



eBay's proxy bidding: A license to shill

Joseph Engelberg^{a,*}, Jared Williams^b

^a Department of Finance, Kenan-Flagler Business School, University of North Carolina, McColl Building CB 3490, Chapel Hill, NC27599-3490, United States

^b Department of Finance, Smeal College of Business, Penn State University, 344 Business Building, University Park, PA, United States

ARTICLE INFO

Article history:

Received 10 November 2006

Received in revised form 15 May 2009

Accepted 18 May 2009

Available online 6 June 2009

JEL classification:

D44

K42

Keywords:

eBay

eBay auctions

Shill bidding

Shilling

Incremental bidding

ABSTRACT

We introduce a bidding strategy which allows the seller to extract the full surplus of the high bidder in eBay auctions. We call this a “Discover-and-Stop” bidding strategy and estimate that 1.39 percent of all bids in eBay auctions are placed by sellers (or accomplices) who execute this strategy. We argue that this kind of shill bidding is unnecessarily effective due to eBay's proxy system and the predictability of other bidders' bids. We also model eBay auctions with shill bidding and find that, in equilibrium, eBay's profits are higher with shilling than without it. Finally, to determine whether bidders have an incentive to bid on their own items, we mimic the bidding behavior of shill bidders in actual eBay auctions and find some evidence of the strategy's success.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Within the last decade, the growth of online auctions as a means of trade has been enormous. eBay, which is nearly a monopolist in the U.S. market, has seen its community grow from 2 million users in 1998 to 181 million users in 2005.¹ More than \$44.3 billion in goods and services traded through eBay's sites in 2005.² To understand the magnitude of this figure, Slovenia, the world's 81st largest economy in 2005 according to the IMF, had GDP of \$43.7 billion in 2005.

Along with the growth in popularity of internet auction trade has come a concern by regulators about various forms of fraud in this environment. One such fraud is shill bidding. Shill bids are bids placed by a seller (or an accomplice) on his own auction. Such activity is a violation of eBay's user agreement³ and, in some cases, illegal. For example, on November 8, 2004, former New York Attorney General Eliot Spitzer announced the details of three cases that involved shill bidding in eBay auctions.⁴ Two of the cases involved civil settlements with the New York Attorney General's Office, while the third resulted in a felony charge of Combination in Restraint of Trade, a violation of New York's antitrust law. In one of the two civil cases, the

* Corresponding author.

E-mail addresses: joseph_engelberg@kenan-flagler.unc.edu (J. Engelberg), jmw52@psu.edu (J. Williams).

¹ Source: eBay's 10-k filing at <http://www.sec.gov/Archives/edgar/data/1065088/000095013406003678/f17187e10vk.htm#102>.

² Source: eBay's 10-k filing at <http://www.sec.gov/Archives/edgar/data/1065088/000095013406003678/f17187e10vk.htm#102>.

³ eBay's position on shill bidding is posted on their website: “Shill Bidding is bidding that artificially increases an item's price or apparent desirability, or bidding by individuals with a level of access to the seller's item information not available to the general Community. Shill Bidding is prohibited on eBay. Violations of this policy may result in a range of actions, including: Listing cancellation, Forfeit of eBay fees on cancelled listings, Limits on account privileges, Loss of PowerSeller status, Account suspension and Referral to Law Enforcement.”

⁴ See http://www.oag.state.ny.us/press/2004/nov/nov8a_04.html.

Attorney General alleged the accused placed a total of 610 bids in 106 auctions for automobiles. In his office's press release, Spitzer comments that, "the use of shill bids in on-line auctions illegally drives up prices and defrauds consumers. These cases and continuing efforts to monitor transactions should help maintain the integrity of on-line auctions." Even before these high publicity cases, the problem of shill bidding was significant enough to warrant the attention of the U.S. Congress. On June 25, 2001, former House Energy and Commerce Committee Chairman Billy Tauzin (R-LA) and former Committee Member Rep. Heather Wilson (R-NM) wrote letters asking the chief executives of Yahoo! Inc., Amazon.com Inc., and eBay Inc. the following questions: "What is the incidence of shill bidding in online auctions? Are online auction companies successful in detecting shill bidders? What steps could the companies take to reduce shill bidding in auctions?"⁵

Despite the government's interest and the increasing relevance of online trade, the economics literature has done little to address the questions posed by Congress or the interesting economic questions concerning shill bidding on eBay: does eBay have an incentive to prevent shill bidding? Are there aspects of eBay's auction mechanism that make it possible for shill bidders to efficiently execute their shill? Is there an optimal shill bidding strategy in eBay's environment? This paper begins to answer these questions by demonstrating the following: (1) there is a simple shill strategy in which a shill bidder can push the price of an item up to the highest bidder's bid without becoming the high bidder, (2) shill bidding in eBay auctions is unnecessarily effective due to eBay's proxy bidding system and the predictability of bidders' bids, (3) there are small changes eBay could make to its mechanism that would reduce shill bidding, and (4) eBay appears to have no incentive to make these changes since in equilibrium its profits are higher with shill bidders. Our results highlight the importance of auction design as a method to reduce shill bidding—a point also made by Ockenfels et al. (2007).

Beyond the political and legal interest in shill bidding, our paper also contributes to the growing internet auction literature. Ockenfels et al. (2007) provide a comprehensive review of the this literature and include a review of the shill bidding literature. There are few papers which find empirical evidence of shill bidding perhaps because it is difficult to detect (its possible illegality provides an incentive to shill bidders to hide their activity). The exception is Kauffmann and Wood (2005) who provide evidence of *reserve price shilling* in which sellers place abnormally large bids in order to avoid eBay's reserve price fees. Our paper is the second to find empirical evidence of shill bidding on eBay and the first to find it in the form of incremental bidding. To date, the literature's most common explanation for incremental bidding is that incremental bidders are naïve and confuse eBay's auction rules with those of an ascending English auction. Roth and Ockenfels (2002) suggest that some incremental bidders are shill bidders, but they provide no evidence that shill bidding actually occurs. Since the shill bidding strategy we detect is an incremental bidding strategy, we provide evidence that some incremental bidders are shill bidders. We estimate that 1.39 percent of all bids placed on eBay are due to the incremental bidding strategy we detect.

The paper is organized as follows: Section 2 provides a description of eBay's auction mechanism and a particular shill strategy that takes advantage of the mechanism and the predictability of bidders' bids. Section 3 provides evidence that shilling occurs and provides an estimate for its prevalence. We then examine why shill bidding occurs by examining its effects on sellers' and eBay's profits. In Section 4, we implement the Discover and Stop strategy in actual eBay auctions and show that it seems to be an optimal strategy for sellers in our sample, and in Section 5 we model eBay auctions with shill bidders and show that their equilibrium profits are higher with shill bidding than without it. Section 6 concludes.

2. eBay's auction mechanism

eBay runs a second-price proxy auction with bid increments and time priority.⁶ A bidder in this environment places a maximum bid which is not observed by other bidders, and eBay's proxy bidder bids on his behalf up to his maximum. In particular, as a response to any acceptable bid b placed by another bidder, the high bidder H 's proxy bidder bids $\min\{b_H, b + \varepsilon(p)\}$ on H 's behalf, where b_H is H 's maximum bid and $\varepsilon(\cdot)$ is the bid increment as a function of price, p .⁷

We illustrate the procedure and nuance of a typical eBay auction by way of example. At time 0, the auction begins with some starting price determined by the seller. In this example, the starting price is \$1.50 and the bid increment is \$0.50. At time 1, bidder A places a bid of \$5.00. Since there is no other bid to beat, the price remains at \$1.50. At time 2, bidder B (who must place a bid no less than the price plus the bid increment, i.e., \$2.00) places a bid of \$3.00. eBay's proxy bidder for bidder A then places the bid $\min\{3+0.50, 5\}=3.50$, so at time 2, A is the high bidder and the price is \$3.50.⁸ At time 3, bidder C places a bid of \$4.75. eBay's proxy bidder for A then places the bid $\min\{4.75 + \$0.50, \$5\} = \$5$, so at time 3, A is the high bidder and the price is \$5.00. It is worth noting that bid histories, which consist of bidders' user names and the bid amounts of all manually entered bids (as opposed to bids entered by a proxy bidder) other than the high bid were typically observable over the course of an auction prior to 2007 (See Fig. 1).⁹ Hence, at time 3, observers realize that bidder A 's maximum bid is \$5,

⁵ Energy and Commerce Committee News Release (2001).

⁶ In the case of a tie, the bidder who submitted the high bid earliest wins the item.

⁷ By "acceptable" bid, we mean any bid greater than or equal to the current price plus the bid increment.

⁸ The \$5.00 bid by A and the \$3.00 bid by B are called "manual" bids since they are manually entered by the bidders. The bid of \$3.50 is called a "proxy" bid since the bid was placed automatically on A 's behalf.

⁹ Sellers can choose to run auctions in which the bid amounts are observed but the bidders' identities are hidden. eBay and sellers often argue that hiding bidders' identities reduces the solicitation of bidders by eBay members and outside observers, i.e., it prevents others from offering to sell similar items to the losing bidders. On November 2, 2006, eBay announced its Safeguarding Member IDs Project. One of the project's objectives was to replace the bidders' user names with aliases (bidder1, bidder2, bidder3, etc.) on the bidder history page. This was implemented in early 2007.

User ID	Bid Amount	Date of bid
honestguy5858 (2)	US \$255.00	Jan-08-05 12:38:18 PST
mickbid (7)	US \$250.00	Jan-07-05 08:01:57 PST
honestguy5858 (2)	US \$250.00	Jan-08-05 12:38:08 PST
honestguy5858 (2)	US \$245.00	Jan-08-05 12:37:58 PST
honestguy5858 (2)	US \$240.00	Jan-08-05 12:37:48 PST
honestguy5858 (2)	US \$235.00	Jan-08-05 12:37:33 PST
honestguy5858 (2)	US \$230.00	Jan-08-05 12:37:21 PST
domant3 (13 ★)	US \$220.00	Jan-07-05 08:50:01 PST
mickbid (7)	US \$205.00	Jan-07-05 08:01:35 PST
domant3 (13 ★)	US \$200.00	Jan-07-05 07:43:35 PST
mickbid (7)	US \$200.00	Jan-07-05 08:01:24 PST
mickbid (7)	US \$175.00	Jan-07-05 08:01:15 PST
mickbid (7)	US \$170.00	Jan-07-05 08:01:02 PST
chs93teach (2)	US \$160.00	Jan-07-05 07:59:09 PST
mmbaxta (0)	US \$125.00	Jan-07-05 07:56:36 PST
mmbaxta (0)	US \$120.00	Jan-07-05 07:56:05 PST
ajd8203 (24 ★)	US \$100.00	Jan-07-05 07:23:35 PST

Fig. 1. Example of an eBay Auction's Bid History Page. Although eBay now hides the identity of bidders in its auctions, eBay reported the manually entered bids of each bidder except the high bidder at the time of our sample collection. The bid reported for the high bidder is the final price of the auction—which is the high bidder's bid only if the difference between the high bidder's bid and the second-highest bid is less than or equal to the increment. eBay sorts the bid entries by bid amount, not time, on its bid history page. We call bidder domant3 a (non-incremental) repeat bidder because he bid multiple times during the auction. We call mmbaxta, mickbid and honestguy5858 incremental bidders since they bid consecutively within a short window of time. Notice also that at 12:38:08 honestguy5858 discovered mickbid's high bid of \$250. Ten seconds later (at 12:38:18) honestguy5858 chose to beat mickbid's high bid of \$250 so that we classify honestguy5858 as a Discover-and-Go bidder. Had honestguy5858 stopped his incremental bidding at \$250, we would have classified him as a Discover-and-Stop bidder.

Table 1

An example of eBay's auction mechanism. The figure describes a typical eBay auction. eBay uses the pricing mechanism $\min\{b_H, b + \varepsilon(p)\}$ where b_H is the high bid, b is the current bid, and $\varepsilon(p)$ is the bid increment (\$0.50 in this example). Sellers also set the starting price (\$1.50 in this example). eBay's price is observed by auction participants while the current high bid is not; however, at time 3, observers can infer that A's maximum bid was \$5 due to eBay's pricing mechanism and the fact that bid histories are publicly observable.

Time	Bidder	Bid	Price	High bidder
0	–	–	\$1.50	–
1	A	\$5.00	\$1.50	A
2	B	\$3.00	\$3.50	A
3	C	\$4.75	\$5.00	A
4	D	\$8.00	\$5.50	D

since $b_3 + \varepsilon > P_3$, where b_3 is the amount of the (publicly observed) bid placed at time 3, ε is the bid increment, and P_3 is the price resulting from C's bid at time 3. At time 4, bidder D places a bid of \$8.00, so D becomes the high bidder, and the price moves to \$5.50. Table 1 illustrates the example.

A classification of bidders in this environment is useful and can be found within Table 2. Table 2 also includes several variables we will use later in our regressions.

Fig. 1 provides an example of the actual bidder history page displayed by eBay.

2.1. An example of a successful shill

Because a seller's revenue is an increasing function of the second highest bid, sellers have an incentive to increase the second highest bid as much as possible without becoming the high bidder. By way of example, we now illustrate how shill bidders can take advantage of eBay's proxy bidding mechanism to increase their revenue.

The environment is the same as in Table 1: the starting price is \$1.50, and the bid increment is \$0.50. Assume for the moment that bidders only bid on the dollar.¹⁰ At time 1, bidder B places a bid of \$4.00, and the price remains at \$1.50. Suppose the seller, S, has a shill account and desires to increase the price as much as possible without becoming the high bidder. At time 2, S places a bid of \$2.00, after which the price goes to \$2.50. Since bidders are assumed to only bid on the dollar, S knows that B's bid is at least \$3.00. At time 3, S bids \$3.00 and the price moves to \$3.50, so at time 4 S bids \$4.00, and the price only moves to \$4.00. S then realizes that \$4.00 is B's maximum bid, so S cannot push the price higher without becoming the high bidder, and he stops bidding. We call this type of bidder a "Discover-and-Stop" bidder because he bids incrementally until he discovers the high bid, and then he stops bidding. Table 3 illustrates the example.

¹⁰ An identical argument can be used if we instead assume that the seller has a high subjective probability that bidders bid only on the dollar. Based on our data of bidders' bids (see Fig. 6 and Table 9), this belief would be consistent with actual bidding behavior.

Table 2

Definitions. Since the language of eBay bidding strategies and bidder types may be unfamiliar to most audiences, below we provide a detailed list of definitions. The definitions include general terms as well as variables used throughout our paper. (Note that bidder types need not be mutually exclusive.)

	Definition
BidderCount	The number of bidders in an auction.
BRating	The eBay user rating of the bidder. An eBay user rating is a system that allows previous auction participants to leave feedback about a bidder or seller. A user's rating is the sum of its positive ratings minus the sum of its negative ratings. Since most feedback is positive, the user rating is a proxy for the number of auctions in which the bidder has previously participated.
Closing price	The price (i.e., winning bid) at which the auction ended.
DiscNGo bidder	A 10 min incremental bidder who bids until he discovers the current high bid and chooses to exceed the high bid. By "discovers the high bid" we mean that one bid in his bidding sequence, b , is less than or equal to the high bid but strictly greater than the high bid minus the increment (i.e., $b_H \geq b > b_H - \epsilon$). This allows him to infer the high bid without becoming the high bidder. By "chooses to exceed the high bid" we mean that the next bid following b is his bid.
DiscNStop bidder	A 10 min incremental bidder who bids until he discovers the current high bid and chooses not to exceed the high bid. By "discovers the high bid" we mean that the last bid in his bidding sequence, b_2 , is less than or equal to the high bid but strictly greater than the high bid minus the increment (i.e., $b_H \geq b_2 > b_H - \epsilon$). This allows him to infer the high bid without becoming the high bidder. By "chooses not to exceed the high bid" we mean that the next bid following his bidding sequence is not his bid.
Follow-up sample	A sample we take after observing the bidders and sellers in our primary sample. In the follow-up sample, we take data from each bidder's bid history page (see Fig. 2) in order to analyze the relationship in each (Bidder, Seller) pair we observe in the primary sample.
FracBidFLWUP	Defined for each bidder seller pair (B, S) found in the primary sample, it is the fraction of B 's bids that were on S 's auctions in the follow-up sample.
FracLoseFLWUP	Defined for each bidder seller pair (B, S) found in the primary sample, it is the fraction of NumAuctionsFLWUP in which B lost.
Incremental bidder	A repeat bidder whose multiple bids are consecutive, within a short period of time (and often differ by twice the bid increment). When we refer to a 10 Minute Incremental bidder, we mean that the window in which the consecutive bids must fall is 10 min.
INCR30 bidder	A 30 s incremental bidder whose bids are not in the last 5 h of the auction.
NoDisc bidder	A 10 min incremental bidder who stops bidding before he discovers or exceeds the current high bid.
NumAuctionsFLWUP	Defined for each bidder seller pair (B, S) found in the primary sample, it is the number of S 's auctions B bid on in the follow-up sample.
Primary sample	Our sample of 39,212 Event Ticket auctions which ended that ended between September 8 and September 24, 2004. Here we take data (e.g., bidder IDs, seller ID, etc.) from the home page of each auction as well as the bid history (see Fig. 1) page for each auction.
Repeat bidder	A bidder who bids two or more times in the same auction.
Shill Bidder	A seller (or an accomplice) who uses an alternate user account to bid on his own items.
Sniping	Bidding in the closing seconds of an auction.

Table 3

Example of Discover-and-Stop bidding. The table describes a particular type of shill bidding called Discover-and-Stop bidding. In this example, the bid increment is \$0.50 and the starting price is \$1.50. Assuming bidders only bid on the dollar, a shill bidder S can raise the price to bidder A 's maximum bid without becoming the high bidder by placing incremental bids and inferring A 's bid via eBay's pricing mechanism. At times 2 and 3, bidder S continues bidding because the new price exceeds S 's bid by the entire increment. At time 4, the price does not exceed S 's bid, so S knows A 's bid is \$4.

Time	Bidder	Bid	Price	High bidder
0	–	–	\$1.50	–
1	A	\$4.00	\$1.50	A
2	S	\$2.00	\$2.50	A
3	S	\$3.00	\$3.50	A
4	S	\$4.00	\$4.00	A

3. Existence of shill bidders

3.1. Our data

We divide our data into two samples: a "primary sample" and a "follow-up sample." Our primary sample consists of data collected from 39,212 Event Ticket auctions that ended between September 8, 2004, and September 23, 2004.¹¹ From each auction in our primary sample we observe the starting time, ending time, starting price, and seller ID. Also, from each auction's bid history page we collect information about each losing bid – including the bid amount,¹² the time of bid, the bidder's ID, and bidder's rating – placed in the auction (see Fig. 1).¹³

¹¹ We chose the Event Ticket category for two reasons: (1) we wanted substantial dispersion in bid amounts so that we could measure the predictability of bids across a wide range of values and (2) we believed there to be a large degree of heterogeneity of private values within this category.

¹² eBay hides the amount of the winning bidder's maximum bid; for each winning bid we observe the ID and rating of the bidder and the time of the bid.

¹³ Our sample of auctions does not include Buy-it-Now auctions, auctions in foreign currency, auctions which ended with no bids, or auctions where bidder identities were hidden.

There are two types of "Buy-it-Now" auctions: one allows bidders to either bid or choose to pay the "Buy-it-Now" price and thereby end the auction and win the item, and the other only allows bidders to buy the item at the "Buy-it-Now" price (i.e., there is no bidding). We exclude both from our data set.

Current auctions bid on by [rachel7399](#) (63 ★)

For auction items, bold price means at least one bid has been received.

1 - 10 of 10 total. Click on the column headers to sort

Item	Start	End	Price	Title	High Bidder	Seller
4547716832	May-03-05	May-10-05 13:42:52	US \$4,050.00	Ford : Crown Victoria	ivan_j	faustillo
4547840834	May-03-05	May-10-05 20:03:48	US \$1,225.00	1983 Chevrolet Chevy Corvette Custom Go Kart pristine	oldmark610zw9	4everfords
4547845962	May-03-05	May-10-05 20:18:33	US \$3,650.00	Ford : Crown Victoria	sinr_g	quyzebo.com
4548049257	May-04-05	May-11-05 13:09:19	US \$1,825.00	Ford : Crown Victoria	tennconst	statewideford
6531084380	May-08-05	May-11-05 13:15:50	US \$10.50	1 DALLAS MAVS vs PHX SUN PLAYOFF TICKET CENTER RD2 GM 4	local83drew	pcparts_tx
6531072938	May-08-05	May-11-05 17:20:00	US \$107.50	1 TICKETS DALLAS MAVERICKS PHOENIX SUNS 5/15 SEC 105 A	topline433	aw1518
4548236532	May-05-05	May-12-05 06:31:11	US \$2,850.00	Ford : Crown Victoria	ikenneavy	axeman610
4548325228	May-05-05	May-15-05 13:23:27	US \$3,500.00	Ford : Crown Victoria	jcosinfo	kcpolicecars
4548496808	May-06-05	May-16-05 08:08:09	US \$127,522.00	2004 Fountain 35' Lightning 2x525 EFI's Bravo XR's	rachel7399	importcaralley

Fig. 2. Example of an eBay Bidder's History Page. The figure is an example of eBay's bidder history page for a particular bidder (not in our sample). The page records the auctions that the bidder has bid on in the past 30 days, the seller of each auction, and the auction's high bidder.




Fig. 3. Primary Sample and Follow-up Sample In our primary sample of 39,212 Event Ticket auctions we collect auction data – starting time, ending time, starting price and seller ID – from each auction. From the bid history pages of each auction we also collect bidding data – each bid amount, time of bid, and bidder ID. From the bid history pages of our primary sample of auctions we find 89,917 unique bidders and 161,895 unique pairs of (Bidder, Seller). Using the 89,917 unique bidder IDs we perform a follow-up sample to examine each bidder's behavior outside the primary sample window. Using eBay's bidder history function for each bidder we observe how many auctions each bidder bid on, the seller in each auction and the high bidder in each auction. The follow-up sample region is completely outside the primary sample region to avoid contamination in our regression results. Also, the window of time in the follow-up sample preceding the primary sample is small because eBay's bidder history function only allows users to observe completed auctions within the last 30 days (it took 6 days to gather and compile our primary data). The window of time in the follow-up sample proceeding the primary sample is large since we could use the bidder history tool multiple times after our primary sample was gathered.

Our follow-up sample consists of data collected from the bidder history pages (see Fig. 2) of each of the 89,917 bidders who bid in one of the 39,212 auctions in our primary sample. From the bidder history pages we are able to determine on which auctions the bidder bid, the seller of each of the auctions, and the winning bidder of each of the auctions. We exclude from our follow-up sample the bidders' behavior during the primary sample timeframe because we do not want the samples to have overlapping data. Fig. 3 illustrates our sampling.

Table 4 has summary statistics for the data gathered in both the primary sample and the follow-up sample. The average ending price in an Event Ticket auction was \$182.11 with a standard deviation of 309.47 (Table 4). This is consistent with the large dispersion in ticket prices we would expect across a variety of events. Our 89,917 unique bidders and 19,193 unique sellers generate 161,895 unique pairs in our primary sample. Our auctions also confirm the prevalence of incremental bidding in eBay auctions: of the 433,818 bids we observe 259,954 (59.92 percent) are part of a 10-min incremental bidding sequence and 12,364 (2.85 percent) are part of a Discover-and-Stop bidding sequence. We also find that 29.6 percent of bidders who bid in the primary sample did not bid in the follow-up sample and that a majority of bidder seller pairs (95.8 percent) from the primary sample do not pair up in the follow-up sample.

3.2. Shilling variables

From the follow-up sample we create “shilling variables” – variables that indicate suspicious behavior between a bidder B and a seller S . In the extreme case, we expect a shill bidder to be a bidder B who bids on a large number of S 's auctions, who bids exclusively on S 's auctions, and who never wins.¹⁴ Using this intuition we construct the following shilling variables defined on (B, S) pairs:

NumAuctionsFLWUP – the number of S 's auctions B bid on in the follow-up sample.

FracBidFLWUP – the fraction of B 's bids that were on S 's auctions in the follow-up sample.

FracLoseFLWUP – the fraction of *NumAuctionsFLWUP* in which B lost.

¹⁴ As we will see in Section 4, however, it is not clear we should expect shill bidders to rarely win the auctions on which they bid: we won 7 out of the 30 auctions in which we mimicked the behavior of a shill bidder.




Fig. 4. Distributions of B_i , B_i among Discover-and-Stop Bidders and B_i given $B_i = V_i$. The graphs illustrate the empirical distributions of normalized measures of the distance between the current price and the current high bid [i.e., $B_i = (\text{current high bid} - \text{current price})/(\text{closing price})$] among incremental bidding sequences. We estimate these distributions using our sample of 39,212 Event Ticket auctions and restrict attention to bidding sequences (1) in the last day of the auction and (2) where the high bid and the current price differ by at least three bid increments. We make these restrictions because (1) the majority of Discover-and-Stop bidding occurs on the last day of an auction (and we suspect shill bidders are most likely to bid when higher bids have been placed later in the auction) and (2) the current high bid and the current price must differ by at least three bid increments for a Discover-and-Stop incremental bidder to execute his strategy (see Fig. 3). The first graph is the empirical distribution of B_i among all non-Discover-and-Stop incremental bidding sequences (before the winning bid is placed) – a measure of the distance a bidder must travel in order to reach the high bid. The second graph is our estimate of the distribution of B_i given $B_i = V_i$ – a measure of how far a bidder traveled given his valuation is identical to that of the high bidder. We estimate the histograms using one-sided, uniform kernel estimation with a bandwidth of .02. The third graph illustrates the distribution of B_i among Discover-and-Stop bidders – a measure of how far a bidder traveled given he is a Discover-and-Stop bidder. Given our assumptions on bidder behavior, the last distribution is a mixture of the first two.

Table 4

Primary sample and follow-up sample summary statistics. Below are summary statistics for our primary sample and our follow-up sample. Data from the follow-up sample relate to the 161,895 unique pairs of bidder B and seller S we find in the primary sample. NumAuctionsFLWUP is the number of S's auctions B bid on in the follow-up sample; FracBidFLWUP is the fraction of B's bids that were on S's auctions in the follow-up sample; and FracLoseFLWUP is the fraction of NumAuctionsFLWUP in which B lost. FracBidFLWUP is undefined when we observe no bids by B in any auctions in the follow-up sample. FracLoseFLWUP is undefined when we observe no bids by B on S's auctions in the follow-up sample.

	# of unique				
Primary sample					
Bidders	89,197				
Sellers	19,133				
(B, S) pairs	161,895				
	Mean	Median	S.D.		
Primary sample					
Bidder rating	53.05	8	251.89		
Seller rating	168.48	34	884.61		
Bids per auction	11.06	9	9.65		
Winning bid	182.11	132.5	309.47		
	Auctions w/at least one		Total # of bids placed by		
Primary sample					
Bidder(s)	39,212		433,818		
Repeat bidder(s)	29,542		340,627		
Incremental bidder(s)	27,820		259,944		
Disc-N-Go bidder(s)	10,077		18,122		
Disc-N-Stop bidder(s)	3,918		12,364		
	Mean	Median	S.D.	Percent of (B, S) pairs where variable is	
				Undefined (percent)	0 given defined (percent)
Follow-up sample					
NumAuctionsFLWUP	0.099	0	0.901	0.0	95.8
FracBidFLWUP	0.019	0	0.114	29.6	94.0
FracLoseFLWUP	0.834	1	0.336	95.8	11.9

Table 4 provides summary statistics for these variables. Ultimately, we would like some function of these shilling variables – which we call a Shill Score – to measure the shilling characteristics between a bidder and a seller. Such a score would be most useful as a ranking tool for enforcers (eBay or the government) so that they can best identify which (Bidder, Seller) pairs are most suspicious and warrant further investigation. Also, later in the paper we demonstrate the usefulness of a Shill Score as a predictive tool. In particular, the Shill Score of the second-highest bidder in an auction helps predict whether the high bidder will pay his entire bid, which is consistent with the idea that shill bidders aim to extract as much consumer surplus from the high bidder as possible.

3.3. Results

We formally define a “Discover-and-Stop” (or “DiscNStop”) bidder as a bidder who bids consecutively within a 10 min interval, discovers the high bid, and chooses not to exceed the high bid.¹⁵ We use a probit model to see whether the three shilling variables gathered from the follow-up sample help predict whether a bidder is a Discover-and-Stop bidder inside our primary sample. Our results are reported in Table 5. An observation in this setting corresponds to one (Bidder, Seller) pair in one auction. Hence, an auction with $n > 1$ unique bidders would generate n observations (we exclude all auctions with only one bidder since the Discover-and-Stop strategy requires more than one bidder). We use the follow-up sample data to predict behavior within the primary sample so that our predictions are out-of-sample. To further avoid contamination in our results we only consider (Bidder, Seller) pairs in which the bidder lost the auction in the primary sample in specifications (4) and (5) in Table 5.¹⁶

We find that bidders with lower user ratings are more likely to participate in the Discover-and-Stop bidding strategy which is consistent with the idea that these bidders are either shill bidders or inexperienced bidders. We also find that two

¹⁵ By “discovers the high bid” we mean that the last bid in his bidding sequence, b_2 , is less than or equal to the high bid but strictly greater than the high bid minus the increment (i.e., $b_H \geq b_2 > b_H - \varepsilon$). This allows him to infer the high bid without becoming the high bidder. By “chooses not to exceed the high bid” we mean that the next bid following his bidding sequence is not his bid.

¹⁶ We do this because a Discover-and-Stop strategy is, by definition, a *losing strategy*. We might expect a bidder who lost in a seller's auction to be more likely to bid again in another auction by that same seller (if he is selling a similar item). By restricting the observations in the regression to only losing (Bidder, Seller) pairs, we can be more confident our shilling variables are related to a shill bidder's behavior and not to a losing bidder's behavior.

Table 5

PROBIT models for DiscNStop = 1. Here we use a probit to model the binary event that a pair (Bidder, Seller) participates in a Discover-and-Stop strategy. For an auction run by seller S we create a (Bidder, Seller) pair for each bidder B who bid in S's auction. We set the variable DiscNStop = 1 for a (B, S) pair when B is the bidder in at least one Discover-and-Stop bidding sequence. We include in the regression three shilling-related variables: NumAuctionsFLWUP (the number of S's auctions B bid on in the follow-up sample), FracLoseFLWUP (the fraction of NumAuctionsFLWUP in which B lost) and FracBidFLWUP (the fraction of B's bids that were on S's auctions in the follow-up sample). BidderCount is the number of bidders in the auction, WBid is the size of the winning bid in the auction and BRating is B's eBay rating in the auction. We lose observations from specification (2) to (3) because FracBidFLWUP is only defined for (B, S) pairs in which B has observations in the follow-up sample. We lose observations from specification (3) to (4) because we only consider losing (B, S) pairs from the primary sample. We lose observations from specification (4) to (5) because FracLoseFLWUP is only defined for (B, S) pairs in which we observe B bidding on at least one of S's auctions in the follow-up sample. From all specifications we exclude auctions with only one bidder (a Discover-and-Stop strategy requires more than one bidder).

	(1)	(2)	(3)	(4)	(5)
Intercept	−2.172*** (0.016)	−2.177*** (0.016)	−2.219*** (0.020)	−2.093*** (0.021)	−2.158*** (1.078)
BidderCount	0.030*** (0.002)	0.030*** (0.002)	0.031*** (0.002)	0.021*** (0.003)	0.022** (0.009)
Closing Price	3.40 (16.0)	4.92 (15.1)	−0.201 (19.6)	−4.40 (21.3)	−10.50 (16.6)
BRating	−0.482*** (0.062)	−0.487*** (0.062)	−0.386*** (0.061)	−0.3691*** (0.061)	−0.2145 (0.163)
NumAuctionsFLWUP		8.12*** (1.70)	6.54*** (1.87)	6.52** (2.03)	7.71*** (2.09)
FracBidFLWUP			0.211*** (0.049)	0.221*** (0.051)	0.350*** (0.071)
FracLoseFLWUP					−0.033 (0.090)
# of Obs.	174331	174331	125228	104545	10087
# of Obs. w/DiscNStop = 1	4241	4241	2807	2711	295
Only losing (B, S) pairs	No	No	No	Yes	Yes

Standard errors are in parentheses.*** and** indicate significance at the 10 percent, 5 percent and 1 percent level respectively. Closing Price is multiplied by 1,000,000, BRating is multiplied by 1000, and NumAuctionsFLWUP is multiplied by 1000 to reduce the number of decimal places.

of the three shilling variables, NumAuctionsFLWUP and FracBidFLWUP, are positive in the regression and significant at the 1 percent level. Moreover, our finding that FracLoseFLWUP is insignificant is not surprising given that we won 7 of the 30 auctions in which we mimicked the behavior of a shill bidder (See Section 4). In fact, since there are an average of 4.6 bidders per auction in our sample, the likelihood of a non-shill bidder winning an auction (1 out of 4.6) is surprisingly close to the number of auctions we won when we mimicked the behavior of a shill bidder (7 out of 30). Concerning economic significance, we consider the model for large values of our shilling variables. When all variables are set to their means in specification (4), the probit predicts Prob (DiscNStop = 1) = 2.55percent. If we observe a bidder bidding exclusively on a seller's auctions (i.e., FracBidFLWUP = 1) then Prob (DiscNStop = 1) = 4.12 percent. Alternatively, for values of NumAuctionsFLWUP above the 99th percentile (i.e., above 7) the predicted Prob(DiscNStop = 1) increases to 2.82 percent. Although FracLoseFLWUP has a negative sign in specification (5), it is neither statistically significant (with a *p*-value of .72) nor economically significant (with a coefficient near zero).

Since the event DiscNStop = 1 is a tail event, we run OLS as a robustness check. The results are in Table 6 and are qualitatively similar to those found in Table 5. These results provide evidence that some shill bidders indeed use the Discover-and-Stop strategy to execute their shill.

From this probit we can recover a Linear Shill Score (LSSCORE) by using the coefficients on the significant shilling variables under specification (4):

$$LSSCORE = 0.00652 \times \text{NumAuctionsFLWUP} + 0.221 \times \text{FracBidFLWUP}$$

With this shill score in mind, we make two observations: (1) a shill score is a coarse representation of a binary event (either a bidder is a shill bidder or he is not) so that the tails of the shill score's distribution are of particular significance and (2) if shill bidding is successful, then when the second-highest bidder in an auction is suspicious (i.e., has a high shill score) it is more likely that the high bidder in the auction paid his entire valuation. Although we generally cannot observe the high bidder's bid in an auction, we can observe the auctions in which the high bidder paid the full amount of his bid. As illustrated in Section 2, we know a high bidder paid the full amount of his bid when the second highest bid plus the increment is greater than the highest bid. To get a sense of the economic significance of our shilling variables with respect to the high bidder's bid, we sort our auctions in the primary sample based on the LSSCORE of the second highest bidder. In Table 6 we report the frequency in which the high bidder paid the full amount of his bid for different percentiles of LSSCORE. We find that when the second highest bidder has an LSSCORE above the 95th percentile (i.e., he is "suspicious"), the empirical probability of the high bidder paying his full valuation is 12.4 percent (Table 7). When the second highest bidder is below the 95th percentile, the empirical probability of the high bidder paying his full valuation is 7.9 percent. The difference is statistically significant and surprisingly strong: *in our sample a high bidder is 56 percent more likely to pay his entire bid when the second-highest bidder*

Table 6

OLS models for DiscNStop = 1. Here we run OLS regressions to model the binary event that a pair (Bidder, Seller) participates in a Discover-and-Stop strategy. For an auction run by seller S we create a (Bidder, Seller) pair for each bidder B who bid in S's auction. We set the variable DiscNStop = 1 for a (B, S) pair when B is the bidder in at least one Discover-and-Stop bidding sequence. We include in the regression three shilling-related variables: NumAuctionsFLWUP (the number of S's auctions B bid on in the follow-up sample), FracLoseFLWUP (the fraction of NumAuctionsFLWUP in which B lost) and FracBidFLWUP (the fraction of B's bids that were on S's auctions in the follow-up sample). BidderCount is the number of bidders in the auction, WBid is the size of the winning bid in the auction and BRating is B's eBay rating in the auction. We lose observations from specification (2) to (3) because FracBidFLWUP is only defined for (B, S) pairs in which B has observations in the follow-up sample. We lose observations from specification (3) to (4) because we only consider losing (B, S) pairs from the primary sample. We lose observations from specification (4) to (5) because FracLoseFLWUP is only defined for (B, S) pairs in which we observe B bidding on at least one of S's auctions in the follow-up sample. From all specifications we exclude auctions with only one bidder (a Discover-and-Stop strategy requires more than one bidder).

	(1)	(2)	(3)	(4)	(5)
Intercept	0.013*** (0.001)	0.012*** (0.0001)	0.011*** (0.001)	0.016*** (0.001)	0.010 (0.007)
BidderCount	0.002*** (0.000)	0.002*** (0.000)	0.002*** (0.000)	0.001*** (0.000)	0.00143** (0.0006)
Closing Price	0.204 (0.982)	0.309 (0.982)	0.001 (1.02)	−0.229 (1.17)	−5.279 (10.08)
BRating	−0.010*** (0.001)	−0.010*** (0.001)	−0.008*** (0.001)	−0.010*** (0.002)	−0.006 (0.005)
NumAuctionsFLWUP		0.657*** (0.132)	0.518*** (0.136)	0.602*** (0.166)	0.751*** (0.184)
FracBidFLWUP			0.013*** (0.003)	0.016*** (0.004)	0.026*** (0.005)
FracLoseFLWUP					−0.002 (0.006)
# of Obs.	174331	174331	125228	104545	10087
# of Obs. w/DiscNStop = 1	4241	4241	2807	2711	295
Only losing (B, S) pairs	No	No	No	Yes	Yes

Standard errors are in parentheses.***and*** indicate significance at the 10 percent, 5 percent and 1 percent level respectively. Closing Price is multiplied by 1000000, BRating is multiplied by 1000, and NumAuctionsFLWUP is multiplied by 1000 to reduce the number of decimal places.

Table 7

Max-payment frequency and the Shill Score of the 2nd highest bidder. The table describes the frequency in which the high bidder pays his entire bid for various percentiles of the LSSCORE between seller S and the second highest B₂. LSSCORE = $0.0065223 \times \text{NumAuctionsFLWUP} + 0.2215 \times \text{FracBidFLWUP}$ and is taken from specification (4) in Table 5. The number of auctions listed here is less than our full sample because we only include observations for which the LSSCORE is defined for (B₂, S). The column labeled "High Bidder Paid Max" refers to the percentage of auctions in which we observe the high bidder paying his entire bid; we can infer that the high bidder has paid his maximum when the second highest bid plus the increment is greater than the highest bid.

LSSCORE percentiles of (B ₂ , S)	Unique (B ₂ , S) pairs	Mean LSSCORE	# of auctions	High bidder paid max (percent)
Below 90	16981	0	17609	7.9
90–95	864	0.026	978	7.7
Above 95	691	0.212	978	12.4

is suspicious (has a high LSSCORE with the seller) than when the second-highest bidder is unsuspicious (has a lower LSSCORE with the seller).¹⁷

Even more evidence of shilling behavior via incremental bidding comes from the type of auction that sellers choose. At the time of our sample, sellers had the option of making the bidders' identities public or private.¹⁸ We might think that private auctions are a shill bidder's paradise since the seller can create an alternate user account, bid on his own item and blend in perfectly with the crowd of other bidders. Thus – all other things equal – we expect to see more incremental bidding in these markets. However, deciding what is an incremental bid in a private market (with hidden identities) is difficult. When we see two consecutive bids, how are we to know whether these bids were placed by the same bidder? We treat this issue by making the following observation of bidding behavior in public markets. We count the number of times consecutive bids are placed within 30 s of each other before the last 5 h of the auction (when bidding activity significantly increases), and we name such consecutive bids INCR30 bids. In our original sample of 39,212 public auctions we observe this behavior 84,654 times. Out of the 84,654 instances of INCR30 bids in public bidder auctions, all but 8 of the consecutive bids were made by the same bidder. Simply put, in a public auction in which we observe consecutive bids placed within 30 s of each other before the last 5 h of an auction, we are 99.99 percent sure those bids were placed by the same bidder. If we are nearly sure that such bids are placed by the same bidder, it is natural to ask how many INCR30 bids we observe in private auctions. Table 8

¹⁷ One objection to our results is the following: sellers who engage in shill bidding will be careful to avoid detection, and therefore the evidence we find must be a result of some other phenomenon. We believe some shill bidders try to hide their shilling, but we observed several (Bidder, Seller) pairs with high LSSCORES such that the bidder name was a simple permutation of the seller name. For example, one (Bidder, Seller) pair with a high shill score was (ABC,CBA) where A and B represent 5-letter words and C represents a 4-letter word. We disguise the actual (Bidder, Seller) identities for privacy reasons.

¹⁸ In current eBay auctions eBay now hides the identities of its bidders.

Table 8

Incremental bidding in public and private auctions. The table records bidding in public auctions (where bidder identities are known) and private auctions (where bidder identities are unknown). INCR30 are 30-s incremental bids (consecutive bids placed within a 30 s window).

Type	N	Bids	Bids/auction	INCR30 bids	INCR30 bids/auction
Private	1,859	23,522	12.65	6,416	3.45
Public	39,212	433,818	11.04	84,654	2.15

summarizes these results from a sample of 1859 private auctions in the Event Ticket category between the same dates as our public sample.

We can see from the figure that while the average number of bids per auction is only slightly higher in the private case, the number of INCR30 bids per auction is significantly higher. This suggests there is considerably more incremental bidding in private auctions, a conclusion that is consistent with the existence of shill bidders using an incremental bidding strategy to drive up the price in private auctions while remaining anonymous.

3.4. More evidence of shilling and an estimate of its prevalence

The previous results suggest competitive shill bidding occurs, but they do not give us any indication of how many bids are placed by shill bidders. We now present additional evidence that shilling occurs and in the process obtain an estimate for the amount of shilling that occurs.

First, consider the following definitions of (mutually exclusive) bidder types:

DiscNStop bidder — a 10 min incremental bidder who bids until he discovers the current high bid, at which point he stops bidding to avoid becoming the high bidder.

DiscNGo bidder — a 10 min incremental bidder who discovers the high bid and then bids again so that he becomes the high bidder.

NoDisc bidder — a 10 min incremental bidder who stops bidding before he discovers or exceeds the current high bid.

Consider an incremental bidder i with valuation v_i for an item in an auction with current price p and current high bid b . Let p_T denote the closing price of the item. (The variables v_i and p are known by the incremental bidder when he starts his incremental bidding, whereas b and p_T are not.¹⁹) Define the variables: $V_i = (v_i - p)/(p_T)$ and $B_i = B_i(b) = (b - p)/(p_T)$. V_i and B_i are relative measures of the distance between the current price and bidder i 's valuation and the current price and the current high bid, respectively.²⁰

In this setting, the reader can think of non-shill incremental bidders as bidders who bid incrementally because they are naive and confuse eBay auctions with an ascending English auction as in Roth and Ockenfels (2002). Assuming non-shill incremental bidders do not stop bidding until they become the high bidder or the price exceeds their valuation, and that they do not bid above their valuation, for each DiscNGo bidder we observe B_i and conclude that $V_i > B_i$, and for each NoDisc bidder we observe V_i and conclude that $B_i > V_i$.²¹ If there were no shill bidders, for each DiscNStop bidder we would observe V_i and B_i and conclude that $V_i = B_i$. However, based on our earlier regression results, it is reasonable to conclude that there are two types of DiscNStop bidders: non-shill bidders such that $b = v_i$, and shill bidders. In addition to our assumption on non-shill incremental bidding behavior, we also assume that shill bidders bid incrementally until they discover the current high bid and stop before they surpass it. Given this assumption, the distribution of the variable $B_i = (b - p)/p_T$ among DiscNStop bidders is a mixture of the conditional distribution of the variable B_i given $B_i = V_i$ (the distribution corresponding to the bidding behavior of non-shill DiscNStop bidders) and the distribution of the variable B_i (the distribution corresponding to the bidding behavior of shill DiscNStop bidders).²² By calculating the weights of this mixture, we are able to estimate the number of DiscNStop shill bids that are placed in eBay auctions.

Because we do not have access to data on the actual winning bid amounts, we assume that $\{B_i(b) : b \text{ is the winning bid}\}$ has the same distribution as $\{B_i(b) : b \text{ is not the winning bid}\}$. Using this assumption, we analyze the bids placed by non-DiscNStop incremental bidders before the winning bid is submitted to approximate the distribution of B_i . To approximate the distribution of B_i given $B_i = V_i$, we analyze the bids placed by NoDisc incremental bidders before the winning bid is submitted and apply uniform 1-sided kernel estimation.²³ Finally, we observe B_i among the set of DiscNStop bidders. Histograms for the three empirical distributions follow (Fig. 4).

¹⁹ Technically, i knows b when he begins his bidding if and only if the bidder who bid before him discovered the high bid without becoming the high bidder.

²⁰ We divide the expressions $v_i - p$ and $b - p$ by p_T because a \$20 difference between the current high bid and the current price seems significant in auctions for which the closing price is \$25, while the \$20 difference is not as significant in auctions for which the closing price is \$5000.

²¹ Actually, we only obtain estimates for B_i and V_i due to eBay's bid increments. This estimate is within $(\varepsilon)/(p_T)$ units of the actual value, where ε is the bid increment and p_T is the closing price.

²² This assumes the variable p_T is independent of whether the seller is a shill bidder.

²³ We do not include DiscNGo bidders in this dataset (i.e., we do not use a two-sided kernel estimation) because we are not confident that their bids are equal to their valuations. DiscNGo bidders often discover the high bid and beat it by the increment. This is consistent with the literature's claim that some incremental bidders confuse eBay's proxy system with an English auction (Roth and Ockenfels, 2002).




Fig. 5. CMFs of B_i , B_i among Discover-and-Stop Bidders and B_i given $B_i = V_i$. This figure plots the cumulative mass functions (CMFs) of the distributions plotted in Fig. 6. We find 400 points of each CMF using bin sizes of .0025.

Fig. 5 plots the cumulative mass functions (CMFs) of all three distributions. As expected, the CMF of B_i among DiscNStop bidders lies almost entirely between the two other distributions, which supports the claim that the distribution of B_i among DiscNStop bidders is a weighted average of the distribution of B_i and the distribution of B_i given $B_i = V_i$.

Define the variable B_i^{DS} to have the same distribution as $\{B_i : i \text{ is a DiscNStop bidder}\}$ and the variable $B_i^{B=V}$ to have the same distribution as $\{B_i : i \text{ 's valuation equals the current high bid (i.e., } b = v_i)\}$. Let $Y_i = B_i^{DS} - B_i^{B=V}$ and $X_i = B_i - B_i^{B=V}$ so that $Y_i = \alpha X_i$ if and only if $B_i^{DS} = \alpha B_i + (1 - \alpha)B_i^{B=V}$, i.e., $Y_i = \alpha X_i$ if and only if α is the proportion of DiscNStop bids that are shill bids.

We partition each CMF into 400 bins (each of size .0025) so that we have 400 points on each CMF. We find an estimate for α by minimizing the following objective function under square loss:

$$\min_{\alpha} \sum_{i=1}^{400} [Y_i - \alpha X_i]^2$$

This yields an estimate for α of .4877. Therefore, we estimate that 48.77 percent of DiscNStop bids are shill bids. Since 2.85 percent of all bids are DiscNStop bids, one back-of-the-envelope estimation for the number of shill bids generated by the Discover-and-Stop strategy is $0.4877 \times 0.0285 = 1.39$ percent.

In reality this is a conservative estimate for the prevalence of shilling; it only estimates the number of bids placed by sellers who employ our Discover-and-Stop shill strategy. Some shill bidders may bid incrementally but mistakenly surpass the high bid, and others might stop their incremental bidding before reaching the current high bid. Our results do not estimate the prevalence of these or other forms of shilling.²⁴

Notice that this argument for the existence of shillers focuses on the distribution of bidding sequences among bidder types. In particular, unlike the previous argument, it does not analyze how often a particular bidder bids on a particular seller's auctions. Thus, our two arguments are independent of one another, and yet both provide evidence of shill bidding in eBay auctions.

3.5. Predictability of Bids²⁵

Recall that Discover-and-Stop bidding is an effective shill technique when bidders bid in predictable units—especially when those bids are multiples of the bid increment. Fig. 6 describes the distribution of bids less than \$100, and Table 9 describes the frequency at which bidders bid on multiples of the increment at each price range. Fig. 6 and Table 9 suggest that sellers are indeed able to execute the Discover-and-Stop shill strategy.

²⁴ Kauffmann and Wood (2005) find evidence of a different form of shill bidding. Since eBay charges sellers a fee for using a secret reserve price, sellers who wish to use a secret reserve might have an incentive to set a low minimum bid and use a shill account to place a bid of their desired reserve amount. Kauffman and Wood call this *reserve price shilling* and document it in their paper. Our paper is the first to find evidence of *competitive shill bidding*: shilling where the seller waits for bidders to bid and then tries to drive the price up to the highest bidder's bid without becoming the high bidder.

²⁵ Our evidence of predictable bidding in eBay auctions indirectly adds to an existing literature concerning data clustering in financial markets (see Ball et al., 1985; Harris, 1991; Goodhart and Curcio, 1991; Hornick et al., 1994; Christie and Schultz, 1994; Grossman et al., 1997; Ikenberry and Weston, 2008).




Fig. 6. Distribution of Bids less than \$100. The figure plots the total number of bids observed on each dollar between \$0 and \$99 in our sample of 39,212 auctions. 83.18 percent of bids over this interval were on the dollar; the figure also suggests that bidders bid on familiar multiples (\$5, \$10, \$50, etc.).

Table 9

Bidding frequency on multiples of the bid increment. The table describes eBay's bid increments as a function of price. Discover-and-Stop bidders can effectively execute their strategy if bidders bid on a multiple of the increment. Here we record bidders' propensity to bid on multiples of the increment. We notice that a majority of submitted bids are indeed multiples of the increment and some price ranges present better opportunities for shill bidders than others. For example, shill bidders will find it easier to push a high bidder to his maximum in the \$100.00–249.99 price range than in the \$500–999.99 price range. The anomaly in this table is clearly the first tier where only 15.02 percent of bidders bid on a multiple \$0.10. This comes from the fact that several eBay sellers set a minimum bid of \$0.99 and bidders often bid the minimum in auctions.

Price range	Bid increment	# of observed bids	Percent of observed bids on the increment (percent)
\$0.01–0.99	\$0.05	1,125	22.13
\$1.00–4.99	\$0.25	5,498	75.50
\$5.00–24.99	\$0.50	38,121	79.28
\$25.00–99.99	\$1.00	141,691	84.80
\$100.00–249.99	\$2.50	155,751	77.10
\$250.00–499.99	\$5.00	53,010	80.34
\$500.00–999.99	\$10.00	10,203	66.98
\$1000.00–2499.99	\$25.00	1,421	75.65
\$2500.00–4999.99	\$50.00	140	66.43
\$5000.00 and up	\$100.00	47	82.98

It is remarkable that bidders choose to bid on the dollar even though the vast majority of other bidders bid on the dollar, too. If a bidder were considering bidding \$26, it seems that a bid of \$25.02 would yield a strictly higher expected payoff since with high probability no other bidder will bid between \$25.01 and \$25.99; essentially, the bidder would win the good with a bid of \$26 only if he also would have won with a bid of \$25.02, and in the event that the second highest bid were \$25, the bidder would save \$0.98.

4. The effects of shilling on sellers' revenues

We argued in Section 2.1 that sellers have an incentive to bid on their own auctions using shill accounts. To verify this claim and estimate how much sellers can benefit from following this strategy, we employ the DiscNStop bidding strategy in the closing minutes of actual eBay auctions and record the results.

Ideally, we would sell items ourselves and bid on them using a shill account—i.e., we would conduct an actual shill. Since this approach is illegal and against eBay's rules, we employ the DiscNStop bidding strategy in the closing minutes of other users' eBay auctions. That is, we mimic the behavior that the seller of the item might employ if he or she were a shill bidder. Since sellers must use alternate user accounts when conducting their shill bids, to outside observers and auction participants our bidding is observationally equivalent to actual DiscNStop shill bids.

4.1. Data

To accurately estimate our effect on the items' final prices, we restrict our bidding to simultaneous auctions. That is, we restrict our attention to pairs of auctions for identical goods offered by the same seller whose ending times are within

minutes of one another. For each pair of simultaneous auctions, we randomly choose one of the auctions to participate in. We label those auctions “treatment” auctions, and we label the auctions in which we do not bid “control” auctions. By comparing the final prices of the control auctions and the treatment auctions we can estimate the profitability of shilling.

Our sample consists of the bid histories of sixty auctions for bleacher seat tickets to Chicago Cubs home games: thirty in which we participated (the treatment auctions), and thirty in which we did not participate (the control auctions).²⁶ The items within each pair are identical but vary across pairs as the Cubs play different teams on different days at different times. While each pair of auctions will have similar prices (since they are identical items), the average price of each pair varies substantially across pairs. Each pair had an average closing price of \$75.22 with a standard deviation of \$53.45.

The closing times between our treatment auctions and control auctions also vary from 0 s (simultaneous ending times) to 15 min, with a median difference in closing times of 41.5 s.

4.2. Algorithm rules

We bid according to the following algorithm:

- (1) Randomly choose which of the two auctions to bid.
- (2) Five minutes before the earlier auction's close, determine whether the high bidder's maximum bid has already been discovered in the auction on which we plan to bid. If the maximum bid has already been discovered, participate in the other auction instead.
- (3) Five minutes before the treatment auction's close, bid the minimum acceptable bid (the current high bid plus the bid increment). Stop if and only if we discover the high bidder's maximum bid. This can happen in two ways: either we become the high bidder ourselves or the high bidder's proxy bidder beats our bid by less than the bid increment. The latter case is the optimal outcome for a shill bidder.
- (4) After completing the incremental bidding strategy outlined in steps 1–3, if another bidder bids and becomes the high bidder, we repeat the incremental bidding sequence in the treatment auction.

We had to repeat the incremental bidding sequence a total of 16 times in 11 different auctions during our bidding in the 30 auctions. There are two ways we were forced to do this. Of the 16 times we repeated the incremental bidding sequence, 8 times we had become the high bidder in our earlier DiscNStop attempt, and the bidder we had outbid came back to outbid us. In the other 8 auctions, an outside bidder entered to outbid us.

The price movements during the last five minutes, the difference in closing times between the treatment and control auctions, and the shipping costs of all 60 auctions are listed in Table 10. We also record indicators for whether we won the auction and whether the treatment auction ended before or after the control auction.

The average price of the treatment (control) auctions 5 min before its close is \$65.17 (\$67.00), and the average closing price of the treatment (control) auctions is \$76.65 (\$73.79). Table 11 provides these and more summary statistics about the sample.

4.3. Results

Recall the goal of a DiscNStop shill strategy: for the seller to increase the closing price of the item by determining the high bidder's maximum bid (without becoming the high bidder himself).

As predicted, the average closing price of the treatment auctions (\$76.65) is higher than the average closing price of the control auctions (\$73.79). In 22 out of the 30 auction pairs, the treatment auction had a higher closing price than its corresponding control auction. The difference in closing prices between the treatment auctions and control auctions is strongly significant under the sign test and Wilcoxon signed-rank test ($p = .008$ and $.001$, respectively) and weakly significant under the paired t -test ($p = .07$). However, these numbers are misleading because the seller would not have been able to sell all 30 items at these prices; we won 7 of the 30 treatment auctions.²⁷ This demonstrates that the risk to a seller involved with a DiscNStop strategy. Restricting attention to the 23 auctions which we did not win, the average closing price of the treatment auctions (\$85.52) is still higher than the average closing price of the control auctions (\$82.99). Of these 23 auction pairs, the treatment auction had a higher closing price than its paired control 17 times. Among these 23 auctions, the difference in closing prices between the treatment auctions and control auctions is strongly significant under the sign test and Wilcoxon signed-rank test ($p = .017$ and $.003$, respectively), but insignificant under the paired t -test ($p = .16$).

We also examine the effectiveness of shill bidding by analyzing the difference between the auctions' closing price and the price five minutes before the auctions' close. For each auction, we define the variable *RUNUP* as the difference between the auction's closing price and its price five minutes before the close. We hypothesize that shill bidding increases the expected

²⁶ The items in the control auctions are identical to the ones in the treatment auctions because the tickets are general admission and for the same game.

²⁷ We paid the sellers for the items.

Table 10

Features of treatment and control auctions. We list features of each of the 60 auctions in our sample. The “Win?” column records whether or not we won the auction. “Treatment First?” records whether or not the treatment auction ended before the control auction. Time Difference is the difference in ending times, in seconds, between the treatment auction and the control auction. The “Price Path” columns list the price movements during the last 5 min of the control auctions and treatment auctions.

Win?	Treatment First?	Time Difference	Price Path (Control): last 5 min → end		Price Path (Treatment): last 5 min → end		Shipping
No	Yes	74	\$90.00	\$90.00	\$91.00	\$100.00	\$10.00
No	No	44	\$106.50	\$106.50	\$112.50	\$125.00	\$6.00
No	Yes	4	\$127.50	\$127.50	\$127.50	\$130.00	\$7.50
No	Yes	3	\$110.00	\$151.00	\$106.94	\$153.50	\$7.50
No	No	4	\$164.72	\$175.61	\$166.56	\$182.50	\$7.50
Yes	No	26	\$29.00	\$29.00	\$31.00	\$40.00	\$12.00
No	Yes	35	\$177.50	\$182.50	\$177.50	\$180.00	\$10.00
No	Yes	266	\$63.02	\$81.00	\$63.02	\$85.02	\$0.00
Yes	Yes	32	\$75.00	\$75.00	\$75.00	\$76.00	\$5.95
No	No	40	\$89.99	\$89.99	\$90.99	\$100.00	\$5.95
No	No	596	\$20.50	\$40.00	\$22.50	\$38.50	\$7.50
No	Yes	122	\$19.99	\$19.99	\$19.99	\$22.99	\$4.00
No	Yes	26	\$152.50	\$155.40	\$157.50	\$160.00	\$4.00
No	No	31	\$10.50	\$25.51	\$20.50	\$25.00	\$4.50
No	Yes	34	\$40.00	\$40.00	\$40.00	\$45.00	\$12.00
No	Yes	36	\$62.00	\$91.00	\$61.00	\$97.00	\$5.00
No	Yes	21	\$66.67	\$66.67	\$56.56	\$65.67	\$5.00
No	No	21	\$36.01	\$40.55	\$34.00	\$45.01	\$5.00
No	No	104	\$189.52	\$189.52	\$174.50	\$189.79	\$8.00
Yes		0	\$30.03	\$30.03	\$31.00	\$36.00	\$9.99
No	No	169	\$31.00	\$50.00	\$0.99	\$20.51	\$10.00
No	No	62	\$6.00	\$7.50	\$6.50	\$10.00	\$5.00
No	Yes	233	\$0.99	\$0.99	\$0.99	\$10.04	\$10.00
No	Yes	58	\$50.00	\$56.55	\$50.00	\$58.00	\$5.00
Yes	Yes	29	\$19.99	\$19.99	\$19.99	\$37.49	\$15.00
Yes	Yes	43	\$19.99	\$51.00	\$23.01	\$50.01	\$15.00
No	No	514	\$41.00	\$41.00	\$32.00	\$41.00	\$10.00
Yes	No	283	\$40.00	\$40.00	\$30.00	\$31.00	\$10.00
Yes	Yes	900	\$60.00	\$60.00	\$51.00	\$62.01	\$7.00
No	Yes	84	\$80.00	\$80.00	\$81.00	\$82.50	\$10.00

Table 11

Summary statistics of Treatment and Control Auctions. Reported are summary statistics for the treatment and control auctions.

	Treatment Auctions ($N = 30$)	Control Auctions ($N = 30$)
Price 5 min before close		
Mean	\$65.17	\$67.00
Min	\$0.99	\$0.99
Median	\$50.50	\$55.00
Max	\$177.50	\$189.52
Closing price		
Mean	\$76.65	\$73.79
Min	\$10.00	\$0.99
Median	\$60.01	\$58.28
Max	\$189.79	\$189.52
Average number of bidders	5.6	4
Average number of bids	11.83	6.43

price run-up in the final five minutes of an auction.²⁸ In our entire sample of 30 auction pairs, the average *RUNUP* in the treatment group (\$11.48) is greater than it is in the control group (\$6.80). This difference is statistically significant under the paired *t*-test, the sign test, and the Wilcoxon signed-rank test ($p < .001$ for each). We obtain similar results when we restrict the sample to the 23 auctions in which we did not win the item. The average *RUNUP* in the treatment group is greater (\$11.87) than it is in the control group (\$7.52), and the differences are highly significant under all three tests ($p = .001, .003$, and $.001$).

We assume for this analysis that the sellers in our experiment are not shill bidding in these auctions. If they were and if they were successful in their bidding, the high bidder would pay his high bid and the expected difference in closing prices

²⁸ Since the prices of the items in the treatment and control groups are not identical 5 min prior to the auctions' close, this is not equivalent to comparing the closing prices of the auctions.

Table 12

DiscNStop bidding length and the frequency of success. We report the distribution of the number of bids in our 46 incremental bidding sequences. The “Successes” column lists the sequences in which we actually discovered the high bidder’s maximum bid without becoming the high bidder ourselves, whereas the “High Bid” column lists the attempts in which we became the high bidder before discovering the high bidder’s maximum bid. These two columns do not sum to the total number of attempts because 3 of our attempts were cut short by the auction’s ending time and 1 was interrupted by another bidder. The cumulative probability distributions are in parentheses.

Length	All ($N = 46$)	Successes ($N = 18$)	High bid ($N = 24$)
1	16 (0.35)	7 (0.39)	7 (0.29)
2	5 (0.46)	2 (0.50)	2 (0.38)
3	10 (0.67)	4 (0.72)	6 (0.63)
4	3 (0.74)	2 (0.83)	1 (0.67)
5	6 (0.87)	2 (0.94)	4 (0.83)
6	1 (0.89)	0 (0.94)	1 (0.88)
7	1 (0.91)	0 (0.94)	1 (0.92)
8	0 (0.91)	0 (0.94)	0 (0.92)
9	1 (0.93)	0 (0.94)	1 (0.96)
10	0 (0.93)	0 (0.94)	0 (0.96)
11	2 (0.98)	0 (0.94)	1 (1.00)
12	0 (0.98)	0 (0.94)	0 (1.00)
13	0 (0.98)	0 (0.94)	0 (1.00)
14	1 (1.00)	1 (1.00)	0 (1.00)
> 14	0 (1.00)	0 (1.00)	0 (1.00)

between the treatment and control groups would be less than the bid increment. Since we find that the average difference is greater than the bid increment, we are not concerned by the assumption.

The preceding analysis provides some evidence that sellers can increase the final price by following our shill bidding strategy. It did not, however, address the effectiveness of the individual DiscNStop bidding sequences. We engaged in 46 sequences of DiscNStop bids in our 30 treatment auctions.²⁹ We were able to determine the high bidder’s maximum bid without becoming the high bidder ourselves in 18 of the 46 DiscNStop attempts. We became the high bidder in 24 of our attempts. Three of our attempts were cut short by the time constraint,³⁰ and one attempt was interrupted by another bidder entering a bid exceeding the previous bidder’s maximum bid.

Table 12 shows that most of the DiscNStop sequences were short in duration: 16 of the 46 attempts consisted of only one bid, and 40 of them were five bids or fewer. Out of the 18 that were successful, 7 consisted of only one bid and 17 of them were five bids or fewer.

The columns labeled “Successes” and “High Bid” of Table 12 suggest that for all n , the likelihood of the n th bid becoming the high bid is approximately equal to the likelihood of the n th bid discovering the high bid without becoming the high bidder. This should be expected. To see this, consider the case in which the bid increment is \$0.50 and where the DiscNStop bidding strategy starts at \$3.00. Then it is easily seen that the strategy will cause the bidder to become the high bidder if and only if the current maximum bid is congruent to $x \bmod \$1$, where $x \in [50, 99]$ (e.g., if the current maximum bid is \$4.78), and it will be a success if and only if the current maximum bid is congruent to $x \bmod \$1$, where $x \in [0, 50]$ (e.g., if the current maximum bid is \$4.26).³¹ The same logic applies for starting prices other than \$3 and bid increments other than \$0.50.

4.4. Is shilling an optimal strategy for sellers?

When choosing whether or not to bid on one’s own auction, the seller must weigh the benefits of the shill – higher expected selling price – against the costs – having to locate another buyer to purchase the item if the seller’s shill account wins the auction. To get an estimate of how large this cost would have to be in our sample to make a risk-neutral seller indifferent between shilling and not, we solve for the c that satisfies the following equation:

$$c \times \Pr(\text{win auction}|\text{shill}) = \Pr(\text{lose auction}|\text{shill}) \times (\mathbf{E}[p|\text{shill, lose auction}] - \mathbf{E}[p|\text{do not shill, lose auction}]),$$

where $\Pr(\cdot)$ denotes probability and p represents the item’s closing price.

In our sample, we won 7 of the 30 auctions in which we bid. Hence, our estimates for $\Pr(\text{win auction} - \text{shill})$ and $\Pr(\text{lose auction} - \text{shill})$ are 7/30 and 23/30, respectively. Our estimates for $\mathbf{E}[p - \text{shill, lose auction}]$ and $\mathbf{E}[p - \text{do not shill, lose auction}]$ are simply the average closing prices of the treatment and control auctions, respectively, among the 23 pairs of auctions which we did not win. These numbers are \$85.52 and \$82.99, respectively. The cutoff cost, c , that solves this equation is then \$8.31. To put this value in perspective, the average closing price of the 7 control auctions for the items we

²⁹ Recall our bidding strategy (outlined in Section 4.2) has us repeating the DiscNStop strategy if a new member becomes the high bidder within the last five minutes of the auction. Because this happened 16 times, we engaged in a total of $30 + 16 = 46$ DiscNStop bidding sequences.

³⁰ These attempts followed snipers’ bids in the closing seconds of the auctions.

³¹ To see this, note that the proxy bidder for the maximum bidder will beat a bid by the DiscNStop bidder by less than the bid increment in the latter case but not in the former.

Table 13

Bid increment and the frequency of success. Here we report the number of times the high bidder's maximum bid was a multiple of \$10 (\$5, \$1). "All Discovered" refers to our bidding sequences in which the high bidder's maximum bid was discovered, which occurred for all but the three attempts in which we ran out of time. The "Become High Bidder" column refers to the attempts in which we became the high bidder.

High bid multiple of:	All Discovered (N = 43)	Become High Bidder (N = 24)
\$10	15	8
\$5	22	12
\$1	27	16

won was only \$43.57.³² The computed cutoff cost, \$8.31, which is 19.1 percent of the item's expected selling price conditional on winning the item, appears large when considering the seller's options should he win the auction: he may contact the second highest bidder to see if he's still interested in purchasing the item, or, if he does not like the highest bidder's bid, he may relist the item. Doing so would incur an additional transaction fee which is 8.75 percent of the first \$25 of the item's selling price and 3.5 percent of the remainder of its price.³³ The cutoff cost c appears high relative to these relisting costs which provides some evidence that shilling would have benefited the sellers in our sample. However, the relisting fees do not include some unobserved costs such as the cost of a suspended account (for a shiller who is caught) or the cost of a seller's time to relist the item. We cannot be certain that shilling is optimal in the face of such unobserved costs.

When deciding whether to shill or not, a seller should also optimize over all possible shilling strategies. In our bidding (and hence in the calculation for the cutoff cost, c), we followed a simple strategy of never submitting a bid above the lowest possible amount. We have reason to believe, however, that such a strategy is not optimal. Recall from Table 9 that bidders have a tendency to bid on common multiples, e.g., \$5, \$10, etc. Because of this habit, the bidding behavior of actual shill bidders would probably deviate from the simple strategy we defined in Section 4.2. For example, suppose the bid increment is \$0.50 and the current price is \$9.00. Then the DiscNStop strategy we employ prescribes the bidder to bid \$9.50, which is the minimum acceptable bid. A savvy shiller might instead bid \$9.51 due to the likelihood that the high bidder's maximum bid is \$10.00. In this scenario, the shill bidder who modified the DiscNStop strategy would have been able to discover the high bidder's maximum bid without becoming the high bidder, whereas one who followed the algorithm as described in Section 4.2 would have become the high bidder.³⁴

The information reported in Table 13 suggests such a modification of the strategy would have been effective in our sample: of the 24 DiscNStop sequences that resulted in us becoming the high bidder, the previous high bidder's maximum bid was a multiple of \$5 in 12 of them. Since the bid increment is not higher than \$2.50 for any prices in our sample, the modification described above would have been able to exploit the bidders' tendency to bid in multiples of \$5.00.

Recall that we do not compare the bidders' first bids in their incremental bidding sequences to the item's price at the start of the sequence in Section 3. Hence, our results in that section are robust to the possibility that shill bidders' first bids differ from the ones specified by our strategy in this section.

5. The effects of shilling on eBay's profits

eBay's auction policy explicitly prohibits shill bidding and, according to the NY Attorney General's Office, eBay has been helpful in providing data to investigate shill bidding in its auctions.³⁵ However, eBay's incentive to reduce shill bidding in their auctions is unclear. Dr. Hampton Finer, the economist who worked for the NY Attorney General's Office on the aforementioned cases, sums it up this way: "Obviously, shill bidding tends to move the prices up and eBay gets a fee that's based on price, then one might imagine that they would tacitly encourage shill bidding. At the same time, you don't want to discourage people from coming in and thinking that it's fair."³⁶

To estimate how shilling affects eBay's profits, we consider the equilibrium behavior of bidders in a second price quasi-sealed bid auction (to be described below) in which the seller is a shill bidder with probability p . Suppose there are n bidders in the auction, and suppose the bidders' valuations are distributed independently and identically from the uniform distribution over the unit interval $[0, 1]$. Let v_i denote i 's valuation for the good. Suppose that the seller is a competitive shill bidder with probability $p \in [0, 1]$; i.e., with probability p the seller will bid the price up to the highest bidder's bid. In this context, a shill bidder is "successful" in his shill attempt if and only if he does not become the high bidder. Let $q \in [0, 1]$ be the probability

³² This is quite small compared to the average closing costs of the auctions we did not win, \$82.99.

³³ For more on eBay's rules regarding second chance offers and their fees, see http://pages.ebay.com/help/sell/second_chance_offer.html and <http://pages.ebay.com/help/sell/fees.html>.

³⁴ To see this, notice that if the high bidder's maximum bid were in fact \$10.00, a bid of \$9.50 would cause the high bidder's proxy bidder to bid \$10.00. If the bidder continued the DiscNStop strategy he would bid \$10.50 and become the high bidder. If instead the shiller's first bid were \$9.51, the high bidder's proxy bidder would have beaten the shiller's bid by \$0.49, which is less than the bid increment, so the shiller would have discovered that the high bidder's maximum bid was in fact \$10.00 without becoming the high bidder.

³⁵ For eBay's policy on shill bidding see footnote 9. Comments about eBay's cooperation in the shill bidding cases can be found at <http://www.oag.state.ny.us/press/2004/nov/nov8a.04.html> and http://www.ftc.gov/be/workshops/internetauction/Auction.Transcript_public.pdf.

³⁶ See http://www.ftc.gov/be/workshops/internetauction/Auction.Transcript_public.pdf.

that a shill bidder is successful in his shill attempt. Note that if a shill bidder is *not* successful in his shill attempt, he becomes the high bidder and the price of the item becomes the amount of the highest non-shill bid. Finally, assume that the seller is able to respond to all bidders' bids, but that bidders are not able to respond to other bidders' bids.

Proposition 1. All players bidding αv_i , where $\alpha = (n-1)(pq+1-p)/npq + (1-p)(n-1)$, is a Bayesian Nash Equilibrium.

Proof. Let $\Gamma(\beta)$ be bidder i 's expected payoff from bidding βv_i given that the other bidders are bidding αv_j . Note that:

$$\begin{aligned}\Gamma(\beta) &= \Pr(\beta v_i > \alpha v_j \forall j \neq i) \times [pq(1-\beta)v_i + (1-p)(v_i - \mathbf{E}[\max\{\alpha v_j : j \neq i\} | \beta v_i > \alpha v_j \forall j \neq i])] \\ &= \left(\frac{\beta v_i}{\alpha}\right)^{n-1} \left[pq(1-\beta)v_i + (1-p)\left(v_i - \frac{n-1}{n}\beta v_i\right) \right] \mathbf{1}_{\{\beta v_i \leq \alpha\}} \\ &\quad + \left[pq(1-\beta)v_i + (1-p)\left(v_i - \frac{n-1}{n}\alpha\right) \right] \mathbf{1}_{\{\beta v_i > \alpha\}} \\ &= \left(\frac{v_i^n}{\alpha^{n-1}}\right) \left[\beta^{n-1}(pq+1-p) + \beta^n\left(-pq - \frac{n-1}{n} + p\frac{n-1}{n}\right) \right] \mathbf{1}_{\{\beta v_i \leq \alpha\}} \\ &\quad + \left[pq(1-\beta)v_i + (1-p)\left(v_i - \frac{n-1}{n}\alpha\right) \right] \mathbf{1}_{\{\beta v_i > \alpha\}}.\end{aligned}$$

Thus,

$$\Gamma'(\beta) = \left(\frac{v_i^n}{\alpha^{n-1}}\right) \left[(n-1)\beta^{n-2}(pq+1-p) + n\beta^{n-1}\left(-pq - \frac{n-1}{n} + p\frac{n-1}{n}\right) \right] \mathbf{1}_{\{\beta v_i \leq \alpha\}} - pqv_i \mathbf{1}_{\{\beta v_i > \alpha\}}.$$

Clearly, $\Gamma'(\beta) \leq 0$ for all $\beta > \alpha/v_i \geq \alpha$. Moreover, $\Gamma'(\beta) \geq 0$ for all $\beta \leq \alpha$, $\Gamma'(\beta) \leq 0$ for all $\beta \in [\alpha, \alpha/v_i]$, and $\Gamma'(\alpha) = 0$. Hence, bidding αv_i is a best response for bidder i . \square

Unlike second price auctions, bidding one's valuation in this environment is not a dominant strategy because the amount the seller bids depends on the maximum of the other bidders' bids. Because of this feature bidders bid below their valuations in equilibrium.

If $p = q = 1$, the auction is equivalent to a first price sealed bid auction, and $\alpha = n-1/n$; in this case, the above equilibrium is equivalent to the symmetric equilibrium of first price sealed bid auctions.

If $p = 0$, the auction is equivalent to a second price sealed bid auction, and $\alpha = 1$, which is the unique equilibrium in this environment.

By the revenue equivalence theorem, if $q = 1$, the seller's expected revenue equals $n-1/n+1$ for all $p \in [0, 1]$.³⁷ Because the fee eBay collects from its sellers is a function of the final price, eBay does not benefit from shill bidding if shilling is perfectly efficient.

To address the $q < 1$ cases, notice that

$$\frac{\partial \alpha(n, p, q)}{\partial q} = \frac{(n-1)(p-1)p}{[npq + (1-p)(n-1)]^2} < 0 \quad \text{for all } (p, q) \in (0, 1) \times [0, 1].$$

Clearly, for fixed p (shilling probability), the expected final price of the good is increasing in α . For any fixed shilling probability $p \in (0, 1)$, the above inequality shows that eBay's profits are higher with parameter values (p, q) , $q < 1$, than with $(p, 1)$. Earlier, we argued that the expected final price with $q = 1$ (perfectly efficient shilling) is the same as the expected price with $p = 0$ (no shilling). Hence, this model suggests that eBay's profits are higher in the presence of shill bidding if and only if the shilling is not perfectly efficient ($q < 1$).³⁸

6. Conclusion

We have provided evidence of shill bidding via a strategy we call Discover-and-Stop bidding. Our evidence came from three independent analyses: (1) a probit model which demonstrates that (Bidder, Seller) pairs with intimate relationships are more likely to participate in Discover-and-Stop bidding, (2) an observation that incremental bidding (a fundamental property of the Discover-and-Stop strategy) is more popular in auctions where bidders have hidden identities and (3) distributional analysis of bidding sequences of Discover-and-Stop bidders. We also argued that shill bidders can implement a Discover-and-Stop bidding strategy due to eBay's auction design and the predictability of bidders' bids. In our experimental section,

³⁷ This can also be seen by a direct calculation.

³⁸ Some readers might question the assumption that bidders play equilibrium strategies and/or that bidders are even aware that some sellers are shillers. If bidders were unaware of shilling, then they would likely bid their entire valuation, and eBay would benefit even more from the presence of shill bidding. The key point is that it appears eBay has little incentive to prevent shill bidding despite the fact that in equilibrium bidders bid below their valuations as a response to shilling.

we participated in real eBay auctions to demonstrate the potential profitability as well as the risk of overbidding from a Discover-and-Stop strategy.

Given our evidence of Discover-and-Stop shill bidding, a natural question is how the auction mechanism can be changed to reduce its prevalence. One change eBay could make to reduce shilling involves its bid increments. Recall that under eBay's current mechanism, the proxy bidder of a current high bidder H bids $\min\{b_H, b + \varepsilon\}$ whenever another bidder places an acceptable bid b , where b_H equals H 's maximum bid and ε is the bid increment.³⁹ Recall that this spread between the new bidder's bid, b , and the resulting price, $\min\{b_H, b + \varepsilon\}$, is what allows shill bidders to drive the price of the item to b_H without becoming the high bidder. If eBay were to change its proxy bidding system to one in which the current high bidder's proxy bidder *matches* any bid $b \leq b_H$ (so that the resulting price is b , rather than $\min\{b_H, b + \varepsilon\}$), the seller would have a more difficult time moving the price to b_H without becoming the high bidder. Hence, shilling would be less profitable for sellers, so shilling would likely decrease.⁴⁰ Another way eBay could deter this – and other forms of shill bidding – is to make sellers unable to respond to bidders' bids. For example, eBay could offer bidders the option of having their bids processed at the close of the auction. As a result, sellers would have fewer opportunities (and less incentive) to shill, which would likely result in less shilling.⁴¹ While these are plausible changes to the auction mechanism, the results of Section 5 question whether eBay actually has an incentive to reduce the final price of a good sold in its exchange if its fee is tied to that price.

Acknowledgements

We thank the associate editor for suggesting the natural experiment implemented in Section 4. We also thank two anonymous referees, Mike Emmet, Grant Farnsworth, Mike Fishman, Andrew Hertzberg, Jörg Oechssler, Mitchell Petersen, Rob Porter, Jim Schummer, Annette Vissing-Jorgensen, Mike Whinston, and the Northwestern IO seminar participants for their helpful comments and suggestions.

References

- Ball, C., Torus, W., Tschoegl, A., 1985. The degree of price resolution: the case of the gold market. *Journal of Futures Markets* 5, 29–43.
- Christie, W., Schultz, P., 1994. Why do NASDAQ market makers avoid odd-eighth quotes? *Journal of Finance* 49, 1813–1840.
- Goodhart, C., Curcio, R., 1991. The clustering of bid/ask prices and the spread in the foreign exchange market. Discussion Paper No. 110, LSE Financial Markets Group Discussion Paper.
- Grossman, S., Miller, M., Cone, K., Fischel, D., Ross, D., 1997. Clustering and competition in asset markets. *Journal of Law and Economics* 40, 23–60.
- Harris, L., 1991. Price clustering and discreteness. *Review of Financial Studies* 4, 389–415.
- Hornick, Cherian, Zakay, 1994. The influence of prototypic values on the validity of studies using time estimates. *Journal of Market Research Society* 36, 145–147.
- Ikenberry, D., Weston, J., 2008. Clustering in US stock prices after decimalization. *European Financial Management* 14, 30–54.
- Kauffmann, R., Wood, C., 2005. The effects of shilling on final bid prices in online auctions. *Electronic Commerce Research and Applications* 4, 21–34.
- Ockenfels, A., Reiley, D.H., Sadrieh, A., Online auctions. In: Hendershott, T. (Ed.), *Handbook of Information Systems and Economics*. Elsevier Science, New York, 2007, pp. 571–628.
- Roth, A.E., Ockenfels, A., 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.
- The Committee on Energy and Commerce, 2001. Tauzin, Wilson want details on shilling and online auction fraud. News Release, 26 June 2001. Available at: http://energycommerce.house.gov/107/News/06262001_304print.htm.

³⁹ By “acceptable” bid, we mean a bid that is greater than or equal to the current price plus the bid increment.

⁴⁰ The seller could still shill effectively by taking advantage of eBay's bid retraction option. In particular, eBay allows bidders to quickly retract a bid if they “mistakenly” add an extra zero to their bid; so by “mistakenly” placing a shill bid that is ten times larger than what he “intended” to bid, a seller could determine the current high bid, retract that shill bid, and use a different shill account to submit a bid equal to the current high bid. Hence, eBay might need to eliminate its bid retraction option, too, in order to prevent shilling.

⁴¹ This change would also likely improve social welfare. Sniping, or last second bidding, is a strategy that is commonly used by bidders. The problem is that not all snipe attempts are processed: Roth and Ockenfels (2002) surveyed late bidders and reported that “more than 80 percent of the bidders who successfully bid at least once in the last minute of an eBay auction replied that it happened at least once to them that they started to make a bid, but the auction was closed before the bid was received.” Esnipe.com, an automated sniper, stopped reporting the frequency at which their bids are successfully submitted, but in September 2000 they reported that 4.5 percent of their bids failed to be successfully transmitted. Clearly when bids are not processed the good may not be sold to the person with the highest valuation, resulting in a welfare loss.