

Commentary

Do Bids Equal Values on eBay?

Ali Hortaçsu

Department of Economics, University of Chicago, Chicago, Illinois 60637; and National Bureau of Economic Research, Cambridge, Massachusetts 02138, hortacsu@uchicago.edu

Eric R. Nielsen

Department of Economics, University of Chicago, Chicago, Illinois 60637, nielsen@uchicago.edu

We argue that the Adams and Zeithammer paper [Zeithammer, R., C. Adams. 2010. The sealed-bid abstraction in online auctions. *Marketing Sci.*, ePub ahead of print August 11, <http://mktsci.journal.informs.org/cgi/content/abstract/mksc.1100.0561v1>] successfully documents consistent patterns in eBay bidding data that cast doubt on the common assumption that bidders in such auctions follow a “bid = value” strategy. These anomalies lend support to the authors’ alternative model in which some bidders bid reactively and consequently bid below their valuation most of the time. The consistency of the authors’ findings as well as the ability of their alternative explanation to account for all of their test results lends great support to their thesis. However, we think that several of their empirical tests examine ancillary assumptions about bidder behavior and do not test the bid = value assumption directly. Furthermore, although their reduced-form model incorporating “reactive” bidders is a good first attempt at expanding the canonical framework, we worry that their counterfactual pricing analysis using the reactive model is suspect because the parameters they estimate are not structural. Overall, the Zeithammer and Adams paper is a carefully argued critique of empirical methods used to study online auctions and provides valuable ideas to improve on these methods.

Key words: eBay; auction demand estimation; iid bidder valuation; optimal reserve price; sealed-bid assumption

History: Received: May 8, 2010; accepted: May 8, 2010. Published online in *Articles in Advance*.

Online auctions have become a ubiquitous part of everyday life and understanding underlying bidder demand, as revealed through bids, has proved to be an important growth area for marketing and economic research. The unifying idea of this research program is to “invert” observed bids to infer the unobserved distribution of bidder valuations and to use the inferred bidder valuations to conduct counterfactual pricing/policy exercises. A successful real-world application of this approach with large revenue consequences has recently been performed in the sponsored search setting by Ostrovsky and Schwarz (2009).

Of course, this inversion exercise depends crucially on identifying the correct mapping between valuations and bids. Unfortunately, online auctions typically utilize an open ascending format (but do not follow the “button” auction abstraction of Milgrom and Weber 1982). Moreover, the presence of concurrent competing auctions of often identical objects complicates the characterization of bidder behavior. Furthermore, assuming that players are following Nash equilibrium strategies might not be very

compelling in this setting because the bidders typically are very heterogeneous in terms of their sophistication and experience with the auction format.

Against this theoretical background, early economic work in online auctions has largely abstracted away from the complications generated by the extensive form of the bidding game. Within the private values framework (which, with a few exceptions, is standard in the online auction literature), researchers have relied on the behavioral assumption that the final observed bid of each bidder is equal to the bidder’s valuation. All players bidding their private values constitutes an equilibrium in weakly dominant strategies in a sealed-bid second-price auction, but it could also be supported as equilibrium strategies in a number of nonsealed-bid settings as well. Zeithammer and Adams (2010), in fact, is a critique of an even weaker behavioral assumption: that the top two bids in an auction reflect bidders’ true valuations.

As Zeithammer and Adams point out, this behavioral assumption is not without its justifications, though, undeniably, one of its main attractions is the simplicity it affords to the “inversion” strategy

mentioned above: if bids equal values, then the distribution of bids identifies the distribution of bidder valuations directly. However, if only the top bids are to be used in estimation, one would have to account for the fact that these bids are order statistics. An important challenge in the early literature was the lack of direct measures of the number of (potential) bidders in the auction, which made the inversion of order statistic distributions difficult. An important advance made by Song (2004) appeared to address this issue by a “nonlinear fixed effect”-type approach; Song showed that observing any two order statistics from the bid/value distribution is enough to “difference out” the unknown number of bidders in the auction, allowing for the identification of the underlying distribution.

Song’s approach, although still regrettably unpublished, has remained the standard to beat in the empirical online auctions literature for quite some time. The significance of Zeithammer and Adams (2010) is to assail the behavioral assumption motivating Song’s method: that the (top) two observed bids indeed equal bidders’ valuations. Their paper thus proceeds in two parts: first, the authors implement a battery of empirical tests to assess the validity of the “bid = value” assumption for the top two bids. Second, they propose a model that departs from this behavioral assumption in a way that is in line with the results of the tests in the first part, and they compare the results of a structural estimation exercise that nests a model akin to that of Song.

Before we delve into details, we would like to commend Zeithammer and Adams for a very well thought-out and carefully executed paper that left few stones unturned. The “testing” part of their paper in fact proposes seven different tests of the bid = value assumption. Although some of these tests are somewhat less directly connected to the behavioral bid = value assumption than to other ancillary assumptions used in the literature (such as the assumption that bidder valuations are iid), we think, overall, that the authors present clear evidence that there are patterns in the bidding data that are difficult to reconcile with the bid = value assumption. Furthermore, the behavioral assumption that the authors propose to extend the bid = value paradigm—namely, that there are some “reactive” bidders who bid less than their value and return to the auction periodically and “top” the extant highest bid should it be less than their value—does an excellent job of explaining all of the empirical anomalies the testing portion of their paper documents. Although one could try to posit ad hoc explanations of their test results on a case-by-case basis, the fact that one intuitive behavioral modification can, in principle, reproduce all of the anomalies documented lends great support to their thesis.

As practitioners of the dismal science, we are inclined to accentuate the negative; thus we start with the least convincing of the tests. These are the tests based on the timing of bids. For example, the null hypothesis of the first test, T1, is that the highest bid arrives before the second-highest bid 50% of the time. This, however, is a test of the assumption of independence between bid timing and valuations (A1) rather than a test of whether bidders bid their valuations, and these two conditions need not coincide. Indeed, consider the following scenario: suppose that auctions are listed so that they end at evening hours, and bidders who place their bids during the daytime are “traders,” who bid on items as part of their business, and bidders who place their bids in the evenings are “hobbyists,” who can only bid when they are at home. If the hobbyists, on average, have higher valuation for the objects than the traders, we may get the result that the highest bid is placed later than the second-highest bid even if both groups always follow sealed-bid strategies and bid their private valuations (which is what their Table 2 suggests). Of course, this is just an illustrative example, but the point is that there could many plausible mechanisms generating the correlation between the timing of bids and bidder values, and this correlation may not necessarily be due to departures from sealed-bid strategies. Note that the correlation between bidder valuations and bid timings would also be problematic for tests T3 and T4.

A more convincing test that does not rely on the independence between bidder valuations and timing is T5: the distribution of the highest bid given the second-highest bid should have the same right tail for any realization of the second-highest bid. This criterion must be satisfied if bidders are playing the *symmetric* equilibrium of a second-price sealed-bid auction. This criterion, however, may be violated if one allows for (ex ante) asymmetries in bidder valuations. For example, suppose there are two (types of) bidders in the auction, H and T , and the bidders’ valuations are drawn from two-point distributions with mass 0.5 on each point: $v_H \in \{2, 4\}$ and $v_T \in \{1, 3\}$. The bidders play sealed-bid strategies, i.e., bid their valuations. Then conditional on the second-highest bid being 3, the distribution of the highest bid will assign a probability of 1 to 4, but if the second-highest bid is 2, the highest bid will be 3 with a probability of 1. Extending this example to cases with more than two bidders per auction yields similar violations of T5.

The test T5 constitutes a nice contribution along another dimension as well. Zeithammer and Adams (2010) propose a method of implementing the test that appears to be nonparametric. Roughly speaking, they examine pairs of auctions such that the smaller of the two highest bids exceeds the higher of the

two second-place bids. Among such auctions, the auction with the higher price should have the higher top bid half the time. Given a large enough sample of homogeneous auctions, one simply needs to count the number of times this event occurs among all qualifying pairs of auctions and compare the resulting sample fraction to 0.5. Interestingly, the authors find that the distribution of this test statistic does not seem to depend on the valuation distribution. Although they do not attempt to verify the conjecture, it is quite intriguing because it would provide a straightforward test of whether order statistic data are drawn from iid samples. In their empirical work, sample size and concerns about unobserved heterogeneity force them to employ a semiparametric approach in which auction-level observables enter linearly in the logs of the bids. Even in this specification the test has a lot of intuitive appeal and remains easy to implement.

We should note that the examples we used above all invoke some departure from the symmetry/iid assumptions that are standard in the literature and, quite reasonably, are also taken as a given by the authors. However, the theory proposed by Zeithammer and Adams (2010) in §5 to explain their test results is one that allows for *behavioral* asymmetries across bidders: although bidders' valuations are iid, some bid their values, but others use reactive strategies. This is an interesting and, in light of the data, reasonable route toward relaxing the standard framework. However, it may strike some that a similarly fruitful route to take may have been to leave the bid = value assumption intact but relax the iid assumption, in particular, allowing for asymmetries across bidders' valuations.

Fortunately, the results of another test conducted by the authors give further credence toward the presence of behavioral asymmetries. This test, T2, checks whether a significant fraction of first and second bids are separated by exactly one bidding increment. The answer is a resounding yes, which is a violation of the assumption that the bid distribution (symmetric or asymmetric) is continuous. Of course, value distributions do not have to be continuous, but the authors are very careful to check whether this finding may be due to the bid increment being in terms of dollar multiples or due to other explanations reflecting the possibility that bidders' private valuations of objects indeed do contain masses at "round" numbers. One additional possibility that is left unexplored is whether shill bidding by sellers may be the cause of such patterns; on eBay, the seller (or an accomplice) may submit incremental bids that try to extract as much surplus as possible from the buyer. This possibility is examined at some length by Engelberg and Williams (2009), who report a number of empirical patterns consistent with shill bidding on eBay.

The model proposed in the Zeithammer and Adams paper has, once again, two types of bidders, each with iid valuations. A fraction of the bidders, called "sealed" bidders, bid their value, while the others, called reactive bidders, start by bidding a fraction of their value, matching competing bids until they exceed their valuation. The main econometric implication of this model is that if a reactive bidder wins the auction, it is with a bid that is below her valuation. Thus, an important goal of the structural estimation exercise is to recover by how much a reactive bidder's winning bid is exceeded by her actual valuation. The authors approach this difficult problem by estimating the average "bias" in the winning bid of a reactive bidder, which they find to be 47%. This surprisingly large difference in bids and valuations for (winning) reactive bidders also leads to a strong bottom-line conclusion for their paper: not accounting for reactive bidders can lead to very significant biases in the estimation of valuations, as shown in their Figure 2.

Although the identification of the crucial bias parameter is reasonably clear from the intuition provided by Zeithammer and Adams (compare the winning bids across auctions won by reactive and sealed bidders, conditioning on the second-highest bid, and the type of the second-highest bidder), we wonder if they could have performed a few other robustness checks around this main empirical exercise. For example, using the sample of auctions for which the top two bidders have been classified in the first stage as being sealed bidders with very high probability, the result of the Song (2004) specification should be identical to the authors' preferred results. A somewhat more nitpicky robustness check is to investigate the sensitivity of the results to the log-normality specification, as the auction literature at large is typically quite careful about parametric distributional assumptions. Another nitpicky objection is that there is a tension between the language used during the "testing" phase of their paper to describe the reactive bidder model and the empirical model with reactive bidders actually estimated. An auction won by a reactive bidder should always have the top two bids separated by exactly one bid increment. Because some bidders can be identified as reactive because they bid multiple times in one auction, in principle one could test this prediction of the "pure" reactive model in auctions won by a bidder revealed as reactive. Similarly, their model should generate predictions regarding the values of the other test statistics in simulated auctions, and it would be interesting to see how well their model can match these other values.

The authors are careful to qualify their estimate of the bias parameter as being reduced form, as the

behavior of reactive bidders likely depends on a number of factors (such as the importance of competing auctions in the particular market under consideration). Unfortunately, we do not have satisfactory theoretical explanations for this kind of behavior; thus the estimation exercise should not be extrapolated too far. In particular, counterfactual scenarios in which the behavior of reactive bidders may change are difficult to evaluate with the current model. Although we are convinced that a substantial number of online auctions contains bidders following nonsealed-bid strategies, we are not sure that the counterfactual reserve pricing scenarios are sensible because the bias parameter is not structural and may well depend on the reserve price.

We once again commend Zeithammer and Adams for their thoughtful and thought-provoking work, and we are optimistic that their model will lead to future research understanding the causes of non-“sealed-bid” behavior in online auctions and how such models can be incorporated into existing structural econometric models. Their paper provides a wealth of empirical regularities that should prompt

further research on more complex models of bidder behavior. It is clear from the present work that such an undertaking is both needed and consequential. In the meantime, the model proposed in their paper should be taken as a benchmark in future empirical studies of (private value) online auctions.

References

- Engelberg, J., J. Williams. 2009. eBay's proxy bidding: A license to shill. *J. Econom. Behav. Organ.* **72**(1) 509–526.
- Milgrom, P. R., R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50**(5) 1089–1122.
- Ostrovsky, M., M. Schwarz. 2009. Reserve prices in Internet advertising auctions: A field experiment. Stanford Graduate School of Business Research Paper 2054, Stanford University, Stanford, CA.
- Song, U. 2004. Nonparametric estimation of an eBay auction model with an unknown number of bidders. Working paper, University of British Columbia, Vancouver.
- Zeithammer, R., C. Adams. The sealed-bid abstraction in online auctions. *Marketing Sci.*, ePub ahead of print August 11, <http://mktsci.journal.informs.org/cgi/content/abstract/mksc.1100.0561v1>.