

Advances in Economic Analysis & Policy

Volume 6, Issue 2

2006

Article 3

FIELD EXPERIMENTS

...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay

Tanjim Hossain*

John Morgan†

*Hong Kong University of Science & Technology, tanjim@ust.hk

†University of California, Berkeley, morgan@haas.berkeley.edu

Copyright ©2006 The Berkeley Electronic Press. All rights reserved.

...Plus Shipping and Handling: Revenue (Non) Equivalence in Field Experiments on eBay*

Tanjim Hossain and John Morgan

Abstract

Many firms divide the price a consumer pays for a good into two pieces—the price for the item itself and the price for shipping and handling. With fully rational customers, the exact division between the two prices is irrelevant—only the total price matters. We test this hypothesis by selling matched pairs of CDs and Xbox games in a series of field experiments on eBay. In theory, the ending auction price should vary inversely with the shipping charge to leave the total price paid constant. Contrary to the theory, we find that charging a high shipping cost and starting the auction at a low opening price leads to higher numbers of bidders and higher revenues when the shipping charge is not excessive. We show that these results can be accounted for by boundedly rational bidding behavior such as loss-aversion with separate mental accounts for different attributes of the price or disregard for shipping costs.

KEYWORDS: internet auctions, e-retailing, field experiments, price framing, mental accounting, loss aversion

*We thank two anonymous referees, Mike Baye, Raj Chetty, Richard Engelbrecht-Wiggans, Stefano DellaVigna, Antonio Rangel, Jason Shachat, Ken Steiglitz, Steve Tadelis, Jeremy Tobacman and seminar participants at Harvard, SITE conference, UC Berkeley, UC Santa Barbara and Washington University in St. Louis for helpful comments. The second author also thanks Charles Thibaux for his able research assistance in conducting earlier pilot experiments. All remaining errors are our own. The first author gratefully acknowledges the financial support of the Woodrow Wilson Foundation. The second author gratefully acknowledges the financial support of the National Science Foundation. This paper was previously circulated under the title, “A Test of the Revenue Equivalence Theorem Using Field Experiments on eBay.”

Erratum

The equation for bidder i 's payment if she wins on page 5 should read $P_i(r,o,s) = \max_{j \neq i} \beta_j(r,o,s) + s$. Consequently, the subscript on β and γ in the proof of Proposition 1 should be j instead of i . These are typographical errors, and none of the results and conclusions of the paper are affected.

A common marketing practice, particularly in e-retail, is to divide the price a consumer pays for a good into two pieces—the price for the item itself and the price for shipping and handling. A similar practice has been prevalent in certain sectors of the offline retail sector as well. Many items advertised on television and “not sold in stores,” such as Ginsu knives as well as other products sold by the companies K-tel and Ronco, highlight the price for the item itself and then (typically in smaller type) add on a shipping and handling charge. A similar practice occurs routinely in restaurants. In the US, it is common to automatically add an 18% gratuity to the list price for parties of eight or more. Outside the US, such a gratuity (although at a lower percentage) is often automatically added to bills for parties of any size. In theory, the practice of dividing a price into these two pieces should have little effect on overall demand for a good. A perfectly informed and fully rational consumer will merely add together the two parts of a price to obtain the total out of pocket price for an item and then determine whether or not to buy based on this total price.

While for some cases there may be an important cost reason to keep the “list” price and the shipping and handling charge separate. For instance, a customer buying a consumer electronics product from an e-retailer will often be given a choice in the speed in which an item is shipped. Since faster shipping does indeed cost the e-retailer more to provide, it makes sense that the total price is divided in this fashion. Similarly, in book retailing online, the cost to the company of shipping the second and additional items to the same address is typically lower than the cost to ship the first item and the shipping and handling schedule will often reflect this nonlinearity.

In other instances, however, there is no cost-based reason for the division of the price into two parts. For example, it is clear that if there is only one way in which the item is to be shipped or one cannot really even separate the item (say, an item from a menu) from the service (actually serving the item on the menu to the customer at the restaurant), the two part price is not selected for cost-based reasons.

It is, however, anecdotally suggested that dividing the price into two parts leads to an increase in demand. The informal argument suggests that consumers pay attention to the list price while neglecting the “add-on” price associated with shipping or with the gratuity. Therefore, they systematically underestimate the total out of pocket price paid for the items being purchased or consumed and hence the seller is able to sell more of a given item than had that firm simply listed a single total price. Mental accounting (Kahneman and Tversky, 1984 and Thaler, 1985) offers one way to formalize this intuition. Here, the idea is that consumers may keep separate mental “accounts” for each of the components of the provision of an item, such as an account for the item itself and a separate account for the shipping of the item. Under this framework, demand under the

“total price” need not be the same as the demand when the price is decomposed into prices for the various accounts. To estimate how various price “frames” affect demand in a posted price market setting using field data is difficult. One would need a significantly large data set on demand resulting from different price frames. That is, since the two prices making up the frame are fixed, one would have to use variation in quantity demanded to test for an effect. The online auction site, eBay, however, offers an attractive alternative for testing whether variation in the price frame matters. In a typical eBay auction, the quantity available for sale as well as the shipping and handling charge are fixed while the market clearing price for the item (exclusive of shipping) is determined endogenously through eBay’s auction process. By relying on this price variation rather than quantity variation, one is able to much more readily determine the effect of changes in the price frame.

Exploiting this observation, we conducted 80 auctions on eBay. Forty of these auctions were for various popular music CDs while the remaining 40 auctions were for video games for Microsoft’s Xbox gaming console. The items auctioned were chosen to be “simple” in the sense that all of these items are readily available at retail stores; thus differential expertise or information about the quality of a given item among bidders would seem to play little role. Further, since these items tend to be of short-lived popularity, a bidder’s estimate of resale value is unlikely to be an important consideration either. We auctioned all of these items using the standard (in many cases the default) rules governing auctions on eBay—again in hopes of isolating the effect of changing the price frame from other effects that might be present in more complicated settings such as auctions of unique items where expertise might be important.

Auctions on eBay are, in effect, English auctions run over seven day periods.¹ Hence, if a bidder knows his or her value for an item, her optimal bidding strategy is simply to bid up to maximum willingness to pay in that auction as her final bid.² When a bidder has outside options (future auctions or a retail market), a bidder should bid up to the point where she is just indifferent between winning the auction at her bid amount and the value of the outside option. In determining whether to bid at all, a bidder must determine whether her willingness to pay exceeds the value of the “effective reserve price” in an auction. The effective reserve price is the minimum amount any bidder has to pay in the event that no other bidders bid in an auction. For eBay, the effective or total reserve price is simply the sum of the opening bid amount and the shipping and handling charge. A key implication of auction theory is strategic equivalence.

¹ The interested reader should see Roth and Ockenfels (2002) for additional details about the specifics of bidding on eBay.

² For simplicity, we will use valuation and willingness to pay interchangeably.

Two eBay auctions with the same effective reserve price are strategically equivalent in the sense that a bidder's effective maximum bid (her bid on eBay plus the shipping charge) should be invariant to changes in the composition of the effective reserve price. An example helps to fix ideas. Suppose that a bidder had a maximum willingness to pay of \$40 for the delivery of an Xbox game to her home via first-class mail. Suppose that shipping and handling were free for the item. Then the bidder should be willing to bid up to her true willingness to pay, \$40. If, however, the auction had a shipping and handling charge of \$5 for the item, then the bidder should bid only up to \$35 in the auction. Notice that her effective bid—the amount that she would pay to the seller in the event that she won the auction at her amount bid—is exactly the same, \$40, in the two auctions. Thus, the two auctions are strategically equivalent. Moreover, strategic equivalence in turn implies that the expected revenues to the seller will be the same under variations in the shipping charge keeping the effective reserve level fixed—that is, these auctions are also predicted to be revenue equivalent as well. Since the revenue equivalence theorem is widely considered to be the cornerstone of auction theory, our experiments should also be of interest in that literature.³

In our experiments, we auctioned 4 copies each of 10 popular music CDs and 10 Xbox games. Half of these auctions were run with a \$4 effective reserve and the other half with an \$8 effective reserve. For a given reserve amount, we auctioned off matched pairs of CDs where we varied the level of the opening bid and the shipping cost by \$4 while holding fixed the overall reserve level.

Our main findings are as follows: When the effective reserve was \$4, auctions with a low (\$0.01) opening bid and high (\$3.99) shipping charges attracted more bidders, earlier bidding, and yielded higher revenue for both CDs and Xbox games. With \$8 effective reserve, the low opening bid was (\$2) and high shipping charge (\$6) generated higher revenue for Xbox games. In these cases, charging a significant portion of the price as shipping increased revenue. For CDs, the \$8 effective reserve represents over 50% of the retail price of the item. Here, we find no systematic difference in the number of bidders attracted to the auction or revenues as a function of how the effective reserve is allocated between opening bid and shipping charges. An institutional detail of eBay may account for this non-result: A shipping charge of \$6 is uncommon in eBay auctions of music CDs while is not uncommon for auctions of Xbox games. For

³ The main interest in the revenue equivalence theorem centers around differences in the auction form (i.e., first-price versus second-price), which we are not testing. Nonetheless, for the conditions of our experiments, revenue equivalence is predicted to hold regardless of whether bidders have private values or know how many other bidders are competing, are symmetric, and so on. We tried, however, to choose conditions approximating the usual statement of the theorem (see, for example, Proposition 3.1 in Krishna, 2002) to offer the theory its best chance of succeeding.

all other treatments in the experiment, the shipping prices charged were well within the norms observed on eBay.

The nearest antecedent to our paper is Katkar and Reiley (2005), who conducted auctions for Pokémon trading cards on eBay to test the equivalence between public and secret reserve prices. List and Lucking-Reiley (2000), who examine whether the auction theory prediction of demand reduction in multiple unit uniform price auctions is observed in the data, also use field experiments. Lucking-Reiley (1999) and Reiley (2005) conducted field experiments examining the implications of varying auction form and reserve prices using collectible trading cards on Usenet.

The literature on auction laboratory experiments is vast (see, for instance Davis and Holt, 1993; and Kagel, 1995 for surveys of the literature). The main focus on this literature is to study the effects of variation on auction form (i.e. first-price versus second-price, static versus dynamic) as well as bidder information structure (i.e., independent versus correlated signals about value) on auction outcomes. Our main finding that a seller usually earns greater revenue by setting a lower opening bid and a higher shipping charge is similar to the findings of Morwitz, Greenleaf, and Johnson (1998). They show that, relative to a standard first price auction, bidders effectively bid more aggressively when the winner of the auction has to pay 115% of her bid. Unlike in our field experiments where the total cost is the price from the auction plus shipping fee, the total cost to a winner cannot be so easily calculated in that experiment.

Within the empirical literature on e-retailing, our work is related to Ellison and Ellison (2004), who show that “obfuscation” in computer memory is a common (and apparently profitable) retailing strategy. Smith and Brynjolfsson (2001) find that in online book retailing consumers are more sensitive to variation in shipping charges than to variation in price. Again, this is contrary to the notion that only the total price matters in determining demand but in a direction opposite to our findings. There are several key differences between Smith and Brynjolfsson and our work. First, Smith and Brynjolfsson can only observe consumers’ click behavior (last clicks more precisely) rather than actual purchases. Second, as we mentioned above, there are economic reasons for variation in shipping charges in book retailing as well as nonlinearity in shipping schedules that make these prices arguably more salient than in our setting.

The remainder of the paper proceeds as follows: In section 2, we sketch the theory leading to strategic equivalence. In section 3, we detail the experimental methodology. In section 4, we present the results of the experiments. In section 5, we consider a number of possible explanations for our findings. We conclude the paper in section 6.

1. Theory

Consider a second-price sealed bid auction for some object, which one can think of as a stylized version of an eBay auction.⁴ In this auction, a stochastic number of bidders, $i = 1, 2, \dots, N$ submit bids b_i for an object. The high bidder pays the second-highest bid amount plus a shipping charge, s . Further, there is a minimum opening bid, o . A bidder must bid at least o to participate in the auction. Further, if only a single bidder bids, then that bidder pays the amount of the effective reserve price, $r = o + s$. Suppose that for an auction with reserve price r , opening bid o and shipping charge s , Bidder i makes equilibrium bid $\beta_i(r, o, s)$. If that bidder wins, then her payment is:

$$P_i(r, o, s) = \max_{j \neq i} \beta_j(r, o, s) + s.$$

Next, consider a variation where the effective reserve price remains fixed but o and s are varied to o' and s' respectively. That is, there exists some $\varepsilon \neq 0$ such that $o' = o + \varepsilon$ and $s' = s - \varepsilon$.

Proposition 1 *Suppose $\beta_i(r, o, s)$ is an equilibrium bidding profile in a second-price auction under (r, o, s) . Then, $\gamma_i(r, o, s) = \beta_i(r, o, s) + \varepsilon$ is an equilibrium in a second-price auction under (r, o', s') .*

Proof. Notice that the conditions in which each bidder participates are identical under (r, o, s) as under (r, o', s') . Further, under the γ bidding strategies, the conditions in which bidder i wins the auction are identical under (r, o, s) as under (r, o', s') . Finally, notice that if bidder i wins, her expected payment is

$$\begin{aligned} P_i(r, o', s') &= \max_{j \neq i} \gamma_j(r, o', s') + s' \\ &= \max_{j \neq i} (\beta_j(r, o, s) + \varepsilon) + s - \varepsilon \\ &= \max_{j \neq i} \beta_j(r, o, s) + s \\ &= P_i(r, o, s), \end{aligned}$$

which is identical under (r, o, s) as under (r, o', s') . Therefore, if the β strategies comprise an equilibrium under (r, o, s) then the γ strategies comprise an equilibrium under (r, o', s') . ■

⁴ The basic intuition for the main result in this section can be readily extended to other possible models of eBay such as English auctions or dynamic second-price auctions.

In words, Proposition 1 shows that variation in the composition of the effective reserve price between the opening bid and the shipping charge leads to strategically equivalent bidding strategies. This immediately implies:

Corollary 1 *Fix an effective reserve level r . Then all eBay auctions under $\$r$ are revenue equivalent.*

Finally, since we vary the effective reserve level in the experiments, we restate the following well-known result from auction theory (See Krishna (2002)):

Proposition 2 *Consider an eBay auction where bidders have private values. Raising the effective reserve price (regardless of its composition): reduces the number of expected bidders, decreases the probability of a sale, and increases expected revenues conditional on a sale being made.*

2. Procedures

We wanted to auction goods where multiple units of the same good are identical and where markets are thick enough so that our auctions were unlikely to have a marked effect on market prices for these goods overall. First, we chose personal entertainment goods: music CDs and Xbox games as the categories in which we want to participate. Then, we did a survey of eBay sales in each of these markets to select items that are currently popular within each category. By choosing relatively popular CDs, we would expect smaller variance in the sale price from random fluctuations in demand for the good. Further, we hoped that the thickness of the market would disguise the fact that we are running a field experiment. This alleviates the worry that bidders might behave differently when they are aware that they are participating in an experiment.

A number of studies have found the existence of reputation effects on sales price for eBay auctions (see, for instance, Resnick and Zeckhauser, 2002). To control for this effect, we created unique seller accounts for each of the goods we sold. Each of these seller accounts had a zero feedback rating. Thus, the sellers were all identical from a bidder's perspective.

For each type of good (music CDs and Xbox games), we ran four treatments. The first two treatments were constructed as follows: In treatment A, we set an opening bid of \$4 and set the shipping and handling cost of the item at \$0. In treatment B, we set an opening bid of \$0.01 and a shipping cost of \$3.99. We will refer to treatments A and B collectively as "low reserve price" treatments. The second set of two treatments were as follows: In treatment C, the opening bid was set at \$6 and the shipping cost at \$2. Finally, treatment D set the opening bid at \$2 and the shipping and handling cost at \$6. We will refer to

treatments C and D collectively as “high reserve price” treatments. For bidders, treatments A and B are strategically equivalent and treatments C and D are strategically equivalent.

For all treatments, we used an identical description of the good and stipulated the exact method of shipping (USPS first-class). The exact wording used in each of the auctions is given in the Appendix. To control for the possibility that local bidders might assume that shipping and handling charges might be avoided by picking up the item in person, we responded to all bidder queries along these lines that the shipping and handling charge was non-negotiable and that in-person pickup was not possible. Upon receipt of payment, items were shipped to the winning bidder exactly as described.

Our first round of experiments consisted of auctioning four copies of each of the 10 music CDs. Each of these CDs was purchased new (and sealed) from Amazon.com. In every case, we used the standard seven day closing rules on eBay, varying only the opening bid and the shipping and handling amounts as described above. We ran the low reserve price treatments for auctions scheduled to end on 11/19/2001 and 11/20/2001 (a Monday and a Tuesday, respectively). For the auctions scheduled to end on Monday, we randomly chose whether to run treatment A or B with equal probability. The following day, we posted auctions for the same 10 CDs using the opposite treatment. Thus, for each CD, treatments A and B occur (in random order) on consecutive days. The reason for adopting this procedure is that we would expect there to be little difference in demand for auctions for a given CD that are 24 hours apart. Further, if there is a systematic Monday versus Tuesday effect on revenues, our randomization should avoid confounding this with a treatment effect. The following week, we performed an identical procedure for the high reserve price treatments. All of the auctions ended between eight and ten p.m. eastern standard time.

Our second round of experiments, which consisted of auctioning four copies of each of 10 Xbox games, occurred in March 2002. Each of these games was purchased new and sealed from Amazon. We chose the timing of this round of experiments to avoid the Christmas holiday, when we expected demand might be different from other times of the year. We followed exactly the same procedures as with the music CDs. The next section discusses the results from these 80 auctions.

3. Results

3.1 Overview

Tables 1 and 2 summarize the results from music CD and Xbox games auctions respectively. The column “Bids” presents the total number of bids and “Bidders”

is the number of distinct bidders who placed at least a bid in that auction. For the low reserve price treatments, all but one of the CDs auctioned were sold; whereas under the high reserve price treatments, 5 out of the 20 CDs went unsold. This is broadly consistent with the prediction that higher reserve prices are more likely to lead to failures to sell. Conditional on being sold, the average final sale price (including shipping) for the low reserve price treatments was \$9.30. Under the high reserve price treatments, the final price conditional on a sale was \$12.21. This too is consistent with theoretical prediction that prices conditional on a sale occurring are higher under a higher reserve price. Finally, sale prices on eBay tend to be well below the retail price. The retail prices for each of the CDs represent the cost to us from purchasing from Amazon, allocating the total shipping cost equally over all 40 music CDs. Including shipping costs, the average price we paid was \$14.82 per CD; thus, winning bidders in our auctions received “bargains” compared to buying at Amazon. This can be explained by the fact that many consumers with relatively lower valuations shop on eBay to find great deals.

In all treatments of Xbox games auctions, every Xbox game was sold as Table 2 shows. The average final sale price (including shipping fees) for the low reserve price treatments was \$37.47. Under the high reserve price treatments, the final price conditional on a sale was \$39.01. This is consistent with the theoretical prediction that prices conditional on a sale occurring are higher under a higher reserve price. This revenue improvement prediction, however, assumes that at least the higher reserve price is inside the support of the bidders’ private valuations. If both reserve prices are below the lowest possible valuation, they will generate the same expected revenue conditional on a sale. For Xbox games, even \$8 is significantly lower than the pre-tax retail price of \$49.99. Thus, a very small fraction of bidders may have valuation below either of the reserve prices. Not surprisingly, although the mean revenue is higher for \$8 reserve, this difference is not statistically significant. As with the music CDs, sale prices on eBay tend to be below the retail price. Including shipping charges, the average price we paid was \$51.07 per Xbox game; thus, winning bidders in our auctions received “bargains” paying about 72% of the cost at Amazon.

Table 1. Overview of Music CD Auctions

Low Reserve Treatments:

Total Reserve = \$4 CD	Treatment A			Treatment B		
	Revenue	Bids	Bidders	Revenue	Bids	Bidders
Music	5.50	4	2	7.24	6	4
Ooops! I Did it Again	6.50	3	3	7.74	10	4
Serendipity	8.50	5	4	10.49	8	4
O Brother Where Art Thou?	12.50	7	7	11.99	7	4
Greatest Hits - Tim McGraw	11.00	11	8	15.99	12	8
A Day Without Rain	13.50	7	6	14.99	9	6
Automatic for the People	0.00	0	0	9.99	5	3
Everyday	7.28	3	3	9.49	9	7
Joshua Tree	6.07	3	3	8.25	6	3
Unplugged in New York	4.50	3	3	5.24	5	2
<i>Average</i>	<i>7.54</i>	<i>4.6</i>	<i>3.9</i>	<i>10.14</i>	<i>7.7</i>	<i>4.5</i>

High Reserve Treatments:

Total Reserve = \$8 CD	Treatment C			Treatment D		
	Revenue	Bids	Bidders	Revenue	Bids	Bidders
Music	9.00	3	3	8.00	1	1
Ooops! I Did it Again	0.00	0	0	0.00	0	0
Serendipity	12.50	5	4	13.50	5	4
O Brother Where Art Thou?	11.52	5	5	11.00	7	5
Greatest Hits - Tim McGraw	18.00	9	3	17.00	15	7
A Day Without Rain	15.50	11	7	16.00	10	5
Automatic for the People	0.00	0	0	0.00	0	0
Everyday	10.50	6	4	13.50	6	4
Joshua Tree	8.00	1	1	11.10	7	4
Unplugged in New York	8.00	1	1	0.00	0	0
<i>Average</i>	<i>9.30</i>	<i>4.1</i>	<i>2.8</i>	<i>9.01</i>	<i>5.1</i>	<i>3</i>

Table 2. Overview of Xbox Auctions**Low Reserve****Treatments:**

Total Reserve = \$4	Treatment A			Treatment B		
Game	Revenues	Bids	Bidders	Revenues	Bids	Bidders
Halo	34.05	13	9	41.24	11	9
Wreckless	44.01	7	7	33.99	12	7
Circus Maximus	40.99	8	7	39.99	10	6
Max Payne	36.01	24	10	36.99	16	10
Genma Onimusha	41.00	17	8	32.99	14	11
Project Gotham Racing	37.00	13	6	38.12	11	8
NBA 2K2	42.12	20	10	42.99	15	9
NFL 2K2	26.00	14	5	33.99	9	9
NHL 2002	36.00	10	8	37.00	8	8
WWF Raw	33.99	13	8	40.99	21	11
<i>Average</i>	<i>37.12</i>	<i>13.9</i>	<i>7.8</i>	<i>37.83</i>	<i>12.7</i>	<i>8.8</i>

High Reserve**Treatments:**

Total Reserve = \$8	Treatment C			Treatment D		
Game	Revenues	Bids	Bidders	Revenues	Bids	Bidders
Halo	40.01	10	8	43.00	14	10
Wreckless	35.00	11	6	36.00	10	7
Circus Maximus	39.00	8	7	42.53	13	10
Max Payne	37.50	27	8	42.00	24	11
Genma Onimusha	36.00	18	11	37.00	13	8
Project Gotham Racing	35.02	13	7	40.01	12	10
NBA 2K2	41.00	26	7	45.00	15	10
NFL 2K2	33.00	15	8	40.10	21	10
NHL 2002	36.00	16	12	41.00	18	14
WWF Raw	37.00	17	10	44.00	27	11
<i>Average</i>	<i>36.95</i>	<i>16.1</i>	<i>8.4</i>	<i>41.06</i>	<i>16.7</i>	<i>10.1</i>

Finally, we check for the presence of “day of the week” effects in the data. Under the null hypothesis that there are no systematic differences between Mondays and Tuesdays, the revenue ranking for each pair of CDs or games should be the outcome of a fair coin flip. We find that for 20 out of 38 matched pairs of auctions, higher revenue was obtained on Tuesdays compared to Mondays; thus we do not find evidence of a systematic day of the week effect.⁵

We now look more closely into revenue ranking and effect of different treatments on the number of bidders.

3.2 Revenues

Table 3 compares revenues under the low reserve treatments. Notice for music CDs, the average revenue under treatment B is \$10.14 compared to only \$7.54 in treatment A. Of course, part of this difference is accounted for by the fact that one of the CDs under treatment A did not sell. If we exclude this CD from the averages, the average under treatment B still exceeds that under treatment A by \$1.79, or about 21% of the price CDs sold for under treatment A. For Xbox games, average revenues under treatment B still exceed those under treatment A, but by a smaller margin, only \$0.71 (or 2% of the revenues under treatment A). Taken together, however, this suggests that there might be systematic revenue differences between the two treatments.

More formally, we can test the null hypothesis of revenue equality against the one-sided alternative that B outperforms A using a binomial test. This essentially involves a count of the number of auctions in which treatment B yielded higher revenues than did treatment A. In 9 out of 10 matched pairs of auctions for music CDs, this is the case. In 7 out of 10 matched pairs of auctions for Xbox games treatment B outperforms A. Thus, for 16 out of 20 matched pairs of auctions with low reserve prices, we find that treatment B outperforms treatment A. Using a one-sided binomial test, then we get a p -value of 0.005. Alternatively, one could use a Wilcoxon signed-ranks test to test the null hypothesis of no treatment effect against a two-sided alternative. Performing this test, we obtain a test statistic of $z = 2.249$, which is significant at the 97.5% level. We can obviously reject that treatments A and B are revenue equivalent in favor of the hypothesis that treatment B generates higher revenue. Thus, we get significant increase in revenue for auctions with a higher shipping and handling fee than auctions with free shipping even though the auctions end within one day of each other. This may suggest that the trade volume on eBay is large enough that these auctions are quite independent even though they overlap in time.

⁵ In the high reserve price treatment, the CDs “Oops!...I did it again” and “Automatic for the People” did not sell under either treatment; hence there is no revenue ordering for these items. The “day of the week” effect is also insignificant if we look at music CDs and Xbox games separately.

Indeed, since we chose music CDs and Xbox games that are popular on eBay, many auctions of identical objects went on at the same time.

Table 3. Revenues from Low Reserve Treatments

CD Title	Revenues	Revenues	B - A	Percent Difference
	under Treatment A	under Treatment B		
Music	5.50	7.24	1.74	32%
Ooops! I Did it Again	6.50	7.74	1.24	19%
Serendipity	8.50	10.49	1.99	23%
O Brother Where Art Thou?	12.50	11.99	-0.51	-4%
Greatest Hits - Tim McGraw	11.00	15.99	4.99	45%
A Day Without Rain	13.50	14.99	1.49	11%
Automatic for the People	0.00	9.99	9.99	
Everyday	7.28	9.49	2.21	30%
Joshua Tree	6.07	8.25	2.18	36%
Unplugged in New York	4.50	5.24	0.74	16%
<i>Average</i>	<i>7.54</i>	<i>10.14</i>	<i>2.61</i>	<i>35%</i>
<i>Average excluding unsold</i>	<i>8.37</i>	<i>10.16</i>	<i>1.79</i>	<i>21%</i>

Xbox Game Title	Revenues	Revenues	B - A	Percent Difference
	under Treatment A	under Treatment B		
Halo	34.05	41.24	7.19	21%
Wreckless	44.01	33.99	-10.02	-23%
Circus Maximus	40.99	39.99	-1.00	-2%
Max Payne	36.01	36.99	0.98	3%
Genma Onimusha	41.00	32.99	-8.01	-20%
Project Gotham Racing	37.00	38.12	1.12	3%
NBA 2K2	42.12	42.99	0.87	2%
NFL 2K2	26.00	33.99	7.99	31%
NHL 2002	36.00	37.00	1.00	3%
WWF Raw	33.99	40.99	7.00	21%
<i>Average</i>	<i>37.12</i>	<i>37.83</i>	<i>0.71</i>	<i>2%</i>

Table 4 compares revenues under the high reserve treatments. Here, one sees a distinct difference in the ranking of average revenues between CDs and Xbox games. Excluding the two CDs that went unsold under both treatments, the average revenues under treatment C are \$11.63 compared to \$11.26 under treatment D.

Table 4. Revenues from High Reserve Treatments

CD Title	Revenues		D - C	Percent Difference
	under Treatment C	under Treatment D		
Music	9.00	8.00	-1.00	-11%
Oops! I Did it Again	0.00	0.00	0.00	
Serendipity	12.50	13.50	1.00	8%
O Brother Where Art Thou?	11.52	11.00	-0.52	-5%
Greatest Hits - Tim McGraw	18.00	17.00	-1.00	-6%
A Day Without Rain	15.50	16.00	0.50	3%
Automatic for the People	0.00	0.00	0.00	
Everyday	10.50	13.50	3.00	29%
Joshua Tree	8.00	11.10	3.10	39%
Unplugged in New York	8.00	0.00	-8.00	-100%
<i>Average</i>	<i>9.30</i>	<i>9.01</i>	<i>-0.29</i>	<i>-3%</i>
<i>Average excluding unsold</i>	<i>12.15</i>	<i>12.87</i>	<i>0.73</i>	<i>6%</i>

Game Title	Revenues		D - C	Percent Difference
	under Treatment C	under Treatment D		
Halo	40.01	43.00	2.99	7%
Wreckless	35.00	36.00	1.00	3%
Circus Maximus	39.00	42.53	3.53	9%
Max Payne	37.50	42.00	4.50	12%
Genma Onimusha	36.00	37.00	1.00	3%
Project Gotham Racing	35.02	40.01	4.99	14%
NBA 2K2	41.00	45.00	4.00	10%
NFL 2K2	33.00	40.10	7.10	22%
NHL 2002	36.00	41.00	5.00	14%
WWF Raw	37.00	44.00	7.00	19%
<i>Average</i>	<i>36.95</i>	<i>41.06</i>	<i>4.11</i>	<i>11%</i>

Thus, for music CDs, we do not find that increasing shipping costs and reducing the opening bid one-for-one is revenue enhancing. However, this difference in revenues is reversed if we only look at revenues from the seven observations where the same CD sold under both treatments. Here, we find that treatment D yields an average revenue of \$12.87 compared to \$12.15 under treatment C. Regardless, the presence of the treatment effect on revenues is tenuous at best. In 4 matched pairs of auctions, treatment D had higher revenue

than treatment C, in another 4 matched pairs, this revenue ranking is reversed. Finally, in 2 auctions, the CD did not sell under either treatment, so the revenues are the same. Moreover, the test-statistic for Wilcoxon signed-ranks test is 0. Clearly, one fails to reject the null hypothesis of revenue equivalence at any level for this data; that is, a treatment effect is absent.

For Xbox games, the situation more closely resembles that in Table 3. Average revenues under treatment D are higher by \$4.11 compared to treatment C. This is an 11% difference in revenues. Further, in 10 out of 10 matched pairs of auctions, treatment D yields higher revenues than treatment C. The test-statistic for Wilcoxon signed-ranks test is 3.364 implying that we can reject the null hypothesis of revenue equivalence in favor of the one-sided alternative at any significance level.

The key feature that the low reserve price treatments and the high reserve price treatment with Xbox games share is that the reserve is less than 27% of the retail price. In contrast, for CDs under the high reserve price treatments, the reserve represents 53% of the retail price. Thus, the data suggests that the theoretical prediction of revenue equivalence fails systematically for relatively low reserve prices but not for relatively high reserve prices. In section 4, we will return to this idea and offer some *ex post* theoretical rationales drawn from behavioral economics to try to explain this difference.

Table 5 examines the effect of raising the total reserve price across treatments. Recall that treatment C and D are constructed by adding \$2 to the opening bid and \$2 to the shipping cost to treatments A and B, respectively. Having already established that pooling treatments A and B is not justified, we study revenues under C versus A (Table 5) and D versus B (Table 6) to get at reserve price effects.

Looking at CDs in table 5, we earlier observed that the likelihood of not selling the item was higher under treatment C (2 items unsold) compared to treatment A (1 item unsold). Conditional on a sale, revenues should be higher under treatment C compared to A. As the table shows, conditional on the sale of the item, average revenues under C are \$11.63 versus \$8.61 under treatment A. Further, in 7 of 8 auctions where the CDs were sold under both treatments, revenue was higher under treatment C than treatment A. For these 8 CDs, we obtain a *z*-value of 2.380 using the Wilcoxon signed-ranks test accepting the hypothesis that for low shipping treatments, higher effective reserve generated higher revenue conditional on sale at the 99% confidence level.

All items sold under both treatments for Xbox games, so there was no observable difference in the probability of a good going unsold. Interestingly, the average revenues under the high reserve treatment, treatment C, are lower than under treatment A. There is no systematic difference in the revenues for a given game across treatments. Treatment C yields higher revenues than A for 4 of 10

matched pairs, whereas A yields higher revenues than C for 5 out of 10 matched pairs. Revenues are exactly equal for the game NHL 2002. Thus, one cannot reject the null hypothesis of no treatment effect for this data.

Table 5. Effects on Revenues of Changes in Reserve Price: A vs C

CD Title	Revenues	Revenues	C - A	Percent Difference
	under Treatment A	under Treatment C		
Music	5.50	9.00	3.50	64%
Oops! I Did it Again	6.50	0.00	-6.50	-100%
Serendipity	8.50	12.50	4.00	47%
O Brother Where Art Thou?	12.50	11.52	-0.98	-8%
Greatest Hits - Tim McGraw	11.00	18.00	7.00	64%
A Day Without Rain	13.50	15.50	2.00	15%
Automatic for the People	0.00	0.00	0.00	
Everyday	7.28	10.50	3.22	44%
Joshua Tree	6.07	8.00	1.93	32%
Unplugged in New York	4.50	8.00	3.50	78%
<i>Average</i>	<i>7.54</i>	<i>9.30</i>	<i>1.77</i>	<i>23%</i>
<i>Average excluding unsold</i>	<i>8.61</i>	<i>11.63</i>	<i>3.02</i>	<i>35%</i>

Game Title	Revenues	Revenues	C - A	Percent Difference
	under Treatment A	under Treatment C		
Halo	34.05	40.01	5.96	18%
Wreckless	44.01	35.00	-9.01	-20%
Circus Maximus	40.99	39.00	-1.99	-5%
Max Payne	36.01	37.50	1.49	4%
Genma Onimusha	41.00	36.00	-5.00	-12%
Project Gotham Racing	37.00	35.02	-1.98	-5%
NBA 2K2	42.12	41.00	-1.12	-3%
NFL 2K2	26.00	33.00	7.00	27%
NHL 2002	36.00	36.00	0.00	0%
WWF Raw	33.99	37.00	3.01	9%
<i>Average</i>	<i>37.12</i>	<i>36.95</i>	<i>-0.16</i>	<i>0%</i>

Table 6 compares revenues under treatments B and D. Looking at CDs, again we find a higher probability that the item goes unsold with a high reserve price. This happens three times under treatment D and never under treatment B. Conditional on the item being sold, average revenues are higher (\$12.87) under

treatment D than under treatment B (\$11.21). Further, conditional on the item being sold, treatment D yields higher revenues in 6 out of 7 matched pairs. Results are quite similar for Xbox games. Average revenues are higher for treatment D compared to treatment B, and, indeed, in 10 out of 10 matched pairs, D yields higher revenues than B.

Table 6. Effects on Revenues of Changes in Reserve Price: B vs D

CD Title	Revenues	Revenues	D - B	Percent Difference
	under Treatment B	under Treatment D		
Music	7.24	8.00	0.76	10%
Ooops! I Did it Again	7.74	0.00	-7.74	-100%
Serendipity	10.49	13.50	3.01	29%
O Brother Where Art Thou?	11.99	11.00	-0.99	-8%
Greatest Hits - Tim McGraw	15.99	17.00	1.01	6%
A Day Without Rain	14.99	16.00	1.01	7%
Automatic for the People	9.99	0.00	-9.99	-100%
Everyday	9.49	13.50	4.01	42%
Joshua Tree	8.25	11.10	2.85	35%
Unplugged in New York	5.24	0.00	-5.24	-100%
<i>Average</i>	<i>10.14</i>	<i>9.01</i>	<i>-1.13</i>	<i>-11%</i>
<i>Average excluding unsold</i>	<i>11.21</i>	<i>12.87</i>	<i>1.67</i>	<i>15%</i>

Game Title	Revenues	Revenues	D - B	Percent Difference
	under Treatment B	under Treatment D		
Halo	41.24	43.00	1.76	4%
Wreckless	33.99	36.00	2.01	6%
Circus Maximus	39.99	42.53	2.54	6%
Max Payne	36.99	42.00	5.01	14%
Genma Onimusha	32.99	37.00	4.01	12%
Project Gotham Racing	38.12	40.01	1.89	5%
NBA 2K2	42.99	45.00	2.01	5%
NFL 2K2	33.99	40.10	6.11	18%
NHL 2002	37.00	41.00	4.00	11%
WWF Raw	40.99	44.00	3.01	7%
<i>Average</i>	<i>37.83</i>	<i>41.06</i>	<i>3.24</i>	<i>9%</i>

3.3 Bidders

Going back to Table 1, we look at how bidding behavior differed in each of the treatments for the music CD auctions. First, notice that treatment B averages about 0.6 more bidders per auction than treatment A. Treatment D averages only 0.2 more bidders than treatment C; thus, whatever effect transferring the reserve from the opening bid to shipping and handling fee has in inducing bidder participation, it is attenuated in the high reserve treatments. Similar analysis for Xbox games can be done using Table 2. Here there are an average of 8.8 bidders from Xbox games under treatment B, whereas there are only 7.8 under treatment A, which is consistent with what we saw for music CDs. An average of 1.6 more bidders bids under treatment D than under treatment C.

The comparison of number of bidders between low and high shipping treatments, on the surface, seems consistent with the revenue ranking we found in the previous section. Using the Wilcoxon signed-ranks test to test more formally for a treatment effect on the number of bidders, we find that the increase in high shipping treatments is not significant for music CDs and for low reserve treatments of Xbox games (treatments B vs. A). However, the increase is significant at 97.5% confidence-level for high reserve treatments of Xbox games (treatments D vs. C) with a test statistic of $z = 1.962$. Pooling the matched pairs of CD auctions and testing the null hypothesis that the number of bidders in high and low shipping treatments (treatments B vs. A and D vs. C) are equal, we obtain a test statistic of $z = 0.956$, which is not significant. Pooling the matched pairs of Xbox games auctions and testing the hypothesis that the number of bidders in high and low shipping treatments are equal, we obtain a test statistic of $z = 2.624$, which is significant at 99.5% level.

Next, we turn to changes in the level of the reserve price. Recall that increases in the level of the overall reserve price are predicted to decrease the number of bidders by excluding those with values between the two different reserve price levels. Comparing treatments A and C (or B and D) in Table 1, one sees evidence of this effect. Compared to treatment A, the average number of bidders under treatment C falls by 1.1 bidders. Likewise, compared to treatment B, the average number of bidders under treatment D is lower by 1.5 bidders. Qualitatively similar results do not occur for Xbox games as shown in Table 2. Compared to treatment A, the average number of bidders under treatment C increases by 0.6 bidders. Similarly, compared to treatment B, the average number of bidders in treatment D increases by 1.3 bidders. This strengthens the assertion that a very small fraction of bidders may have valuation below \$8 for Xbox games. The increase in bidders can be due to the fact that high reserve auctions were held one week after the low reserve auctions.

The theory predicts that auctions attracting a larger number of bidders will have higher expected revenues. To study this question, we compare the revenues under treatments A and B for each matched pair of CDs and ask how often the auction with the higher revenues attracted (weakly) more bidders. The evidence is consistent with the theory, in 9 out of 10 auctions, the auction with the higher number of bidders obtained the higher revenue. One obtains similar results from comparing treatments A versus B for Xbox games. In this case, the treatment attracting the higher number of bidders obtained the higher revenue in 8 out of 10 auctions. Comparing treatments C versus D, the CD auction with the higher number of bidders obtained the higher revenue in 6 out of 8 cases, while in 9 out of 10 cases, the Xbox auction with the higher number of bidders obtained the higher revenue. Thus, it seems that a likely transmission mechanism for the revenue ranking we obtained above is that auctions with lower opening bid amounts succeed in attracting more bidders than those with higher opening bids, holding the total reserve price constant.

An investigation of timing of the first bid an auction and the revenue suggests a more nuanced explanation, auctions with lower opening bid amounts attract earlier bidding and this, in turn, leads to higher revenues. There is a strong correlation between receiving the earlier first bid and obtaining higher revenue. In 29 out of 38 matched pairs of auctions, the treatment receiving the earlier bid obtained the higher revenue. Thus, it seems that, holding the overall reserve price fixed, auctions with lower opening bid amounts attract earlier bidders, and this in turn leads to greater overall bidding interest in the item, which ultimately results in higher revenues.

4. Discussion

While strategic equivalence predicts that two auctions for identical items with the same total reserve price should yield the same expected revenue, we find evidence that for auctions on eBay, lowering the opening bid and increasing shipping charges while holding the total reserve price constant raises the revenue of the seller. It seems to do this by attracting earlier and more bidders to the auction. In this section, we discuss a number of possible theoretical models that might rationalize the observed results.

4.1 Loss Aversion and Mental Accounting

Two essential characteristics of a value function to evaluate outcomes of risky prospects proposed by Kahneman and Tversky (1979) and Tversky and Kahneman (1991) are reference-dependence and loss aversion. Reference-dependence suggests that an agent defines gains and losses with respect to a

reference point. Loss aversion hypothesize that negative utility received from a loss is greater in size than utility received from a gain of the same amount.

Further, Kahneman and Tversky (1984) and Thaler (1985) postulate that consumers retain separate mental accounts for different aspects of a purchase decision. Experimental evidences show that a decision maker makes different choices when presented with relevant data in different accounting formats. One plausible way this might happen in our setting is that bidders have separate accounts for shipping and for the good itself.

Combining these observations allows us one way to rationalize our findings as follows: Suppose bidders have separate mental accounts for shipping and the good itself with less emphasis on shipping. Then, the treatment with higher shipping costs will generate higher expected revenue than the treatment with lower shipping cost. However, if bidders are loss-averse with respect to some reference level, then a very high shipping cost may lead to revenue equivalence and revenue performance can even be reversed.

We offer a simple model to operationalize this intuition. Suppose a consumer with valuation of v for some object has a reference level of v_s for shipping charges. When she pays a price of p and shipping cost of s in winning an auction, her utility is:

$$U(v, p, s) = \gamma u(v - v_s - p) + \alpha(v_s - s).$$

Here α is smaller than γ and for simplicity we assume $\gamma = 1$. We further assume that $u(0) = 0$ and, if x is positive, x and $-x$ refer to subtraction and addition of wealth respectively. If the utility function $u(\cdot)$ demonstrates loss aversion then $u(x) \leq -u(-x)$ and $u^{-1}(y) \geq -u^{-1}(-y)$ for all positive x, y . For a fixed s , the buyer resets her maximum bid to p^* where

$$p^* = v - v_s - u^{-1}(-\alpha u(v_s - s))$$

and $p^* \leq v - v_s + u^{-1}(+\alpha u(v_s - s)).$

Since the agent is loss-averse, if $\alpha u(v_s - s) > u(s - v_s)$ for all $s > v_s$ then a bidder's total expenditure on the object will be increasing in s for $s \leq v_s$ and decreasing in s for $s > v_s$. Hence, we can have a situation where

$$u(v - v_s - p_1 + \varepsilon) + \alpha u(v_s - s_1) \geq u(v - v_s - p_1) + \alpha u(v_s - s_1 + \varepsilon) \text{ but}$$

$$u(v - v_s - p_2 + \varepsilon) + \alpha u(v_s - s_2) \leq u(v - v_s - p_2) + \alpha u(v_s - s_2 + \varepsilon).$$

For example, suppose $s_1 = v_s$, $s_2 = 2v_s$, $p_2 = p_1 - v_s$, $\alpha > 0.5$, $\frac{\alpha}{1-\alpha}v_s \geq \varepsilon \geq v_s$ and

$$u(x) = \begin{cases} 2x & \text{for } x < 0 \\ x & \text{for } x \geq 0. \end{cases}$$

Plugging in these values, we obtain:

$$\begin{aligned} v - v_s - p_1 + \varepsilon &\geq v - v_s - p_1 + \alpha\varepsilon \text{ and} \\ v - p_1 + \varepsilon - 2v_s &\leq v - p_1 + \alpha(\varepsilon - v_s) \end{aligned}$$

That is,

$$\begin{aligned} U(v, p_1 - \varepsilon, s_1) &\geq U(v, p_1, s_1 - \varepsilon) \text{ but} \\ U(v, p_2 - \varepsilon, s_2) &\leq U(v, p_2, s_2 - \varepsilon). \end{aligned}$$

This implies that a bidder with valuation v bids higher in the auction with higher shipping cost in the first case and she bids higher in the auction with lower shipping cost in the second (high reserve price) case.

This hypothesis is testable. Let us assume that v_s is the same for all agents for some good. If shipping charges are smaller than v_s for treatment C but not for treatment D, treatment C may generate revenue higher than or equivalent to that of treatment D. On the other hand, if shipping charges are smaller than v_s for all treatments A through D then there will not be loss aversion effects for any treatment. Hence, treatments B and D will outperform treatments A and C respectively. It is worth mentioning that on eBay, average shipping and handling fee for new music CDs is between \$3 and \$4 and for new Xbox games is between \$4 and \$6. Thus, our findings from the music CD and Xbox games experiment are consistent with this explanation.

4.2 Salience

The eBay display makes the current second highest bid much more prominent than the shipping cost. Suppose that a fraction of the potential bidders simply ignore shipping costs when making their bids. In this case, it is still a weakly dominant strategy for each bidder to bid up to his or her expected surplus. Thus, when shipping is zero, the expected price is simply the second highest valuation for the object. When shipping is positive, a fraction of bidders discount for the amount of the shipping, whereas bidders who ignore shipping continue to bid up to their values for the object. This leads to a distribution of bids that first-order stochastically dominates the no shipping cost environment and one obtains the

revenue ranking we observe. Gabaix and Laibson (2005) show that when some agents are naïve, it might be optimal for sellers to hide information by using shrouded attributes. In our experiment, the shipping and handling fee is somewhat shrouded as it is not immediately visible when auctions are listed on eBay.

It seems plausible that the salience of shipping costs depends on two things, the ratio of shipping cost to the total cost of the item and the level of shipping cost itself. We postulate that the fraction of bidders ignoring shipping cost is given by $\lambda(\rho, s)$ where ρ is the ratio of shipping cost to the retail price of the good and s is the level of shipping cost. We also assume that $(\partial\lambda/\partial\rho) < 0$, $(\partial\lambda/\partial s) < 0$ and $\lambda(\rho, s) = 0$ if $\rho \geq \underline{\rho}$ for some $\underline{\rho}$.

Under these assumptions, we can conjecture the ancillary prediction that revenue differences, for a given total shipping cost, should be greater for Xbox games than for CDs (since ρ is smaller for Xbox) and revenue differences should be reduced for higher shipping treatments than for lower shipping treatments. First we test that the revenue difference (between the high shipping cost and low shipping cost treatments) for low reserve price treatments will be greater than that for high reserve price treatments. This is a difference-in-differences $((B - A) - (D - C))$ measure and we use the standard one-sided t-statistic for testing the null hypothesis that the mean of a normal population is zero. Here the difference-in-differences in revenue for a given music CD is considered as one observation. Next we test that the revenue difference for Xbox games are greater than that for music CDs. We test this for low reserve price and high reserve price treatments separately. We treat the revenue difference $(B - A$ or $D - C)$ for CDs and Xbox games as two different population samples whose means we would like to compare. Since these two populations may not have the same variance, we use the one-sided Smith-Satterthwaite test for comparing the means of two processes with different variances. The degree of freedom (DF) for this test involves both the number of observations in each sample and the sample variances. The results of these tests, presented in Table 7, are somewhat consistent with the salience mode.

Table 7. Differences in Revenue Differences Under Various Treatments

	Test Statistic	DF	$t(.05, DF)$
Difference-in-Differences $((B-A)-(D-C))$ for CD	2.32	9	1.83
Difference-in-Differences $((B-A)-(D-C))$ for Xbox	- 2.41	9	1.83
Difference in Xbox and CD Revenue Differences for Low Reserve (B-A)	- 0.89	13	1.77
Difference in Xbox and CD Revenue Differences for High Reserve (D-C)	3.73	16	1.75

Specifically, we reject (at 95% confidence level) the two alternative hypotheses that:

- i) for low reserve price treatments, revenue difference for Xbox games is larger than music CDs
- and
- ii) for Xbox games, revenue difference is higher for the low reserve price treatment.

We can reject (at 95% confidence level) the null hypotheses in favor of the one-sided alternatives implied by the salience model that:

- i) for high reserve price treatments, revenue difference for Xbox games is larger than music CDs
- and
- ii) for music CDs, revenue difference is higher for the low reserve price treatment.

4.3 Costly Search

Another alternative explanation postulates that the cost of search might account for our findings.⁶ Suppose that bidders use the following search strategy: Sort the prices for the desired item from lowest to highest then bid only on the low price item. In that case, since auctions run under treatments B and D have lower initial listed prices than those run under treatments A and C, these auctions would be more likely to attract earlier bidders. For example, suppose there are only two auctions and one is run under treatment A and the other under treatment B. Then, treatment B would get the first two bids. Since, with the same number of bidders placing bids in these two auctions, the expected “price” is always lower in treatment B (the expected “price” plus shipping fee are equal for a given number of bidders), treatment B will receive higher number of bids on average. This translates into higher revenue for treatment B. This is also consistent with the result that auctions attracting earlier bid generated higher revenue. Notice that the revenue difference goes away if we have infinite number of bidders. The revenue difference due to costly search will be larger the thinner the market (in terms of number of bidders) is.

While this may explain the revenue ranking and number of bidders we observe in the low reserve treatments for both types of products and in the high reserve treatment for Xbox games, it does not explain why revenue equivalence holds for the high reserve treatments in the CD auctions. In fact, the market size for CD auctions with high effective reserve is relatively thin implying that the revenue difference between treatment D and treatment C is likely to be larger.

⁶ We thank Richard Englebrecht-Wiggans for suggesting this possibility.

One point worth mentioning is that only 15 out of 20 such auctions resulted in sale and the smaller data set could be a driving force behind this equivalence. Costly search also does not explain why Xbox game auctions received higher number of bidders under high reserve treatments. Xbox game auctions with high reserves on average attracted 0.95 bidders per auction more than those with low reserves. This difference is statistically significant at 90% confidence level. However, there was a one-week difference between low reserve and high reserve auctions and that can be a cause behind the increase in bidders.

Finally, we offer several alternative explanations that are less successful in rationalizing our findings.

4.4 Suspicious Bidders

Another possibility is that some bidders have a preference for a more familiar auction format than a less familiar one. While free shipping is sometimes offered on eBay auctions, it is less common than auctions with positive shipping costs, so it may be less familiar to bidders. Further, bidders may view the familiarity of the auction as a signal of the quality of the seller. For instance, bidders may infer from the low shipping charge that the seller is a low quality seller. This might mean that they expect the speed of shipping to be slower (although we described the shipping method using identical language in both types of auctions) or that there is a greater possibility of not receiving the good at all (although we restricted payment to credit cards so that bidders are mostly insulated against this problem). For these reasons, it might be that fewer bidders choose to participate in less familiar auctions and that this accounts for the revenue ranking. This explanation seems less plausible in explaining the same revenue ranking for Xbox games under treatments C and D. There seems little a priori reason to suppose that bidders are less familiar with an auction with a \$2 shipping charge than a \$6 shipping charge. Moreover, if they are, then one should see a similar revenue ranking for music CDs under these two treatments, but this is inconsistent with the data.

There are further reasons to doubt this explanation. First, many online retailers, such as Barnes and Noble and Amazon have gone to a free shipping policy for purchases about some modest amount. Indeed, the trend in online retail seems to be toward more free shipping. So how unfamiliar is free shipping to bidders? Second, field experiments by List and Lucking-Reiley (2000) found confirmation of many of the theory predictions in comparing the uniform price auction to a 2-unit Vickrey auction, which is a very unfamiliar auction format.⁷

⁷ It should be mentioned here that in List and Lucking-Reiley (2000) experiment, the numbers of bidders in different treatments were exogenously given. It is possible that were the bidders free to

4.5 Love of Winning

Suppose that bidders obtain utility simply from winning the auction itself, in addition to their surpluses. In particular, suppose that the utility from winning is a decreasing function of the difference between the winning bid and the retail price of the object. In this case, since bidders are shading down bids for positive shipping treatments, the love of winning “bonus” is greater in when shipping is positive than when it is zero. This leads to revenue non-equivalence of auctions with different shipping and handling fee that is consistent with our findings. However, this does not explain the return of revenue equivalence in the high reserve price treatment for music CDs. Moreover, if love of winning is independent of price level or fully accounts for shipping costs, then revenue equivalence once again obtains. To explain failure of strategic equivalence between first-price and Dutch auctions, Cox, Roberson and Smith (1982) suggest that bidders may gain some utility from winning in a Dutch auction because of the suspense value of the auction structure. This is, in essence, similar to the love of winning. However, Cox, Smith and Walker (1983) reject that explanation using some further experiments.

5. Conclusions

We have shown that when the effective reserve price is not a large fraction of the retail price of the item, the structure of the reserve price affects the number of bidders, the timing of bids, and, most importantly, the revenue derived from the auction. Setting a low opening bid and a high shipping and handling cost yields systematically higher revenue than doing the reverse. When the effective reserve is more than 50% of the retail value of the item, we do not observe these systematic differences. We present a number of theories that attempt to ex post rationalize our findings. One theory suggests that bidders might be loss averse and maintain separate mental accounts for shipping charges and the price of the item itself. Under such a framework, we show that an apparently revenue-neutral variation in the structure of the effective reserve price can lead to revenue differences consistent with those we find. Another potential explanation can be that some fraction of bidders simply overlook the shipping charge when making bids. We would expect that as the shipping charge becomes large relative to the final price of the item, fewer people will ignore it. This theory also leads to systematic variation in revenue as a function of the structure of the reserve price. If many bidders sort the auctions of the desired item in terms of the “current price” and then bid on the lowest price item, that will also lead to auctions with

choose the auction where they want to bid, they would choose to bid in auctions with a more familiar format.

high shipping charges having revenue advantage over auctions with the same effective reserve but lower shipping charges.

We began the discussion of the paper by noting that dividing a price into two components was a common marketing practice in posted price markets as well as auctions. A central contribution of this paper is to show the effect of dividing the price in different attributes in an extremely transparent setting where possible payments were made crystal clear to bidders and determined endogenously. We find important revenue differences even in this relatively ideal setting. This suggests to us that the effects we identified are likely to be quite important in posted-price retailing as well. Indeed, the main conclusion we draw is that different framing of the same price as a sum of different attributes may significantly affect consumer behavior.

Appendix

This section gives the exact wording used in placing our auctions on eBay.

CD Auctions

Title: {Name of CD here} -- Brand New in Original Shrink Wrap

Text: This superhit CD is brand new in original shrink wrap. The songs are: {List of songs here.}

I accept only Paypal or Billpoint payments.

Buyer pays shipping and handling charge of {insert shipping charge here.}
I will ship via first-class mail.

I will ship to US and Canada only.

Happy Bidding.

For the free shipping treatment, this paragraph read:

There is no shipping and handling fee for this CD. I will ship via first-class mail.

I will ship to US and Canada only.

Happy Bidding.

Xbox Auctions

Title: BRAND NEW!!! {Name of Xbox Game here} XBOX GAME!!*SEALED*

Text: This superhit game for Microsoft XBOX system is brand new and is in original shrink wrap.

I accept payment only by Paypal or money orders.

Buyer pays a fixed shipping and handling charge of {insert shipping charge here}, and I will deliver by USPS first-class mail as soon as the payment is received. Payment should be sent within 10 days of the end of the auction.

I will ship to US and Canada only.

Happy Bidding.

For the free shipping treatment, this paragraph read:

There is no shipping and handling fee for this Xbox game. I will deliver by USPS first-class mail as soon as the payment is received. Payment should be sent within 10 days of the end of the auction.

I will ship to US and Canada only.

Happy Bidding.

References

- Cox, James C., Bruce Roberson and Vernon L. Smith (1982): "Theory and Behavior of Single Object Auctions," *Research Experimental Economics*, Vol. 2, Vernon L. Smith, editor, Greenwich, JAI Press.
- Cox, James C., Vernon L. Smith and James M. Walker (1983): "A Test that Discriminates between two models os the Dutch-First Auction Non-isomorphism" *Journal of Economic Behavior and Organization*, 14, 205-219.
- Davis, Douglas D., and Charles A. Holt (1993): *Experimental Economics*, Princeton, Princeton University Press.
- Ellison, Glenn, and Sara F. Ellison (2004): "Search, Obfuscation, and Price Elasticities on the Internet," Working Paper, MIT.
- Gabaix, Xavier, and David Laibson (2005): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, forthcoming.
- Kagel, John H. (1995): "Auctions: A Survey of Experimental Research," *Handbook of Experimental Economics*, John H. Kagel and Alvin E. Roth, editors, Princeton, Princeton University Press, 501-585.
- Kahneman, Daniel, and Amos Tversky (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263-291.
- Kahneman, Daniel, and Amos Tversky (1984): "Choices, Values, and Frames," *American Psychologist*, 39, 341-350.

- Katkar, Rama, and David H. Reiley (2005): "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment," Working Paper, Vanderbilt University.
- Krishna, Vijay (2002): *Auction Theory*, San Diego, Academic Press.
- List, John A., and David Lucking-Reiley (2000): "Demand Reduction in Multi-Unit Auctions: Evidence from a Sportscard Field Experiment," *American Economic Review*, 90, 961-972.
- Lucking-Reiley, David (1999): "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, 89, 1063-1080.
- Morwitz, Vicki G., Eric A. Greenleaf, and Eric J. Johnson (1998): "Divide and Prosper: Consumers' Reaction to Partitioned Prices," *Journal of Marketing Research*, 35, 453-463.
- Reiley, David H. (2005): "Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet," *RAND Journal of Economics*, forthcoming.
- Resnick, Paul, and Richard Zeckhauser (2002): "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System.," *The Economics of the Internet and E-Commerce, Advances in Applied Microeconomics, Vol. 11*, Michael R. Baye, editor, Amsterdam, Elsevier Science, 127-157.
- Roth, Alvin E., and Axel Ockenfels (2002): "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review*, 92, 1093-1103.
- Smith, Michael D. and Erik Brynjolfsson (2001): "Consumer Decision-making at an Internet Shopbot: Brand Still Matters," *Journal of Industrial Economics*, 49, 541-558.
- Thaler, Richard H. (1985): "Mental Accounting and Consumer Choice," *Marketing Science*, 4, 199-214.
- Tversky, Amos, and Daniel Kahneman (1991): "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *Quarterly Journal of Economics*, 106, 1039-1061.