass point.

### 3.3. Results of the Tail Comparison Test T5 Based on the Conditional iid Assumption

The first battery of tests rejected the sealed-bid model as a general property of all auctions. Specifically, about 30% of the auctions are suspect because the top two bids were exactly one increment apart or less. However, 70% of the auctions with top two bids more than an increment apart remain, and the sealed-bid model might fit these auctions well enough to permit nonparametric demand estimation based on conditional order statistics. Test T5 is ideally suited to investigate this possibility.

As explained in §2.3, the first step of test T5 is conditioning on auction-level observables. To give the assumption A3 the best chance of holding, we first focus on one popular product at a time, and then we run the control lognormal regressions $\log b_{1,j} = \theta Z_j + \varepsilon_{1,j}$ truncated at $\log(b_{2,j} + inc_j)$ on remaining auction-level observables $Z$. In the MP3 player and DVD data, we focus on the two most popular products (Diamond Rio 500 and KB Gear JamP3 in players, *Black Hawk Down* and *A Beautiful Mind* in movies). To further reduce unobserved heterogeneity, we also restrict the data on the most popular MP3 player to the listings stating "new" in the description. In the car data, we focus on the C5 Corvette. Note that in the case of C5 Corvettes, we have more observations for test T5 (732) than reported in Table 2 (543) because our data set contains bid data but no timing data from two additional years. Different observables are available in the three different data sets. (Table A.2 in the appendix shows both the variables and results of the control regressions.) As explained in §2.3, we can run each regression in two ways. Table 3 reports the results of T5 using both approaches as well as the raw bid data without any conditioning. For C5 Corvettes, the theoretically preferable truncated regression of $\log b_{1,j}$ on $Z_j$ truncated at $\log b_{2,j}$ did not converge, so we only report the results based on the OLS price regression and the raw data.

The T5 test statistic rejects the sealed-bid abstraction in the MP3 player and car data but not in the DVD movie data. Therefore we cannot reject the sealed-bid abstraction in the movie data based on bids alone, and the corresponding truncated control regressions can be interpreted as estimates of demand.

Interestingly, the rejections all involve $T5 < 1/2$, namely, a negative correlation between the prices and the high bids among feasible pairs. We will argue next

**Table 3** Results of the Tail Comparison Test (T5) on the OverInc Subset of Auctions

| Product | Number of auctions | Truncated regression residuals (%) | OLS regression residuals (%) | Raw data (%) |
|---|---|---|---|---|
| MP3 player: Diamond Rio 500 (new) | 785 | **46.2** | 48.6 | **44.9** |
| MP3 player: Diamond Rio 500 | 922 | **46.3** | 47.3 | **43.5** |
| MP3 player: KB Gear JamP3 | 418 | 52.8 | 52.0 | 46.3 |
| DVD movie: *Black Hawk Down* | 229 | 50.8 | 51.7 | 49.8 |
| DVD movie: *A Beautiful Mind* | 198 | 49.0 | 51.0 | 48.4 |
| Car: Chevrolet Corvette C5 | 732 | n/a | **35.1** | **33.9** |

*Notes.* Each test statistic is the probability that in a feasible pair of auctions, the auction with the higher price also has a higher top bid. The OverInc data include all auctions with $b_1 > b_2 + inc$. Test statistics significantly different from 1/2 at 5% level are shown in bold.

(in §4.1) that T5 < 1/2 rules out unobserved (to the econometrician) heterogeneity of the auctioned goods as an alternative explanation of these rejections. The reason behind T5 < 1/2 is the clearest in the Corvette C5 data: the joint distribution of $(b_1, b_2)$ is so concentrated near the "diagonal" (near $b_1 = b_2$) that finding feasible pairs of auctions is difficult, and the tails of the conditional distributions of $(b_1 \mid b_2)$ are so steep that within the feasible pairs, the higher price makes a lower top bid more likely. The top bid is thus a function of the second-highest bid, and there is no additional information contained in observing the top bid if the second-highest bid is known. Intuitively, this coupling of order statistics cripples the conditional order statistic approach that relies on the information about demand contained in $b_1$ beyond what is contained in $b_2$.

## 4. Robustness Checks

### 4.1. Unobserved Heterogeneity of Auctioned Goods Does Not Explain T5 Results

The biggest weakness of the tail comparison test (T5) is the assumption (A3) of no auction-level demand shocks unobservable to the econometrician. Suppose instead that the regression equation is $\log \beta(x_{i,j}) = \theta Z_j + \xi_j + \varepsilon_{i,j}$, and the econometrician does not observe $\xi_j$. Then the test will reject spuriously whenever the variance of $\xi_j$ is similar to the variance of $\varepsilon_{i,j}$. Consider the case of apples and oranges, when the econometrician only observes fruit being auctioned. Suppose the auction is actually sealed, there are no observable differences $Z_j$, and apples are usually privately valued at slightly more than oranges. Conditioning on a lower $b_2$ confounds conditioning on the fruit being an orange with conditioning on the second-highest bidder on either fruit having a relatively low draw from the valuation distribution. The question thus arises: Are the rejections in Table 3 because of unobserved shocks $\xi_j$ or because of the sealed-bid abstraction not holding?

Table 3 already contains the two pieces of evidence against this alternative explanation. First, unobserved heterogeneity of products would imply a positive correlation between the top two bids and hence a T5 statistic above 1/2. This claim is difficult to prove analytically, but our simulations indicate it is true in a wide variety of settings. In contrast, all the significant rejections we find are rejections with T5 below 1/2. Second, T5 is not consistently closer to 1/2 when we consider more homogeneous subsets of the products (i.e., new versus all Rio 500 players). We conclude that unobservable heterogeneity alone cannot explain our rejections of the sealed-bid abstraction; that is, the rejection does not seem to be due to unobservable characteristics in the different auctions but rather because the sealed-bid abstraction does not hold.

### 4.2. Lumpy Distribution of Bids Does Not Explain T2 Results

Another assumption that may be too strong for the reality of eBay is the continuity assumption A2. If many bidders bid in whole dollar or whole increment amounts, the resulting mass points in the distribution of $\Delta b$ may arise even in sealed bidding. The movie data are ideally suited to check this possibility because the support of the bidding distribution only involves 20 whole dollar amounts ($5–$25), and the increment is 50 cents on the entire support, so there are only 19 additional whole increment mass points ($5.50, $6.50, ..., $24.50). The empirical distribution of the highest bid indeed does not look continuous: 37% of the bids are whole dollars and an additional 16% are a whole dollar plus 50 cents (considering the most popular movie, *Black Hawk Down*). To approximate the probability of $K$-increment differences $\Pr(\Delta b = K \cdot inc)$ one would expect from sealed bidding with such a lumpy distribution of bids, we perform a simple simulation of the top two bidders in a counterfactual sealed-bid auction. We simulate a second-price sealed-bid auction with two bidders drawn iid from the empirical cdf of $b_1$ (and then ordered). By simulating a million repetitions of such an auction, we find the expected mass points are much lower than those observed in the data. We then increase the number of simulated bidders all the way to 10, keeping the distribution the same, and again examine the distribution of $\Delta b$. The results of this simulation exercise are shown in Table 4.

Table 4 shows that mass points in the distribution of signals may be able to explain the empirical probabilities of exact ties as well as the probability of the top two bids being exactly two increments apart. However, the expected $\Pr(\Delta b = inc)$ remains below 5% even for high numbers of simulated bidders—far below the empirical value of 16.8%. Therefore, a distribution of signals that is as lumpy as the distribution of the highest bid cannot explain the empirical probability of the top two bids being exactly one increment apart.

### 4.3. The Top Two Bids Are Not Special in Rejecting Based on Magnitude and Timing

The sealed-bid abstraction may be violated by the top bid while holding for lower-order bids. One reason why the top bid might be "special" comes from the alternative purely incremental theory: if every bidder bids in a purely incremental fashion, all lower-order bid statistics $b_2, b_3, b_4, \ldots$, correspond to the valuations of the bidders who placed them, whereas the top bid is exactly one increment above the second bid

**Table 4 Can a Lumpy Bid Distribution Explain the Top Two Bids Exactly One Increment Apart?**

| | | Simulated sealed-bid auction, by number of simulated bidders | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| $\Delta b/inc$ | Data | 2 bidders | 3 bidders | 4 bidders | 5 bidders | 6 bidders | 7 bidders | 8 bidders | 9 bidders | 10 bidders |
| Exactly 0 | **3.2** | **2.5** | **3.7** | **4.6** | **5.0** | **5.3** | **5.7** | **5.7** | **6.1** | **6.2** |
| 0–1 | 18.9 | 10.6 | 14.7 | 16.4 | 17.3 | 17.7 | 17.6 | 17.4 | 17.2 | 16.8 |
| Exactly 1 | **16.8** | **3.4** | **4.6** | **4.9** | **5.2** | **4.8** | **4.4** | **4.0** | **3.6** | **3.3** |
| 1–2 | 16.8 | 9.7 | 11.9 | 12.2 | 12.0 | 12.0 | 11.9 | 12.0 | 12.0 | 12.2 |
| Exactly 2 | **7.5** | **4.0** | **4.9** | **5.0** | **5.0** | **5.0** | **4.9** | **5.0** | **4.8** | **5.0** |
| More than 2 | 36.8 | 69.9 | 60.2 | 56.8 | 55.4 | 55.2 | 55.6 | 55.9 | 56.3 | 56.5 |

*Notes*. The table shows both the actual and the predicted distributions of the difference between the top two bids $\Delta b$ measured in multiples of the increment. The distribution of bids used in simulations is the empirical distribution of $b_1$ for *Black Hawk Down*, the increment is 50 cents throughout. The distribution of $\Delta b/inc$ is similar for other movies as well as in the car and MP3 player data.

($b_1 = b_2 + inc$). To check whether the top bid is special, we computed all the T1–T4 tests on the second and third bids $b_2$ and $b_3$. These data are only available for the movie and MP3 player data sets.

Under the incremental model, the second bid should always come after the third bid ($t_2 > t_3$), the two bids should never be exactly an increment apart, and the distribution of $\Delta b = b_2 - b_3$ should be invariant to timing and order. In contrast, the sealed-bid predictions remain analogous to those presented for the ($b_1, b_2$) bid pair. Table 5 shows that the ($b_2, b_3$) bid pair rejects the sealed-bid abstraction with consistency and strength similar to the ($b_1$, $b_2$) pair: the higher bid follows the lower bid only about 80% of the time, they are exactly one increment apart about 10% of the time, and the distribution of $\Delta b$ depends on time left in the auction. These patterns persist in the corresponding OverInc and HighFirst subsets (not reported). We conclude that the T1–T4's rejection of the sealed-bid abstraction is not due to some special nature of the top bid.

**Table 5 Results of Tests T1–T4 on the Second- and the Third-Highest Bids ($b_2$ and $b_3$)**

| | MP3 players | | | DVDs | | |
|---|---|---|---|---|---|---|
| Selected product | All | Diamond Rio 500 | KB Gear JamP3 | All | *Black Hawk Down* | *A Beautiful Mind* |
| All auctions: Number of observations | 3,465 | 1,325 | 462 | 2,922 | 297 | 231 |
| T1: $\Pr(t_2 > t_3)$ | **80.6** | **80.0** | **79.6** | **84.4** | **81.5** | **80.1** |
| T2: $\Pr(\Delta b = inc)$ | **7.2** | **7.0** | **8.0** | **10.7** | **14.5** | **10.4** |
| T3′: $\Pr(\Delta b = inc \mid t_2 < t_3)$ baseline | 4.7 | 4.1 | 10.0 | 7.9 | 18.9 | 14.3 |
| $\Delta \Pr(\Delta b = inc \mid$ "$>$" versus "$<$") | **3.1** | 3.3 | −2.6 | **3.3** | −5.2 | −4.5 |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | $<$**1.0** | $<$**1.0** | $<$**1.0** | $<$**1.0** | 13.4 | **3.9** |
| T4′: $\Pr(\Delta b = inc \mid$ *early*) baseline | 5.9 | 4.1 | 7.4 | 8.9 | 14.8 | 9.6 |
| $\Delta\Pr(\Delta b = inc \mid$ *late* versus *early*) | **2.5** | **5.6** | 1.3 | **3.5** | −0.7 | 1.7 |

*Notes*. Bold values are rejections significant at the 5% level (16 out of the 30 tests reject even under the strictest possible Bonferroni threshold). All numbers are probabilities or differences in probabilities scaled between 0 and 100. T1, T2, T3′, and T4′ entries show the test statistics (themselves probabilities). T4 entries show the $p$-values of the WMW rank-sum test.

## 5. How Does Nonsealed Behavior Bias Demand Estimates Based on the Sealed-Bid Abstraction? A New Model with Reactive Bidders

Taken jointly, tests T1–T5 show that eBay bidding does not satisfy the sealed-bid abstraction. What kind of biases can one expect from IPV demand estimates based on the abstraction? The direction of the bias is clear: if the rejection is due to some sort of incremental bidding, then $\varepsilon_1$ is often less than the highest valuation. Therefore the demand estimates are biased downward. To gauge how large this bias can be, this section compares the estimates to our a priori beliefs as well as to estimates from an alternative model of nonsealed bidding behavior.

### 5.1. Thin Tails of the Demand Estimates Lack Face Validity

We start by estimating demand under the sealed-bid assumption and comparing the estimated amount of variation in preferences with common sense. To illustrate the nonparametric demand estimates based on the sealed-bid abstraction, we ignore the auction-specific observables and compute the implied empirical distribution of valuations of the top product in the DVD movie and the MP3 player categories. For each product, we compute an estimate of the distribution of valuations above the median price, as well as an estimate of the distribution of valuations above a slightly higher price. These computations are quite simple—each estimate is a histogram of $b_1$ such that $b_1 > cutoff$ and $b_2 < cutoff - inc$, where *cutoff* is an arbitrary number—for example, the median price.

Table 6 shows the estimated distributions of valuations, and we propose that the distributions have unrealistically thin tails. For example, the probability that a valuation of a DVD is within \$2 of the cutoff (either the median price $+ inc =$ \$11 or \$12), given

**Table 6 Nonparametric Estimates of Demand Based on the Sealed-Bid Abstraction**

| DVD movie (*Black Hawk Down*) | | | MP3 player (Diamond Rio 500, new) | | |
|---|---|---|---|---|---|
| | *cutoff* = \$11 | *cutoff* = \$12 | | *cutoff* = \$127.50 | *cutoff* = \$137.50 |
| $x$ | $\Pr(v \le x \mid v > \$11)$ | $\Pr(v \le x \mid v > \$12)$ | $x$ | $\Pr(v \le x \mid v > \$127.50)$ | $\Pr(v \le x \mid v > \$137.50)$ |
| \$12 (%) | 54 | n/a | \$135 | 70 | n/a |
| \$13 (%) | **92** | 85 | \$145 | **85** | 55 |
| \$14 (%) | 94 | **91** | \$155 | 94 | **87** |
| \$15 (%) | 96 | 95 | \$165 | 95 | 92 |
| \$16 (%) | 98 | 96 | \$175 | 97 | 95 |
| \$17 (%) | 98 | 98 | \$185 | 97 | 95 |
| No. of obs | 52 | 55 | No. of obs | 122 | 121 |

*Notes.* Each column is the cdf of the valuation within the population of bidders, conditional on the valuation exceeding a specific cutoff. Bold numbers are discussed in the text. No. of obs is the number of observations used to construct the estimate.

that it exceeds the cutoff, is estimated to be more than 90%. Analogously, the probability that a valuation of an MP3 player is within \$20 of the cutoff (either the median price + *inc* = \$127.50 or \$137.50), given that it exceeds the cutoff, is estimated to be more than 85%. We propose that these estimates lack face validity because the valuations are unlikely to be so concentrated within the population.

### 5.2. Motivation for a New Model of Nonsealed Behavior: Generalized Reactive Bidders

Whereas the previous subsection casts doubt on the validity of demand estimates based on the sealed-bid abstraction, we need an alternative model to quantify the extent of bias and investigate the effect on managerial decisions. In constructing the new model, we restrict attention to the OverInc data because auctions that end with $b_{1,j} \le b_{2,j} + inc_j$ are best explained with the high bidder being an incremental bidder, who teaches us nothing new in the conditional order statistic sense.[6] Therefore, if we can learn about demand using the conditional order statistic approach, we need to learn from the OverInc data subset.

To accommodate the rejection of the sealed-bid abstraction, we propose that there are two types of bidding styles in addition to the incremental style mentioned above: *sealed* bidders who bid in a sealed fashion (just once per auction and without regard for other bids) and *reactive* bidders who initially bid only a fraction of their valuation. When they get outbid, reactive bidders return to the auction and bid again as long as their valuation exceeds the minimum acceptable bid at the moment and the auction has not ended. One empirical reason for focusing our model on differences between bidders is our finding that bidder heterogeneity, rather than auction timing, product, or seller heterogeneity, is likely to drive T1's and T2's rejections (see §3.2 for details). Preliminary empirical evidence for the two bidder types comes from comparing high "multibidders" who bid multiple times in a single auction with high "single bidders" who bid only once. In the movie data, the difference between the top two bids is smaller when the high bidder is a multibidder than when the high bidder is a single bidder (see Table A.3 in the appendix). Therefore the conditional distributions of $b_1$ given $b_2$ are different, and the two types of bidders should not be pooled as they are in the sealed-bid model. Moreover, the multibidders' top bids seem to underestimate their true valuations.

From a theoretical perspective, several possible rationalizations of the reactive bidding style exist. Hossain (2008) presents a model of a bidder who does not know his valuation but can tell whether it exceeds any given price once he sees the price. Hossain shows that such a bidder would bid low and often to learn. A related theoretical possibility is that bidders are uncertain about their willingness to pay for the given item, but they can collect costly information to reduce their uncertainty. Rasmusen (2006, p. 1) examines this possibility and concludes that because bidders want to incur the learning cost only if they are going to win, a bidder may reasonably "increase his bid ceiling in the course of an auction" and react to other bidders. However, neither bidder would bid all the way up to his valuation unless competition forced him to do so.

One way to integrate the sealed and reactive bidders under the same theoretical umbrella is to assume bidders do not like to relinquish control of bidding to the proxy-bidding agent eBay provides. Letting the proxy-bidding agent control bidding has the downside of exposing oneself to potentially paying a high price for the item. This exposure may reduce utility during the auction, either for purely psychological reasons or because the bidder would like to keep the

[6] When the high bidder in auction $j$ is an incremental bidder, we do not learn anything about the highest valuation in the auction other than that it exceeds the second-highest valuation—something we knew already.

option of backing out of the auction whenever outbid. At the same time, however, coming back later in the auction is costly, so the bidder might as well put in a bid upon arriving at the auction. Therefore the bidder may have to weigh the disutility of exposure against the transaction costs necessary to remain completely in control (by committing to paying only an increment above the current highest price and likely having to return often). By varying the size of the cost of returning to the site across bidders, all three bidding styles discussed so far emerge as optimal for different types of bidders: when the transaction cost is so high that bidders will not return back to the same auction, bidders bid on arrival, and the sealed-bid abstraction characterizes bidding. When the transaction cost is so small that bidders can return to the auction often, the purely incremental model arises. Finally, when the transaction cost of coming back to the auction site is neither irrelevantly small nor prohibitively high, the bidders use the proxy-bidding agent at least somewhat to reduce the number of times they have to come back to the auction site. For example, the bidders may behave like our reactive bidders and initially submit only a fraction of their valuation. We do not claim to structurally identify any particular theory in our data. Instead, we merely assess the amount of bias that arises when some bidders bid reactively, and we use the above theories to motivate our reduced-form statistical model of reactive bidding. The next subsection lays out the model assumptions.

### 5.3. Model Assumptions

There are $J$ auctions indexed by $j = 1, 2, \ldots, J$ and $I$ bidders indexed by $i = 1, 2, \ldots, K$. The bidders choose which auctions to participate in. The valuation $v$ of bidder $i$ in auction $j$ is a function of the auction-level characteristics (characteristics of the product and the seller) denoted $Z$, and a Normal$(0, \sigma^2)$ private value shock $\varepsilon$ independent across bidders and auctions:

$$\log v_{j,i} = \theta Z_j + \varepsilon_{j,i}. \tag{8}$$

The independence of $\varepsilon_{j,i}$ across bidders within each auction is the assumption of IPV, whereas independence within a bidder across auctions merely simplifies the pooling of data across auctions. Our estimation method will rely only on the observation of the highest valuation draw in each auction, so we will only need independence of the highest valuation draws across auctions, that is, the valuations of the auction winners. Because each bidder's private value shocks are likely positively correlated across auctions and their variance may change as the same bidder participates in repeated auctions for similar goods (Park and Bradlow 2005), we only use one auction for each high bidder in our estimation. The remaining assumption we need is therefore independence of $\varepsilon_{j,i}$ across different bidders who happen to win different auctions.

Bidders have latent bidding styles $\{R = \text{reactive}, S = \text{sealed}\}$. Sealed bidders bid only once, and they bid their valuation. The timing of their bid is irrelevant; some of them may even bid at the last moment of the auction. Reactive bidders bid on arrival, but they do not bid all the way up to their valuation. Instead, they start out by bidding low and come back to the auction to bid again whenever they are outbid. When they bid again, they raise their bid toward their valuation, but not all the way to it. When two reactive bidders 1 and 2 with valuations $v_1 > v_2$ meet in an auction, bidder 2 is thus eventually driven up to his valuation, but the final bid of bidder 1 will be lower than his valuation: $b_1 < v_1$. For tractability, we assume the final bids of reactive winners are, on average, some proportion $0 < \alpha \leq 1$ of their valuations: $b_1 = \alpha v_1$. This reduced form can be motivated by underlying utility maximization; please see the previous section for details.

Estimation of $(\alpha, \theta, \sigma)$ would be straightforward if we observed each bidder's style. Unfortunately, we observe bidding styles only partially: a reactive bidder reveals himself by bidding multiple times in an auction; we call such bidders "definitely reactive." However, a bidder who happens to bid only once in an auction is not necessarily a sealed bidder; he could be a reactive bidder who did not get a chance to increase his bid and reveal his style. Accounting for detection of reactive bidders is important because we find that the more auctions a bidder participates in, the more likely he is to be detected as reactive. By considering eventual winners who participated in more than five auctions within our DVD data set, we find that about two-thirds of them have bid multiple times in at least one of their auctions, so about two-thirds of the bidders are reactive. However, only 22% of eventual winners who participated in just one auction revealed themselves as reactive by bidding multiple

**Table 7 Partial Detection of Bidding Styles**

| No. of auctions in which the bidder is observed | No. of such bidders | Proportion of bidders revealed as reactive: Actual (%) | Proportion of bidders revealed as reactive: Predicted (%) |
|---|---|---|---|
| 1 | 2,602 | 22 | 22 |
| 2 | 795 | 39 | 38 |
| 3 | 312 | 51 | 55 |
| 4 | 116 | 60 | 63 |
| 5 | 60 | 66 | 63 |
| 6 | 31 | 70 | 68 |

*Notes.* The table considers all 3,971 bidders in our DVD data set who won at least one auction. A bidder reveals himself as reactive by bidding multiple times within at least one auction in which he participates. The "Actual" column shows the data, whereas "Predicted" shows predictions of our model.

times (see Table 7). Therefore we henceforth call people observed bidding just once "possibly sealed" and we need a model to assess their probability of actually being reactive.

To explain the empirical regularity in the "Actual" column of Table 7, we first relate the probability of being reactive to bidder observables and then wrap it inside a simple binomial model of detection. Each bidder $i$ has a propensity $M_i$ to be a reactive bidder. This propensity is a function of individual characteristics $W_i$ and a shock $\tau_i$ that captures individual-specific deviations from the trend: $M_i = \gamma W_i + \tau_i$. The shock $\tau$ is distributed Gumbel, leading to the familiar logistic regression $\Pr(R \mid W) = (\exp(\gamma W))/(1 + \exp(\gamma W))$. In the application of this model, we focus on individual characteristics ($W$) related to bidder experience. Our investigation of likely causes behind the rejections by tests T1 and T2 discussed in §3.2 motivates this choice.

If we observed bidding styles, we could estimate a logistic regression of style on $W$. However, we only observe styles partially, so the logistic regression model needs to be augmented with a model of detection. Let the detection probability of a reactive bidder in any given auction be $\delta$. Given the partial observability of bidding styles, the probability of actually observing a multibidder in $T$ auctions is

$$\Pr(multi = 1 \mid W, T) = [1 - (1-\delta)^T]\Pr(R \mid W) = [1 - (1-\delta)^T]\frac{\exp(\gamma W)}{1 + \exp(\gamma W)}. \quad (9)$$

The resulting likelihood is easily maximized to obtain an estimate of $(\gamma, \delta)$.

To allow the separation of inference about style from the inference about valuations, we assume the two shocks $\tau$ and $\varepsilon$ are independent of each other. This assumption rules out a relationship between private values and bidding style (other than that captured by observables). For example, bidders with higher valuations may tend to choose the reactive bidding style to protect their surplus, and hence $\mathrm{corr}(\varepsilon, \tau) > 0$. Alternatively, bidders with higher valuations may tend to snipe at the end of the auction to avoid inviting competition from naïve incremental bidders, and hence $\mathrm{corr}(\varepsilon, \tau) < 0$ because sniping is inherently a sealed strategy. To check that our assumption of independence is reasonable for our data, we consider proxy bids of definitely reactive winners (winners observed bidding multiple times) and calculate the correlation between these bids and the $\Pr(S \mid W) = 1 - \Pr(R \mid W)$. Because these bidders reveal themselves as definitely reactive, a larger $\Pr(S \mid W)$ indicates a larger shock $\tau$. Because these bidders are assumed to bid $b_1 = \alpha v_1$, their proxy bids are correlated with their valuation shocks $\varepsilon$. Therefore, $\mathrm{corr}(\Pr(S \mid W), b_1 \mid \text{definitely } R)$ is a measure of $\mathrm{corr}(\varepsilon, \tau)$. For the top five DVD titles that all sell for approximately equal prices, on average, this correlation is $-0.07$. In other words, higher bids are only slightly negatively associated with the reactive bidding style. We conclude that assuming $\tau$ and $\varepsilon$ are independent is reasonable.

### 5.4. Identification and Estimation Strategy

We do not use bid magnitude to identify the bidding style. Instead, each bidder's bidding style is inferred from his multibidding status and observables $W$ using the Bayes' theorem

$$P_R \equiv \Pr(R \mid W, T, multi) = \begin{cases} 1 & \text{if } multi = 1, \\ \dfrac{(1-\delta)^T \exp(\gamma W)}{1 + (1-\delta)^T \exp(\gamma W)} & \text{if } multi = 0. \end{cases} \quad (10)$$

To make sure $(\gamma, \delta)$ and the implied $P_R$ are precisely estimated, we use a large sample of several thousand eventual winners. We simulated from the asymptotic distribution of $(\gamma, \delta)$ and found that these several thousand observations result in $\Pr(R \mid W, T, multi = 0)$ estimates with an error of about 1% on average, and less than 1.5% for every bidder. Therefore we henceforth treat $P_R$ as data.

Given $P_R$ for every bidder, the demand parameters are nonparametrically identified from the conditional distribution of $b_1$ given $b_2$, generalizing the approach of Song (2004). The crucial parameter $\alpha$ that captures bid shading by reactive bidders is identified from the assumption that the underlying distribution of valuations is the same for reactive and sealed bidders. If $\alpha$ were equal to 1, the conditional distribution of $b_1$ given $b_2$ should not depend on the bidding style because each bidding style would bid $b_k = v_k$. (Table A.3 in the appendix provides model-free evidence that the conditional distribution of $b_1$ given $b_2$ does depend on the bidding style and that reactive bidders shade their bids down ($\alpha < 1$): the difference $\Delta b = b_1 - b_2$ is significantly lower among auctions won by definitely reactive bidders as compared with the auctions won by possibly sealed bidders.) Moreover, the second-highest bids $b_2$ (and hence prices) are marginally higher with definitely reactive winners, so the aforementioned smaller $\Delta b$ is not merely due to auction heterogeneity. The bidding model with $\alpha < 1$ predicts both of these inequalities: reactive winners ensure $b_2 = v_2$ because they do not snipe, whereas sealed winners imply that $b_2 < v_2$ whenever they snipe and meet a reactive runner-up. Hence, prices should be somewhat higher whenever a winner is reactive. $\Delta b$ should be smaller with a reactive winner not only because of the higher $b_2$ but also because of $b_1 = \alpha v_1 < v_1$ with reactive winners, whereas $b_1 = v_1$ with sealed winners.

To estimate the demand parameters, we use the maximum likelihood approach. The likelihood of a single bid observation is a weighted average of the two possible style-specific likelihoods: $L(\alpha, \theta, \sigma \mid P_R, b_1, b_2) = P_R L_R(\alpha, \theta, \sigma \mid b_1, b_2) + (1-P_R)L_S(\theta, \sigma \mid b_1, b_2)$. The two style-specific likelihoods are, in turn, implied by the assumptions about the bidding strategies of the two bidder types. Song (2004) shows that when both top bidders are sealed, $b_1$ is a draw from the distribution of valuations truncated below at $b_2 + inc$ (adding $inc$ to the truncation point $b_2$ is necessary in our setting to account for the fact that we are selecting only OverInc data):

$$L_S(\theta, \sigma \mid \log b_1, \log(b_2+inc)) = \frac{(1/\sigma)\phi((\log b_1 - Z\theta)/\sigma)}{1-\Phi((\log(b_2+inc)-Z\theta)/\sigma)}, \tag{11}$$

where $\phi$ and $\Phi$ and the standard Normal pdf and cdf, respectively. As long as the high bidder is sealed, the same $L_S$ likelihood applies even when the second-highest bidder is reactive. Consider all bidders other than the two observed bidding $b_1$ and $b_2$. Obviously, other sealed bidders above $b_2 + inc$ cannot exist, and the assumption that reactive bidders do not snipe and compete up to their valuations implies other reactive bidders above $b_2 + inc$ cannot exist either. Therefore, even if the second-highest bidder is reactive, $b_1$ is still a draw from the distribution of valuations truncated below at $b_2 + inc$. Note, however, that the observed $b_2$ may be lower than the second-highest valuation whenever the sealed high bidder bids at the last moment of the auction. Therefore, a resulting advantage of our approach is that we do not need to assume $b_2$ is the second-highest valuation—a key assumption of several alternative demand estimation approaches (Athey and Haile 2005, Yao and Mela 2008).

When the high bidder is reactive and the second-highest bidder is sealed, the likelihood $L_R$ corrects for the fact that $b_1 = \alpha v_1 \Rightarrow \log v_1 = \log b_1 - \log \alpha$:

$$L_R(\alpha, \theta, \sigma \mid \log b_1, \log(b_2+inc)) = \frac{(1/\sigma)\phi((\log b_1 - \log\alpha - Z\theta)/\sigma)}{1-\Phi((\log(b_2+inc)-\log\alpha-Z\theta)/\sigma)}. \tag{12}$$

As long as the high bidder is reactive, $b_2$ corresponds to the second-highest valuation because all remaining bidders compete up to their valuations and have enough time to do so by the assumption of no sniping by reactive bidders. Therefore the same $L_R$ likelihood applies even when the second-highest bidder is reactive.

### 5.5. Data and Results: Distribution of Valuations and Model Fit

We focus on the DVD movie data because it conforms to the sealed-bid abstraction the most in that only the tests T1, T2, T3, and T3′ reject. Moreover, test T5 does not reject, so demand estimation based on the conditional distribution of the top two bids is statistically plausible.

**Table 8** **Model of Reactive Bidding Style and Detection (All Winners, $N = 3{,}971$)**

| Variable | Parameter estimate | Standard error | $t$-Value |
|---|---|---|---|
| $\delta$: Detection probability | 0.27 | 0.02 | 12.90 |
| Constant | 4.25 | 0.72 | 5.88 |
| Logfeedback | −0.52 | 0.11 | −4.68 |
| Years on eBay | −0.25 | 0.10 | −2.61 |

*Notes.* Partially observed logistic regression. The standard errors are asymptotic standard errors derived from the Hessian.

To maximize the precision of our bidding-style inference, we pool across all auctions in the data set, effectively considering 3,971 bidders who won at least one auction for a DVD. Focusing only on eventual winners is consistent with the goal of estimating the type of the highest bidder needed for subsequent construction of the demand model likelihood. Our model of partial observability fits the data well: the estimated per-auction probability of detection of $\delta = 0.27$ produces a close match between the actual and predicted probability of detecting a multibidder among bidders who participated in $T$ auctions (see Table 7). Overall, about two-thirds of the bidders are reactive.

The substantive results about bidding style correlates (see Table 8) replicate previous findings that multibidding is associated with less experience as measured by feedback score or by years since registration on eBay (Wilcox 2000, Borle et al. 2006). The correlations of multibidding with bidder experience are substantial. For example, increasing everyone's feedback by one standard deviation (250 points) reduces the average probability of multibidding from 82% to 67%. A bidder who has been registered on eBay for two years but only has the 20th percentile feedback (18 points) has a 91% probability of being a multibidder. In contrast, an 80th feedback-percentile bidder (142 points) with the same tenure has a 75% probability.

To assess the precision of our bidding-style inference, we drew 10,000 draws from the asymptotic distribution of $(\gamma, \delta)$ and computed the posterior probability of a possibly sealed bidder being reactive $\Pr(R \mid W, T, multi = 0)$. This simulation effectively draws from the posterior distribution of $\Pr(R \mid W, T, multi = 0)$, and thus the standard deviation of the probability draws can be interpreted as asymptotic error of $\Pr(R \mid W, T, multi = 0)$. The average error is 0.013, and no bidder's error is above 0.015. We conclude that $\Pr(R \mid W, T, multi = 0)$ is thus estimated very precisely, and thus $P_R$ is estimated very precisely as well.

**Table 9 Two Estimates of Demand (Top Five Movies, OverInc)**

| | Proposed model | | | Sealed-bid model | | |
|---|---|---|---|---|---|---|
| Second stage: Demand for top five movies ($N = 593$) | Parameter estimate | Standard error | $t$-Value | Parameter estimate | Standard error | $t$-Value |
| Constant | **1.87** | 0.30 | 6.14 | **1.20** | 0.44 | 2.71 |
| Movie1 | −0.02 | 0.17 | −0.13 | −0.02 | 0.19 | −0.07 |
| Movie2 | 0.14 | 0.18 | 0.78 | 0.17 | 0.20 | 0.86 |
| Movie3 | 0.06 | 0.20 | 0.30 | 0.12 | 0.23 | 0.54 |
| Movie4 | −0.33 | 0.24 | −1.36 | −0.23 | 0.25 | −0.94 |
| Top seller | −0.03 | 0.11 | −0.26 | −0.05 | 0.13 | −0.40 |
| New disc | 0.06 | 0.10 | 0.62 | 0.07 | 0.11 | 0.59 |
| Seller_has_store | **0.39** | 0.19 | 2.11 | **0.49** | 0.24 | 2.06 |
| $\sigma$ | **0.33** | 0.04 | 7.45 | **0.35** | 0.06 | 6.48 |
| $\alpha$ | **0.53** | 0.09 | 5.43 | Fixed to 1 by assumption | | |
| Second-stage log likelihood | | 781.90 | | | 777.1 | |

*Notes.* The standard errors are asymptotic standard errors derived from the Hessian. Bold parameters are significant at the 5% level.

For demand estimation, we considered all auctions for the top five movie titles while controlling for movie differences with fixed effects (recall that test T5 still did not reject the sealed-bid abstraction under this specification). Our data set holds 615 of such auctions, and we dropped 22 because they involved the same bidder winning a second time.[7] The resulting estimation data set has 593 auctions, each won by a different bidder. (Table A.3 in the appendix shows the descriptive statistics of this data set.) Using the maximum likelihood estimator described in the previous section, we estimated two models of the demand. First, we estimated the proposed model. Second, we set $P_R = 0$ to estimate demand under the sealed-bid assumption, effectively imposing the $L_S$ likelihood on all observations.

The proposed model clearly fits the data better based on a likelihood ratio test because it nests the sealed model and contains only one additional parameter. Specifically, the relevant likelihood ratio test is $1 - \chi^2_{\text{cdf}}[2(\log L_{\text{proposed}} - \log L_{\text{sealed}})] = 0.002$. Moreover, the proposed model fits the data better than the sealed-bid model when one considers the difference between multibidders and single bidders: the sealed model predicts that $\Delta b$ does not depend on multibidding status, whereas the proposed model obviously captures the fact that multibidding winners produce top bids much closer to $b_2$ than single bidders and hence smaller $\Delta b$.

Compared to the sealed-bid model, the proposed model implies that demand has a higher mean and higher variance. The lognormal assumption makes seeing the differences directly from Table 9 difficult, but plotting the two implied distributions reveals the differences clearly. Figure 2 shows the two estimated distributions for a new disc of the most popular movie (*Black Hawk Down*) sold by a top seller. The underestimation because of the sealed-bid assumption that §5.1 suggests is apparent and large. Moreover, the sealed-bid model gives a hard-to-believe estimate that no bidder values the movie more than $10.

Note that the valuations we are estimating are *net* of outside opportunities. As long as the bidder's dollar utility from owning the movie exceeds the lowest price available to the bidder in the outside market, the bidder's "valuation" in the eBay auction is that lowest outside price. Therefore we can interpret the variance of the valuations as a measure of dispersion of outside prices discovered by the bidders. The large variance we find suggests search costs on the Internet are indeed high as Brynjolfsson and Smith (2000) and many others suggest.

The demand estimates based on the proposed model shed light on the underlying process behind the commonly observed effect of seller reputation on

**Figure 2 Customer Valuation of a Movie: Two Estimates**

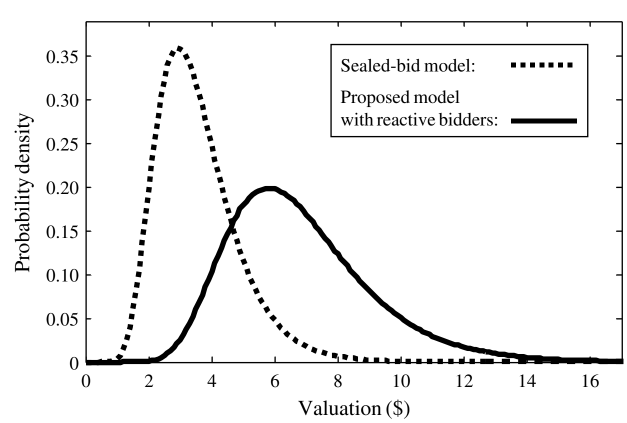


*Notes.* Each line represents the pdf of the population distribution of valuations of a new DVD of *Black Hawk Down*, sold by a "top seller."

[7] Dropping repeated observations of the same bidder ensures $\varepsilon$ is not correlated across auctions (see §5.3).

prices (Resnick et al. 2006). The top seller dummy does not have a significant positive coefficient, so consumers seem to not actually value DVDs from reputable sellers higher. Instead, prices reputable sellers obtain are likely higher because higher reputation increases bidder entry. On the other hand, the results suggest sellers with an eBay store do create value, which could be due to customer service or simply branding by the store owners.

Our model and estimation results are consistent with our tests of the sealed-bid abstraction as applied to the DVD data set in §3. The model is obviously consistent with the prediction that high bids will sometimes (but not always) come after the second-highest bids (T1) and that when they do so, the difference $b_1 - b_2$ should be smaller (T3). Note that we estimate on the OverInc data only, so we effectively eliminate purely incremental bidders that explain the T2 and T3′ rejections. Finally, we estimate on DVD data that did not reject sealed bidding according to T5, and our model does not predict such a rejection.

One cost of the apparent simplicity of the proposed model of reactive bidding is that it cannot explain a T5 test statistic below 1/2 found in our MP3 and car data sets. To evaluate the predicted T5, we simulated 1,000 auctions with auction $j$ having a random number (3–10) of bidders with valuations drawn from a lognormal(0, 1). In each simulated auction, we set $b_2 = v_2$ and $b_1$ according to the model with $\alpha = 0.5$ and half bidders reactive and half sealed. The resulting test statistic was 1/2 on average. Inability to explain a T5 rejection makes the proposed model applicable to data (such as our DVD data set) with milder violations of the sealed-bid abstraction detected by tests T1–T4 but not T5. For data with more serious violations as indicated by T5 (such as our car data set), a different model probably applies. Our simulations indicate that to produce T5 < 1/2, $b_1$ needs to be linked more tightly to $v_2$. For example, letting $b_1 = v_2 + 0.25$ results in a T5 statistic of about 0.44. The problem with such a model is that it reveals no information about $v_1$, and thus the conditional order statistic approach cannot be used to estimate demand. The development of an alternative model that could exploit conditional order statistics even when T5 < 1/2 is beyond the scope of this paper.

### 5.6. Results: Managerial Implications

To assess the impact of the different estimates on managerial decisions, we compute the optimal public reserve prices the two models imply. In the context of eBay, the public reserve price is called the "starting price." Myerson (1981) shows how to compute such optimal public reserves and also demonstrates that they do not depend on the number of bidders. This finding is a critical benefit for the eBay application, where the number of bidders is unobserved. The

**Figure 3** **Optimal Reserves (Starting Prices) a Seller Should Charge**

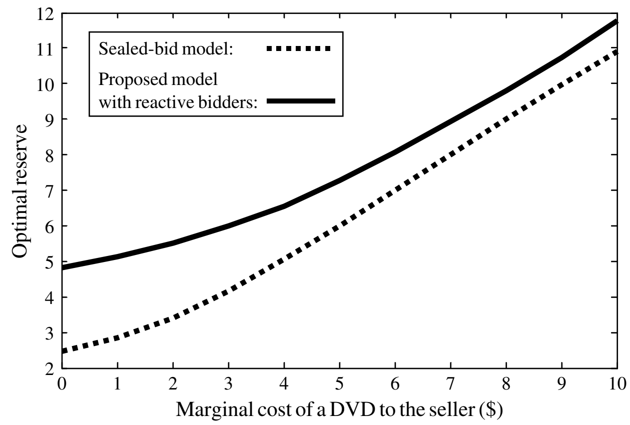


*Notes.* Each line is the optimal public reserve (the "starting price" in eBay terminology) for a new DVD of *Black Hawk Down*, sold by a "top seller" (calculated according to Myerson 1981).

optimal reserve to use does depend on the marginal cost the seller is facing. Because we do not know this cost, we compute the optimal reserves for a range of reasonable costs (between $0 and $10). The differences between the two models are large (see Figure 3): the proposed model suggests much higher reserves. Our best guess of a wholesale cost of a DVD is about $3. For this cost, the optimal reserve based on the proposed model would be about $6.00—43% higher than the $4.20 the sealed-bid model suggests. We conclude that rejecting the sealed-bid abstraction in favor of our proposed model can have large managerial implications.

It is also interesting to compare the optimal reserves with the actual starting prices. As Figure A.2 in the appendix shows, the actual starting prices are quite dispersed. Whether the dispersion is due to differences in marginal cost or lack of seller knowledge about demand is unclear. However, our model suggests that about half of the reserves are too low: 37% of the starting prices are below the $4.80 level we find optimal for marginal cost of zero. Note that our data set excludes the few auctions that had secret reserve prices, so the low starting prices we find are indeed offers to sell at that price.

Besides the different guidance for sellers, the proposed model also paints a different picture of the division of gains from trade on eBay. In a seminal paper, Bapna et al. (2008) used winning bids from their own sniping agent to calculate the difference between the winning proxy bid ($b_1$ in the notation of this paper) and price. This difference is a lower bound on the consumer surplus. Bapna et al. found that the median bound on the surplus was $4—about 26% of the median price of $15. This percentage is lower than the one we find under the proposed model (91%). In contrast, the estimate one would obtain from simply calculating the difference between $b_1$ and the price—the

**Table 10 Different Estimates of Buyer and Seller Surplus**

| Variable | Buyer surplus: Bapna et al. (2008) | Buyer surplus: Sealed model | Buyer surplus: Proposed model | Seller surplus: Lower bound (cost = starting price) | Seller surplus: Upper bound (cost = \$0) |
|---|---|---|---|---|---|
| Surplus (\$) | 4.00 | 0.82 | 10.03 | 5.00 | 11.00 |
| Surplus as percentage of price (%) | 27 | 7 | 91 | 43 | 100 |
| Price (\$) | 15.00 | | | 11.00 | |

*Notes.* Medians of price and surpluses. Buyer surplus is the difference between the valuation of the top bidder and final price. Seller surplus is the difference between the final price and seller cost.

surplus estimate under the sealed-bid assumption—would be much lower (only 7%). See Table 10 for details. We conclude that the sealed-bid model would make eBay seem like a seller's market, whereas the proposed model gives buyers much more surplus, more in line with previous work. It is also clear that the high bids of snipers using a sniping engine (as in Bapna et al. 2008) behave differently from high proxy bids submitted directly to eBay. As a result of the different behavior, the difference between the top two proxy bids in an eBay auction is not a tight bound on consumer surplus; in fact, it is only slightly tighter than saying the consumer surplus is at least as large as eBay's minimum increment (a computation that would yield about a 5% estimate of buyer surplus).

## 6. Discussion

We propose five different tests of the sealed-bid abstraction that bids from online auctions with proxy bidding can be analyzed "as if" they were bids from a second-price sealed-bid auction. When applied to three different eBay data sets, the tests consistently indicate that the sealed-bid abstraction does not describe eBay behavior well. The tests rely on minimal assumptions, with different assumptions implying different tests, so the abstraction can be rejected even when one of the assumptions does not hold. Moreover, the rejections are robust to alternative explanations, such as a lumpy distribution of valuations or the existence of unobserved differences between auctions. Not only does the sealed-bid abstraction not fit the data in general, it does not fit even carefully selected subsets of the auctions in which one might expect it to hold. Specifically, it does not fit the subset of auctions that ended with the top two bids more than one increment apart, or the subset of auctions in which the higher bid was actually submitted first.

Our rejection of the sealed-bid abstraction is not only of theoretical interest; it also suggests demand estimates that rely on the sealed-bid abstraction are likely to be biased. The direction of the bias is downward because we find that some sort of reactive bidding causes the rejection, and thus the top proxy bid is often less than the top valuation and too close to the second-highest bid. Because the top bid is too close to the second bid, models that rely on conditional order statistics for identification (e.g., Song 2004) are the most affected by the nonsealed bidding. Because the second-highest bid is also weakly lower than the second-highest valuation whenever the second-highest bidder bids reactively, even models that rely on the second-highest bid and an inferred number of bidders will be biased down (e.g., Zeithammer 2006, Yin 2006, Chan et al. 2007, Adams 2007, Yao and Mela 2008). To assess the magnitude of the bias, we propose an alternative empirical model of bidding in an eBay auction. Our goal is to build on the conditional order statistic approach because it does not rely on the difficult process of inferring the number of latent bidders in an eBay auction.

The proposed model assumes bidders have different inherent bidding styles and only some bidders conform to the sealed-bid abstraction. Reactive bidders, who do not conform, initially bid only a fraction of their valuation, and they raise their bid whenever they are outbid. They can be at least partially detected by bidding multiple times in a single auction, and we can rely on bidder characteristics, such as experience, to estimate the probability that any given bidder is reactive. The link between experience and multiple bidding replicates previous findings by Wilcox (2000) and Borle et al. (2006). The behavioral literature suggests that bidder behavior may be a function of experience (see List 2003, for example). Consistent with the results presented in List (2003) and Simonsohn and Ariely (2008), we find that experienced traders are more likely to behave "rationally" in the sense of conforming to the sealed-bid abstraction. We hope future research rationalizes the coexistence of multiple bidding styles in the market. Fehr and Tyran (2005) argue that if rational and naïve behaviors are strategic complements (in the game-theoretic sense), the naïve behaviors are more likely to persist in the market. If, on the other hand, rational and naïve behaviors are strategic substitutes, the market is more likely to force out the naïve behaviors. In online auctions, sniping and reactive bidding seem to be strategic substitutes in that sniping is a dominant strategy if one's opponent is a reactive bidder. Empirically, we do not see strong evolutionary pressure against reactive bidders because our estimates indicate that some 70% of bidders behave reactively.

Given the probability of the reactive style for each bidder, we can interpret the observed bid distribution as a weighted combination of bids coming from sealed-like bidders and bids coming from reactive bidders, with auction-specific style probabilities that depend on the winner's experience and multibidding status. We estimate the implied distribution of valuations in the DVD data set, and we compare the result to the estimates the sealed-bid assumption generates following Song (2004). The two demand estimates differ in all aspects of interest: compared to the sealed-bid model, the proposed mixture model implies demand has a higher mean and a higher variance. The differences, namely, the estimates of the magnitude of the bias, are large: the sealed-bid model suggests the distribution of valuations has a mean of \$3.55 and standard deviation of \$1.30; the proposed model suggests a much more realistic mean of \$6.82 and standard deviation of \$2.29. The managerial implications of the two estimates are also very different: the mixture model finds much higher starting prices to be optimal (\$6.00 versus \$4.20 for a marginal cost of \$3.00) and gives a much higher estimate of buyer surplus (91% of final price versus 7% of final price).

Our model is a reduced-form specification in that we do not explicitly derive the different bidding styles from underlying primitives such as utility and information. We hope our evidence spurs the development of such theoretical models that will better capture behavior and consequently allow for more fine-grained structural inference about drivers of behavior in online auctions with proxy bidding. In the rest of this discussion, we summarize how our evidence constrains these theories. Although the sealed-bid abstraction fails, the data do not conform to purely incremental bidding either: the high bid does not always come after the second, and it is not always exactly one increment above. Therefore our findings leave auction theorists with a puzzle: why would a bidder enter an online auction early and bid in some sort of incremental manner instead of waiting for the end and sniping? We can only speculate about the correct model here, and several possibilities emerge.

First, bidders may not like to relinquish control over bidding to the proxy-bidding agent eBay provides, but they also face a cost of returning back to the site and raising their bids when they are outbid. Heterogeneity of this cost across bidders could produce the heterogeneity in bidding styles that we assume here: lower-cost bidders find the reactive style optimal in that it balances the utility of not relinquishing control against the cost of doing so. A second set of assumptions that could generate reactive-like bidding is learning by bidding. Hossain (2008) presents a model of a bidder who does not know his valuation but can tell whether his valuation exceeds any given price once he sees the price. Hossain shows that such a bidder would bid low and often in an eBay auction to learn. If some bidders know their valuations and some do not, a mixture of bidding styles would again result. Third, multiple bids by the same bidder may arise as the bidders consider multiple simultaneous auctions. Peters and Severinov (2006) obtain nonsealed bidding in an equilibrium of a model with sequential arrivals of bidders to many multiple simultaneous auctions. Nekipelov (2007) also focuses on multiple concurrent auctions and shows bidders have an incentive to bid early to deter entry by rival bidders. With incentives for early bidding balanced against the incentive to hide private information, within-auction dynamics that do not satisfy the sealed-bid abstraction emerge in Nekipelov's model.

Our results are applicable beyond eBay, in any market where the institutional details make both sealed and incremental bidding possible. For example, a procurement auction that runs live on the Internet but also accepts proxy bids could be analyzed using our techniques. Another interesting area of application would be detection of collusion in sealed-bid auctions. If the top two bids in a sealed-bid auction are too related to each other, the auctioneer might conclude that the top two bidders are colluding.

**Acknowledgments**
The authors thank Jeremy Fox, Ali Hortaçsu, Patrick McAlvanah, Harry Paarsch, Raphael Thomadsen, and seminar participants at Chicago Graduate School of Business, University of California at Davis, and the University of Alberta for comments and suggestions. The authors also thank eBay for providing the data.

## Appendix

Proof of Proposition 1. Denote $(\varepsilon_{1,j}, \varepsilon_{2,j})$ as the top two order statistics from a sample of $N_j \geq 2$ iid draws from some distribution $F$. It is enough to show that in a data set selected by bounds $B_j$ such that for each $j$, $\varepsilon_{1,j} > B_j \geq \varepsilon_{2,j}$, $\Pr[\varepsilon_{1,j} > \varepsilon_{1,k} \mid \varepsilon_{2,j} > \varepsilon_{2,k}$ and $\min(\varepsilon_{1,j}, \varepsilon_{1,k}) > \max(B_j, B_k)] = 1/2$. Fix $\varepsilon_{2,j}$, and denote $G(z) \equiv \Pr(\varepsilon_{1,j} < z \mid \varepsilon_{2,j}) = (F(z) - F(\varepsilon_{2,j}))/(1 - F(\varepsilon_{2,j}))$. Since $\varepsilon_{2,j} \leq \max(B_j, B_k)$, the distribution of $\varepsilon_{1,j}$, given that it exceeds $\max(B_j, B_k)$, is just a truncation of $G$:

$$\begin{aligned}\Pr(\varepsilon_{1,j} < z \mid \varepsilon_{1,j} > \max(B_j, B_k))\\ &= (G(z) - G(\max(B_j, B_k)))/(1 - G(\max(B_j, B_k)))\\ &= (F(z) - F(\max(B_j, B_k)))/(1 - F(\max(B_j, B_k))).\end{aligned}$$

Analogously, fixing $\varepsilon_{2,k}$ yields $\Pr(\varepsilon_{1,k} < z \mid \varepsilon_{1,k} > \max(B_j, B_k)) = (F(z) - F(\max(B_j, B_k)))/(1 - F(\max(B_j, B_k)))$, so the two distributions do not depend on the respective $\varepsilon_{2,j}$ and are identical to each other: $\Pr(\varepsilon_{1,k} < z \mid \varepsilon_{1,k} > \max(B_j, B_k)) = \Pr(\varepsilon_{1,j} < z \mid \varepsilon_{1,j} > \max(B_j, B_k))$. The probability that two iid draws are ordered in a particular way is 1/2. Q.E.D.

### Example of Tail Comparison Test (T5)

For illustration purposes, this example considers the raw *Black Hawk Down* data, focusing on the OverInc auctions.

**Table A.1 Example of a Joint Distribution of $(b_1, b_2)$ with Decile-Sized Bins**

| | | | | | | $b_1$ | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | 0 | 0.05 | 0.15 | 0.25 | 0.35 | 0.45 | 0.55 | 0.65 | 0.75 | 0.85 | 0.95 |
| $b_2$ | 0.05 | *2* | 13 | 9 | 10 | 2 | 0 | 0 | 1 | 0 | 0 |
| | 0.15 | | *1* | 20 | 12 | 5 | 3 | 0 | 0 | 0 | 0 |
| | 0.25 | | | *9* | 43 | 39 | 5 | 2 | 0 | 0 | 0 |
| | 0.35 | | | | *14* | **68** | **18** | **8** | **2** | **0** | **0** |
| | 0.45 | | | | | *30* | 58 | 35 | 20 | 2 | 1 |
| | 0.55 | | | | | | *14* | 50 | 11 | 1 | 0 |
| | 0.65 | | | | | | | *18* | 31 | 1 | 4 |
| *Histograms: Estimates of pdfs of tails* | $b_1 \mid b_2 \in (0.3, 0.4)$ and $b_1 \geq 0.4$ (%) | | | | | 71 | 19 | 8 | 2 | 0 | 0 |
| | $b_1 \mid b_2 \leq 0.3$ and $b_1 \geq 0.4$ (%) | | | | | 81 | 14 | 4 | 2 | 0 | 0 |

For all auctions $j$, the $(b_{1,j}, b_{2,j})$ data have been rescaled to fit between 0 and 1 by the transformation $b_{i,j} \to (b_{i,j} - \min_j(b_{2,j}))/(\max_j(b_{1,j}) - \min_j(b_{2,j}))$. Table A.1 shows the joint distribution of the rescaled $(b_1, b_2)$ with 10 equal-sized bins. For example, the number 43 in $b_2 = 0.25$ and $b_1 = 0.35$ bin means there are 43 $(b_1, b_2)$ observations such that $b_2$ is in the third decile and the corresponding $b_1$ is in the fourth decile: $b_2 \in (0.2, 0.3)$ and $b_1 \in (0.3, 0.4)$.

The outlined box contains a particular tail comparison test that can be made. The box contains 153 observations, and test T5 says the distribution of these observations

**Table A.2 Control Regressions**

| | | OLS regression of price | | Truncated regression | | C5 Corvettes ($N = 732$) | OLS regression of price | |
|---|---|---|---|---|---|---|---|---|
| Data set | Variable | Parameter | Std. error | Parameter | Std. error | Variable | Parameter | $t$-Value |
| Top player: Diamond | Constant | 4.733 | 0.027 | 3.587 | 0.357 | Log_mileage | −0.012 | −0.71 |
| Rio 500 ($N = 785$) | Top seller | 0.011 | 0.004 | 0.003 | 0.037 | Log_age | −0.272 | −7.28 |
| | Photo | −0.002 | 0.011 | −0.211 | 0.115 | No_age | −1.019 | −6.62 |
| | Bold | 0.079 | 0.028 | −0.268 | 0.327 | Automatic | −0.026 | −0.8 |
| | Mon | 0.089 | 0.020 | 0.322 | 0.200 | Convertible | 0.164 | 5.00 |
| | Tues | 0.078 | 0.021 | 0.056 | 0.218 | New | −0.749 | −6.39 |
| | Wed | 0.102 | 0.020 | 0.279 | 0.196 | Black | −0.072 | −1.78 |
| | Thur | 0.045 | 0.019 | 0.133 | 0.192 | Burgundy | −0.255 | −3.17 |
| | Fri | 0.004 | 0.019 | 0.168 | 0.190 | Red | −0.038 | −0.92 |
| | Sat | 0.049 | 0.020 | 0.060 | 0.204 | Silver | −0.088 | −1.84 |
| | New | −0.001 | 0.016 | −0.250 | 0.149 | Log_feedback | −0.022 | −2.47 |
| | | $R^2$ | 0.060 | $\sigma$ | 0.321 | Neg | −0.053 | −0.77 |
| | | | | | | Year_2002 | −0.211 | −2.01 |
| Top movie: *Black Hawk* | Constant | 2.271 | 0.020 | 0.062 | 1.003 | Year_2003 | −0.263 | −2.52 |
| *Down* ($N = 229$) | Top seller | 0.096 | 0.023 | −0.337 | 0.362 | Month_feb | 0.103 | 1.09 |
| | Store seller | −0.069 | 0.069 | 1.193 | 0.684 | Month_mar | 0.192 | 2.29 |
| | New | 0.073 | 0.023 | 0.414 | 0.322 | Month_apr | 0.113 | 1.37 |
| | | $R^2$ | 0.070 | $\sigma$ | 0.499 | Month_may | 0.107 | 1.33 |
| | | | | | | Month_jun | 0.123 | 1.50 |
| | | | | | | Month_jul | 0.157 | 1.94 |
| | | | | | | Month_aug | 0.228 | 2.78 |
| | | | | | | Month_sep | 0.122 | 1.55 |
| | | | | | | Month_oct | 0.120 | 1.53 |
| | | | | | | Month_nov | 0.027 | 0.35 |
| | | | | | | Month_dec | −0.074 | −0.92 |
| | | | | | | Num_bidders | 0.009 | 3.32 |
| | | | | | | Constant | 10.725 | 52.86 |
| | | | | | | | $R^2$ | 0.52 |

*Notes.* The table shows the control regressions used in the tail comparison test (T5). Only the regressions for the top movie, top MP3 player, and C5 Corvette are shown—the regressions for the other products in the respective categories are similar.

**Figure A.1 Quantile (97.5%) of T5 for 10 Different Distributions**

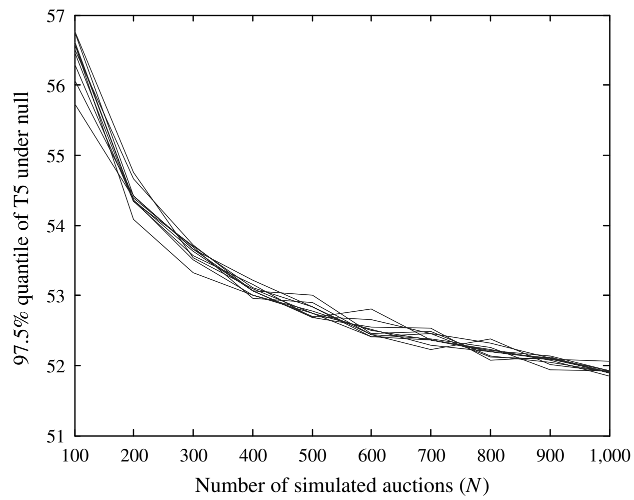


*Notes.* Each line corresponds to a different distribution. For each distribution and $N$, 1,000 simulations were used to compute the quantile.

in any row within the outlined box should be the same. To increase the power of the test, let us compare the shaded row (96 observations) to all the rows above it. The (normalized) histograms of these two tails are shown below the table, and they are clearly not the same: the shaded row seems to have higher central tendency.

### Pairs of Pairs: Properties of the Proposed Nonparametric Test

Suppose there are $N$ pairs of iid draws from some continuous distribution $F$. Denote the higher number in the $n$th pair $b_{1,n}$ and the lower number $b_{2,n}$. The T5 test statistic checks whether higher $b_2$s correspond to higher $b_1$s within the feasible pairs of pairs. The null prediction that $\text{T5} = 1/2$ does not depend on $F$ (Proposition 1), but the sampling distribution of the test statistic around 1/2 is difficult to derive analytically because $F$ influences the chance of finding feasible pairs of pairs as well as the pattern of dependence across those pairs of pairs. To investigate the actual properties of the test statistic, we simulated 1,000 data sets with different $N$ values ranging from 100 to 1,000 and with 10 different $F$ values (Normal, Lognormal, Gamma, Weibull, Beta, and Uniform, some with different parameters). We find that the distribution of the test statistic does not, in fact, seem to depend on $F$: it is approximately Normal even for only $N = 100$, and the 95% confidence interval shrinks down proportionally to $1/\sqrt{N}$. See Figure A.1 for a plot of the upper bound of the 95% confidence interval as a function of $N$. Each line in the figure represents one distribution, and all the lines clearly essentially coincide.

Consider a naïve theory one might have about the distribution of the test statistic is instructive. If there happened to be $K$ *independent* feasible pairs of pairs, the total number of feasible pairs with $[b_{1,k} > b_{1,n} \mid b_{2,k} > b_{2,n}]$ would be distributed Binomial$(p = 1/2, K)$, which would, in turn, be asymptotically approximated by Normal$(pK, p(1-p)K) =$ Normal$(K/2, K/4)$. Therefore, one might naïvely expect the proportion of feasible pairs with $[b_{1,k} > b_{1,n} \mid b_{2,k} > b_{2,n}]$ to be approximately Normal$(1/2, 1/4K)$. We find that the variance of the order statistic is, in fact, far greater than $K$, and it is related to $N$ instead of $K$. Specifically, we regressed the size of the confidence intervals on $4 * 1.96/\sqrt{4N}$ and found the coefficient to be about $0.313 \approx \sqrt{0.1}$ with $R^2 =$ 0.995. This regression implies that the distribution of the test statistic under the null hypothesis is extremely well approximated by Normal$(1/2, 1/40N)$. In other words, the amount of information in feasible pairs of a sample of $N$ pairs can be approximated by the information that *would be* contained in $10N$ *independent* feasible pairs of pairs.

**Figure A.2 Distribution of Actual Starting Prices for the Top Movie**

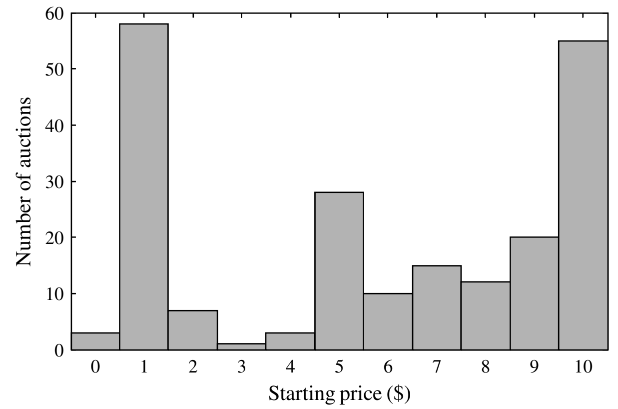


*Notes.* This figure is a histogram of starting prices; each bin is centered on a whole dollar amount. There are no starting prices over $10 in the data sample.

**Table A.3 Descriptive Statistics**

| All data ($N = 593$) | Prob reactive ($\text{Pr}_R$) | $b_1$ ($) | $b_2$ ($) | Top seller | Seller with store | New disc | $b_1 - b_2$ ($) |
|---|---|---|---|---|---|---|---|
| Mean | 0.814 | 12.17 | 10.48 | 0.506 | 0.526 | 0.069 | 1.70 |
| Median | 0.862 | 12.00 | 10.49 | 1.000 | 1.000 | 0.000 | 1.27 |
| Std. dev. | 0.199 | 2.55 | 2.10 | 0.500 | 0.500 | 0.254 | 1.36 |
| Max | 1.000 | 22.00 | 16.00 | 1.000 | 1.000 | 1.000 | 10.34 |
| Min | 0.095 | 5.51 | 4.00 | 0 | 0 | 0 | 0.51 |
| | High bidder definitely reactive ($\text{Pr}_R = 1$, $N = 204$) | | | | | | |
| Mean | 1.000 | 12.07 | 10.68 | 0.505 | 0.490 | 0.059 | 1.38 |
| Median | 1.000 | 12.00 | 10.50 | 1.000 | 0.000 | 0.000 | 1.01 |
| Std. dev. | 0.000 | 2.32 | 2.12 | 0.501 | 0.501 | 0.236 | 0.88 |
| Max | 1.000 | 20.00 | 16.00 | 1.000 | 1.000 | 1.000 | 5.50 |
| Min | 1.000 | 6.00 | 4.99 | 0 | 0 | 0 | 0.51 |
| | High bidder possibly sealed ($\text{Pr}_R < 1$, $N = 389$) | | | | | | |
| Mean | 0.716 | 12.23 | 10.37 | 0.506 | 0.545 | 0.075 | 1.86 |
| Median | 0.751 | 12.00 | 10.49 | 1.000 | 1.000 | 0.000 | 1.49 |
| Std. dev. | 0.181 | 2.66 | 2.09 | 0.501 | 0.499 | 0.263 | 1.52 |
| Max | 0.980 | 22.00 | 15.51 | 1.000 | 1.000 | 1.000 | 10.34 |
| Min | 0.095 | 5.51 | 4.00 | 0 | 0 | 0 | 0.51 |

*Notes.* The only 5% significant difference between the definitely reactive and possibly sealed subsamples is the latter having higher $b_1 - b_2$ ($p < 0.01$). The possibly sealed subsample also has marginally lower $b_2$ ($p = 0.08$). No other differences come close to significance.

**Table A.4** $p$**-Values of T1–T4 Test Statistics in Table 2**

| | MP3 players | | | | | | DVDs | | | | | | Cars | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Selected product | All | | Rio 500 | | KB Gear JamP3 | | All | | *Black Hawk Down* | | *A Beautiful Mind* | | All Fords | | Ford F150 | | C5 Corvette | |
| All auctions | | | | | | | | | | | | | | | | | | |
| No. of obs | 6,239 | $p$ | 1,325 | $p$ | 602 | $p$ | 3,512 | $p$ | 375 | $p$ | 274 | $p$ | 5,621 | $p$ | 510 | $p$ | 543 | $p$ |
| T1: $\Pr(t_1 > t_2)$ | **64.3** | $<10^{-6}$ | **61.1** | $<10^{-6}$ | **64.2** | $<10^{-6}$ | **72.9** | $<10^{-6}$ | **72.5** | $<10^{-6}$ | **75.6** | $<10^{-6}$ | **63.2** | $<10^{-6}$ | **68.9** | $<10^{-6}$ | **67.8** | $<10^{-6}$ |
| T2: $\Pr(\Delta b = inc)$ | **14.2** | $<10^{-6}$ | **12.6** | $<10^{-6}$ | **18.9** | $<10^{-6}$ | **14.4** | $<10^{-6}$ | **16.8** | $<10^{-6}$ | **13.5** | $<10^{-6}$ | **11.4** | $<10^{-6}$ | **14.1** | $<10^{-6}$ | **17.7** | $<10^{-6}$ |
| ($\Delta \Pr \Delta b = inc \mid$ "$>$" versus "$<$") | **14.4** | $<10^{-6}$ | **15.5** | $<10^{-6}$ | **16.1** | $<10^{-6}$ | **12.9** | $<10^{-6}$ | **12.7** | $<10^{-3}$ | **15.7** | $<10^{-6}$ | **10.5** | $<10^{-6}$ | **12.5** | $<10^{-4}$ | **10.5** | 0.002 |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | | $<10^{-6}$ | **<1** | $<10^{-4}$ | **3.6** | 0.0416 | **1.9** | 0.019 | 71.4 | 0.714 | 56.1 | 0.561 | 75.1 | 0.7510 | 30.9 | 0.3090 | 36.5 | 0.365 |
| ($\Delta \Pr \Delta b = inc \mid$ *late* versus *early*) | **2.9** | $<10^{-6}$ | **4.5** | 0.0100 | 3.7 | 0.2512 | **3.0** | 0.011 | 1.6 | 0.679 | 8.1 | 0.052 | −0.9 | 0.2860 | 0 | 0.9990 | 0.4 | 0.910 |
| OverInc ($\Delta b > inc$) | | | | | | | | | | | | | | | | | | |
| No. of obs | 4,420 | | 922 | | 418 | | 2,365 | | 229 | | 198 | | 4,162 | | 352 | | 370 | |
| T1: $\Pr(t_1 > t_2)$ | **64.8** | $<10^{-6}$ | **62.0** | $<10^{-6}$ | **64.2** | $<10^{-6}$ | **73.4** | $<10^{-6}$ | **73.8** | $<10^{-6}$ | **73.7** | $<10^{-6}$ | **60.4** | $<10^{-6}$ | **67.6** | $<10^{-6}$ | **66.7** | $<10^{-6}$ |
| T3: ($\Delta b \mid t_1 > t_2$) versus ($\Delta b \mid t_1 < t_2$) | **<1** | $<10^{-6}$ | 26.0 | 0.0030 | 14.9 | 0.1490 | **<1** | $<10^{-3}$ | 5.4 | 0.054 | 58.2 | 0.058 | **<1** | $<10^{-6}$ | 4.8 | 0.0480 | **3.8** | 0.038 |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | **<1** | | **<1** | 0.0050 | **1.4** | 0.0140 | 35.7 | 0.357 | 88.3 | 0.883 | 78.7 | 0.787 | 20 | 0.2 | 68.7 | 0.6870 | 19.6 | 0.196 |
| HighFirst ($t_1 < t_2$ and $\Delta b > 0$) | | | | | | | | | | | | | | | | | | |
| No. of obs | 1,591 | | 335 | | 154 | | 780 | | 81 | | 57 | | 1,559 | | 125 | | 128 | |
| T2: $\Pr(\Delta b = inc)$ | **5.6** | $<10^{-6}$ | **2.7** | 0.0030 | **10.4** | $<10^{-4}$ | **5.5** | $<10^{-6}$ | **8.6** | 0.007 | 1.8 | 0.321 | **4.9** | $<10^{-6}$ | **4.8** | 0.0137 | **11.7** | $<10^{-6}$ |
| T4: ($\Delta b \mid t$ *late*) versus ($\Delta b \mid t$ *early*) | **<1** | $<10^{-6}$ | **1.3** | 0.0128 | 58.1 | 0.5810 | 17.6 | 0.176 | 98.5 | 0.984 | 85.0 | 0.850 | 7.4 | 0.0736 | 89.2 | 0.8920 | 50.6 | 0.506 |
| ($\Delta \Pr \Delta b = inc \mid$ *late* versus *early*) | **2.7** | 0.022 | **4.2** | 0.0180 | 7.8 | 0.1150 | −0.3 | 0.875 | −2.5 | 0.696 | 3.6 | 0.326 | 0.2 | 0.8050 | 3.2 | 0.4060 | −7.8 | 0.173 |

*Notes*. This information is the same as in Table 2, with each test shown along with its $p$-value. Bold values are rejections of the sealed-bid abstraction significant at the 5% level. All numbers are probabilities or differences in probabilities scaled between 0 and 100. T1, T2, T3′, and T4′ entries show the test statistics (themselves probabilities), whereas T3 and T4 entries show the $p$-values of the WMW rank-sum test.

**Table A.5** **Notation**

T: Test
A: Assumption, a property of a model
$i$: Bidder
$j$: Auction
$x$: Private signal ($x$ is valuation in an IPV setting, also denoted $v$)
$\beta(x)$: Bidding function
$Z$: Auction-level observables (such as product sold or ending time)
$\varepsilon$: Shocks to bids because of private signals (private components of valuations in an IPV setting).
$F$: Population conditional distribution of bids given auction-level observables (distribution of private valuation components $\varepsilon$ in an IPV setting). $f$ is the corresponding pdf.
$b_k$: Magnitude of the $k$th highest bid in an auction, with $\Delta b = b_1 - b_2$
$t_k$: Time when the $k$th highest bid in an auction was submitted
$inc$: Minimum bidding increment
$W$: Bidder characteristics (such as experience)
$M$: Propensity of a bidder to bid in a reactive style
$\tau$: Individual shock to reactive-style propensity
$\delta$: Per-auction probability that a reactive bidder is detected by submitting multiple bids
$\alpha$: Shading parameter of reactive bidders; i.e., $\beta(x) = \alpha x$ with $0 < \alpha < 1$
$\theta$: Demand parameters, i.e., marginal effects of auction-level observables on valuations
$\Phi$: Standard Normal cumulative distribution function
$\phi$: Standard Normal probability density function
$\sigma$: Standard deviation

## References


Adams, C. P. 2007. Estimating demand from eBay prices. *Internat. J. Indust. Organ.* **25**(6) 1213–1232.

Ariely, D., A. Ockenfels, A. E. Roth. 2005. An experimental analysis of ending rules in Internet auctions. *RAND J. Econom.* **36**(4) 890–907.

Athey, S., P. A. Haile. 2002. Identification of standard auction models. *Econometrica* **70**(6) 2107–2140.

Athey, S., P. Haile. 2005. Nonparametric approaches to auctions. J. J. Heckman, E. E. Leamer, eds. *Handbook of Econometrics*, Vol. 6A. North-Holland, Amsterdam, 3847–3966.

Bajari, P., A. Hortaçsu. 2003. The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND J. Econom.* **34**(2) 329–355.

Bajari, P., A. Hortaçsu. 2004. Economic insights from Internet auctions. *J. Econom. Literature* **42**(4) 457–486.

Bapna, R., W. Jank, G. Shmueli. 2008. Consumer surplus in online auctions. *Inform. Systems Res.* **19**(4) 400–416.

Borle, S., P. Boatwright, J. B. Kadane. 2006. The timing of bid placement and extent of multiple bidding. *Statist. Sci.* **21**(2) 194–205.

Bradlow, E. T., Y.-H. Park. 2007. Bayesian estimation of bid sequences in Internet auctions using a generalized record-breaking model. *Marketing Sci.* **26**(2) 218–229.

Brynjolfsson, E., M. D. Smith. 2000. Frictionless commerce? A comparison of Internet and conventional retailers. *Management Sci.* **46**(4) 563–585.

Chan, T. Y., V. Kadiyali, Y.-H. Park. 2007. Willingness to pay and competition in online auctions. *J. Marketing Res.* **44**(2) 324–333.

Efron, B., R. J. Tibshirani. 1994. *An Introduction to the Bootstrap*. CRC Press, Boca Raton, FL.

Fehr, E., J.-R. Tyran. 2005. Individual irrationality and aggregate outcomes. *J. Econom. Perspect.* **19**(4) 43–66.

Games, P. A. 1977. An improved $t$ table for simultaneous control on $g$ contrasts. *J. Amer. Statist. Assoc.* **72**(359) 531–534.

Haile, P. A., E. T. Tamer. 2003. Inference with an incomplete model of English auctions. *J. Political Econom.* **111**(1) 1–51.

Hossain, T. 2008. Learning by bidding. *RAND J. Econom.* **39**(2) 509–529.

List, J. 2003. Does market experience eliminate anomalies? *Quart. J. Econom.* **118**(1) 41–71.

Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* **6**(1) 58–73.

Nekipelov, D. 2007. Entry deterrence and learning prevention on eBay. Working paper, Duke University, Durham, NC.

Park, Y.-H., E. T. Bradlow. 2005. An integrated model for bidding behavior in Internet auctions: Whether, who, when, and how much. *J. Marketing Res.* **42**(4) 470–482.

Peters, M., S. Severinov. 2006. Internet auctions with many traders. *J. Econom. Theory* **130**(1) 220–245.

Rasmusen, E. B. 2006. Strategic implications of uncertainty over one's own private value in auctions. *Advances Theoret. Econom.* **6**(1), Article 7.

Resnick, P., R. Zeckhauser, J. Swanson, K. Lockwood. 2006. The value of reputation on eBay: A controlled experiment. *Experiment. Econom.* **9**(2) 79–101.

Roth, A. E., A. Ockenfels. 2002. Last-minute bidding and rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *Amer. Econom. Rev.* **92**(4) 1093–1103.

Simonsohn, U., D. Ariely. 2008. When rational sellers face nonrational consumers: Evidence from herding on eBay. *Management Sci.* **54**(9) 1624–1637.

Song, U. 2004. Nonparametric estimation of an eBay auction model with an unknown number of bidders. Working paper, University of British Columbia, Vancouver.

Wilcox, R. T. 2000. Experts and amateurs: The role of experience in Internet auctions. *Marketing Lett.* **11**(4) 363–374.

Yao, S., C. F. Mela. 2008. Online auction demand. *Marketing Sci.* **27**(5) 861–885.

Yin, P.-L. 2006. Information dispersion and auction prices. Working Paper 06–011, Harvard Business School, Cambridge, MA.

Zeithammer, R. 2006. Forward-looking bidding in online auctions. *J. Marketing Res.* **43**(3) 462–476.