

Selling a Dollar for More Than a Dollar? Evidence from Online Penny Auctions*

Zhongmin Wang[†] Minbo Xu[‡]

Abstract

Online penny auctions, emerged recently, are seen as an adaptation of the famous dollar auction and as “the evil stepchild of game theory and behavioral economics.” We use the complete bid and bidder history at such a website to study if penny auctions can sustain excessive profits over time. The website we study is characterized by a revolving door of new bidders who lose money. A very small percentage of bidders are experienced and strategically sophisticated, but they earn substantial profits. Our evidence thus suggests that penny auctions cannot sustain excessive profits without attracting a revolving door of new customers.

Keywords: behavioral game theory, behavioral industrial organization, auction, learning, strategic sophistication

JEL Classification: D03, D44, L81

*A previous version of the paper was circulated under the title "Learning and Strategic Sophistication in Games: Evidence from Online Penny Auctions." We are grateful to Christoph Bauner, Christian Rojas, and especially John Conlon, Jim Dana, and Ian Gale for their comments and suggestions that improved the paper. We thank seminar participants at College of William & Mary, University of Mississippi, University of Wisconsin Whitewater, and the 10th International Industrial Organization Conference for helpful discussions. Chengyan Gu and Xi Shen provided excellent research assistance. Any errors are ours only.

[†]Corresponding author. Resources for the Future, 1616 P St. NW, Washington, DC 20036, wang@rff.org, 202-328-5036

[‡]Department of Economics, Georgetown University, Washington, DC 20057, mx7@georgetown.edu.

1 Introduction

Martin Shubik’s (1971) famous dollar auction suggests the possibility of selling a dollar for more than a dollar. Overbidding may occur due to such reasons as the sunk cost fallacy or bidding fever. Can a firm adapt the dollar auction into a selling mechanism that sustains excessive profits over time? A new auction format recently emerged on the Internet, called the penny auction, might be seen as such an attempt. Penny auctions, also known as pay-to-bid auctions, were described by Richard Thaler in the *New York Times* as a “diabolically inventive” adaptation of the dollar auction.¹ An article in the *Washington Post* claims that penny auction is “the evil stepchild of game theory and behavioral economics” because it “fiendishly plays on every irrational impulse buyers have.”² In this paper, we use the complete bid and bidder history at a major penny auction website to study if penny auctions can sustain excessive profits over time.

We find that the website is characterized by a revolving door of new bidders: the overwhelming majority of new bidders who join the website on a given day play in only a few auctions, place a small number of bids, lose some money, and then permanently leave the site within a week or so. A very small percentage of bidders are experienced and strategically sophisticated, but they win most of the auctions and earn substantial profits from the website. Our evidence thus suggests that penny auction websites cannot sustain excessive profits without attracting a revolving door of new customers who will lose money.³ This main conclusion of our paper is strongly supported by a subsequent independent lab study of penny auctions (Caldara 2012).

Unlike eBay, penny auction websites sell products themselves, using rules similar to the following. First, a bidder must pay a small non-refundable fee (e.g., \$0.75) to place a bid. A bid is an offer to buy the product at the current

¹Richard H. Thaler, “Paying a Price for the Thrill of the Hunt,” *New York Times*, November 15, 2009.

²Mark Gimein, “The Big Money: The Pennies Add Up at Swoopo.com,” *Washington Post*, July 12, 2009.

³This feature is shared somewhat by Ponzi schemes. We are not claiming that penny auctions are Ponzi schemes or necessarily scams.

auction price. The auction price for any product is initially 0 and is increased by a fixed amount whenever a bid is placed. The increment is typically one penny, thus the name of penny auction. Second, the winner is the *last* bidder, the person whose bid is not followed by any other bid before a timer (e.g., of 30 seconds) expires. The timer is reset whenever a new bid is placed. The auction winner receives the product and pays the auction price. Consider an example in our dataset. A bidder won an iPad auction after placing 70 bids, and the auction price was \$64.97. The winner paid a total cost of \$117.47 ($= 70 \times 0.75 + 64.97$) for the iPad, and the website's revenue was \$4,937.72 ($= 6,497 \times 0.75 + 64.97$)! A penny auction thus combines elements of an all-pay auction with a series of lotteries.⁴

Our evidence comes from a nearly ideal bid-level dataset collected from a major penny auction website (BigDeal.com). The dataset covers all of the over 22 million bids placed by more than 200,000 bidders in over 100,000 auctions for a period of over 20 months, starting from the website's first day of operation to two days before the site's closure. The dataset records the complete bid history of each bidder as well as the precise timing of each bid. We use a product's retail price at Amazon as an estimate of the product's market value. We define the auctioneer's excessive profit as its revenue minus the market value of the products sold. Similarly, we define a bidder's profit or loss as the market value of the products she won minus her cost of bidding. Given these definitions, our conclusion that penny auctions cannot sustain excessive profits does not mean that penny auctions cannot sustain normal profits.

In this paper, we also address the question of whether experienced bidders are heterogeneous. We find strong evidence that experienced bidders differ in their strategic sophistication and learning. Highly sophisticated experienced bidders make significant amount of profits, but unsophisticated ones lose sig-

⁴A penny auction is not a standard auction in which the bidder who bids the most wins (Krishna 2002, p. 29). The winner of a penny auction is often not the bidder who places the most bids. Another nonstandard auction format is the lowest unique bid auction (e.g., Raviv and Virag 2009; Houba et al. 2011) or the lowest unique positive integer game (e.g., Ostling et al. 2011). A penny auction is clearly very different from eBay auctions. See Bajari and Hortaçsu (2004) for a review of the literature on online auctions, and Einav et al. (2011) for a recent example.

nificant amount of money. Sophisticated bidders learn to play better, but unsophisticated ones do not.

Our paper relates to the behavioral industrial organization literature that focuses on how profit-maximizing firms exploit consumer biases. See sections of Ellison (2006) and DellaVigna (2009) for reviews of the literature.⁵ See Malmendier and Lee (2011) and the references therein for empirical studies of overbidding in auctions. Our finding of a revolving door of new bidders reflects the simple logic of individual rationality: no matter how effective the penny auction might be in exploiting bidder biases, it offers bidders immediate outcome (win or lose) feedback so that losing bidders can quickly learn to stop participating. Our finding thus suggests that learning can limit overbidding, at least in auctions with clear feedback, and that firms' ability to exploit consumer biases is constrained by consumer learning.⁶

Our paper also contributes to the behavioral game theory literature, which finds that subjects' behavior in experimental games often deviates from equilibrium because of limited strategic sophistication or lack of prior experience/learning (e.g., Camerer 2003; Crawford et al. 2010). Our results provide field evidence for Crawford et al.'s (2010, p. 28) observation that strategic sophistication "is heterogeneous, . . . so that no model that imposes homogeneity . . . will do full justice to [players'] behavior." Our paper adds to an emerging literature that uses the behavioral game theory approach to study strategic interactions in field settings. Brown et al. (2012) study the implications of consumers' limited strategic thinking in the movie industry. Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) find managers' strategic sophistication affects firms' performance. Both papers measure managers' strategic sophistication by the number of iterations of best response they perform in selecting an action in a static game, as in level-k/cognitive hierarchy models (e.g., Camerer et al. 2004; Costa-Gomes and Crawford, 2006). We measure an experienced bidder's lack of strategic sophistication by the frequency with

⁵DellaVigna and Malmendier (2006) is an excellent example of empirical behavioral industrial organization study.

⁶See List (2003) for evidence that market experiences may eliminate some forms of market anomalies.

which she places a bid in the middle of the timer. Bids in the middle of the timer, we shall argue, are inferior to last-seconds bids.

Four papers on penny auctions (Augenblick 2011; Platt et al. 2010; Hinnoosaar 2010; and Byers et al. 2010) appeared before our paper. All four papers use data from Swoopo, the first penny auction website, and find that the website made excessive profit during some periods. Platt et al. (2010) study auction-level data and attribute auctioneer profits to bidders' risk-loving preference. Byers et al. (2010) focus on bidder asymmetry as an explanation for why Swoopo made excessive profits. Hinnoosaar (2010) deals largely with a technical issue in modeling an individual penny auction. Augenblick (2011, p. 2) is the only paper that addresses bidder learning, but he concludes that overbidding at Swoopo is consistent with "a *naive sunk cost fallacy* ... Surprisingly, profiting off of this behavioral tendency appears to be a sustainable business strategy...consumer learning occurs but is extremely slow, allowing the auctioneer to profit during the learning process." (emphasis in original). Our finding of a revolving door of new bidders indicates that penny auctions may generate excessive profits in the short run, but not in the long run. Indeed, Swoop, BigDeal, and many other penny auction websites have closed. Because of data constraints, the previous literature did not make the critical observation of a revolving door of new bidders. Augenblick (2011) does note that most of the bidders in this sample play in a small number of auctions and place a small number of bids. However, Augenblick does not study the timing of bidder entry and exit, which is a critical part of our observation of a *revolving* door of *new* bidders, a statement largely about bidder dynamics over time. Augenblick does not identify the time when a bidder first enters Swoopo or permanently leaves Swoopo perhaps because his bid-level data covers only a period of about four months.

Two papers on penny auctions (Caldara 2012; Goodman 2012) appeared after our paper. Caldara (2012) is the first to conduct lab experiments to study penny auctions. He writes (p. 6) that his lab findings support "the [conclusion] of Wang and Xu (2011) that pay-to-bid auction websites profit from a 'revolving door of new bidders,'" and he also concludes (p. 32) that

“excessive revenues will only last as long as pay-to-bid auction websites can attract new, inexperienced bidders.” Penny auctions are also known as *pay-to-bid* or *bidding fee* auctions. His findings also strongly support our measure of bidder strategic sophistication: proportion of middle middles is strongly correlated with bidder performance in the lab. Goodman (2012) is similar in some aspects to Augenblick (2011) but focuses on the role of reputation in penny auctions.

In terms of bidder learning, our findings are very different from that in Augenblick (2011). Our results suggest that bidders are heterogeneous in strategic sophistication and learning. Augenblick, on the other hand, considers all bidders in his sample together in his learning regressions, presuming that all bidders have the same learning function. Augenblick does not measure a bidder’s strategic sophistication; instead, he attempts to measure the sophistication of individual bids, irrespective of the bidder who places the bids. If our learning regression includes all bidders, we would also find extremely slow learning: bidders in our sample, on average, do not start to earn a positive profit until they have already played nearly 200 auctions. However, this is a spurious finding, resulting from the selection bias. The sophisticated and experienced bidders in our sample start to make positive profits from their first few auctions, but their behavior is lost in a learning regression that considers all bidders together: the overwhelming majority of bidders are inexperienced and lose money. Only when the experience variable is large enough then the learning regressions reflect the behavior of the experienced bidders.

The remainder of this paper proceeds as follows. Section 2 describes the penny auction industry, the auction rules, and the data. Section 3 provides some theoretical considerations. Section 4 presents our empirical results. Section 5 concludes.

2 Background, Auction Rules, and Data

2.1 The Penny Auction Industry

The first penny auction firm, Swoopo, was founded in Germany in 2005, and it started its U.S. website in 2008. By November 2010, at least 125 penny auction websites targeting U.S. consumers were being monitored by Compete.com, a web traffic monitoring company. The total number of unique monthly visitors to these penny auction websites reached 25.1% of that to eBay in November 2010, but has since declined sharply. Table 1 lists the 11 websites whose traffic was ranked in the top 5 of all penny auction sites for any two consecutive months from February 2010 through April 2011. We emphasize that among the 9 sites in Table 1 that were in existence in February 2010, 3 were closed in 2011, 2 barely attracted any visitors in October 2011, 1 was closed in 2012 (Bidrivals), and the other 3 sites experienced a dramatic traffic decline in 2011. Most penny auction websites attract little traffic and do not last for long.

Penny auctions are highly controversial. The Better Business Bureau (BBB) has received many consumer complaints against penny auction websites.⁷ In fact, it named penny auctions one of the top 10 scams of 2011.⁸ Three sites in Table 1 (i.e., Bidsauce, Swoopo, and Wavee) have an F rating, the worst BBB rating. Lawsuits have been filed against various penny auction websites, claiming penny auctions are a form of gambling. The industry brands itself as an *entertainment shopping* industry. Penny auction websites advertise that auction winners obtain products at deep discounts. It has been reported that penny auction sites “have driven up the price of advertising keywords on *Google* such as ‘cheap iPad.’ Buying keywords on search sites is the primary way the auction sites advertise products for sale.”⁹

Nearly all penny auction websites have two additional salient rules: win limits and a Buy-It-Now (BIN) option. Win limits restrict the number of auc-

⁷“Online Penny Auctions: Friend or Foe?” <http://www.bbb.org/blog/2010/10/online-penny-auctions-friend-or-foe/>.

⁸<http://www.bbb.org/us/article/bbb-names-top-ten-scams-of-2011-31711>.

⁹Brad Stone, “Penny Auction Sites Hurt by Glut of Competitors”, *Bloomberg Businessweek*, August 12, 2010.

Table 1: Monthly Traffic on the Largest Penny Auction Websites

Website	Number of unique visitors				BIN	Win limit
	Feb. 2010	Nov. 2010	Apr. 2011	Oct. 2011		
BigDeal.com	480,230	1,324,947	943,327	Closed	Yes	Yes
Bidcactus.com	1,428,316	3,411,705	1,979,846	740,981	Yes	Yes
Beezid.com	1,110,859	755,917	549,908	432,352	Yes	Yes
Bidsauce.com	356,811	690,014	344,514	9,052	Yes	Yes
Swoopo.com	286,142	171,141	Closed	Closed	Yes	Yes
Quibids.com	173,142	4,541,783	4,586,523	2,638,490	Yes	Yes
Bidrivals.com	63,329	419,945	490,751	144,468	Yes	Yes
Wavee.com	26,863	1,696,803	62,214	Closed	Yes	?
Bidhere.com	17,359	542,079	750,175	3,731	Yes	Yes
Zbidly.com	0	0	945,149	1,772,935	Yes	Yes
Biggerbidder.net	0	0	120,078	664,636	No	No
Total number of sites	47	125	158	116		
All sites	4,710,541	16,866,475	12,524,625	9,234,509		
eBay.com	64,766,668	67,197,011	69,929,590	77,232,991		
% of eBay traffic	7.3%	25.1%	17.9%	12.0%		

Notes: The 11 websites shown in this table include all the penny auction sites whose traffic was ranked in the top 5 of all penny auction sites in any two consecutive months from February 2010 through April 2011. We obtained the traffic data from Compete.com, and the Buy-It-Now (BIN) and win limit information from each individual penny auction website. For websites that still exist, the BIN and win limit information is as of March 2013.

tions a bidder can win. An individual bidder at BigDeal, for example, was restricted to at most 10 wins during a 30-day period. Once a bidder reached the win limit, she was prohibited from bidding in any auction until the 30-day period expires. Some websites impose much more stringent win limits. For example, bidders at Zbidly.com, a relatively new entrant, are allowed to win only one product with a retail price of \$999 or higher during a 28-day period and to win only one product with a retail price of \$499 or higher during a 7-day period.

The BIN option in penny auctions works differently from that found on eBay. A bidder who exercises the BIN option in penny auctions does not stop the auction. Instead, she stops her own bidding and obtains a product that is the same as the one under auction by paying the difference between the posted

retail price for the product and the cost of her bids. Penny auction websites post a retail price for any product to be auctioned. For example, the posted retail price for an iPad auction with the BIN option in our dataset is \$899.99. A losing bidder in this auction placed 1,067 bids, so her cost of bids is \$800.25 ($= 1,067 \times 0.75$). This bidder only needs to pay \$99.74 ($= 899.99 - 800.25$) more to exercise the BIN option and obtain an iPad that is the same as the one being auctioned. With the BIN option, this bidder pays the posted retail price of \$899.99 to buy an iPad. Without the BIN option, this bidder would have paid \$800.25 for nothing. The BIN option allows losing bidders who placed a large number of bids to recover some of their costs, which has the effect of reducing the profitability of penny auction websites. On the other hand, by eliminating the risk of losing a large amount of bids, the BIN option may allow a website to attract more bidders, which is perhaps why almost all penny auction websites now offer the BIN option.

2.2 BigDeal

BigDeal was one of the largest penny auction websites and appeared to be a serious business endeavor. It received \$4.5 million initial funding from well-known venture capital firms.¹⁰ It posted on its website photos and biographies of its management team and board members. BigDeal had a BBB rating of A-. Perhaps to mitigate potential concerns of shill bidding, BigDeal displayed the bid history of all live and past auctions on its website. Bidders could easily see the bid history of live and recently finished auctions, but it was time-consuming to see the bid history of auctions finished more than a few days earlier.¹¹

The rules of BigDeal auctions were representative of all penny auctions. Prior to bidding in any auction, bidders had to buy packs of bid tokens. Each

¹⁰Brad Stone, “BigDeal Puts a New Spin on ‘Entertainment Shopping’,” *New York Times* Bits Blog, December 19, 2009.

¹¹BigDeal created a separate web page for each auction that contained the general information and bid history of the auction. By clicking link buttons on the homepage or the “winner page” of BigDeal, one could have access to such web pages. It required increasingly larger numbers of clicks to access web pages of auctions finished earlier.

bid token cost \$0.75. The auction price for any product started at \$0, and each bid cost a single nonrefundable token and raised the auction price by a fixed increment. The price increment was \$0.01 in most auctions, and was \$0.05 or \$0.15 in a large number of auctions in the early part of our sample.

BigDeal typically released an auction with an initial countdown clock that last for 36 hours. If a bid was placed when more than 30 seconds were left on the initial countdown clock, the clock continued to run down. If a bid was placed when less than 30 seconds were left, however, the timer would always be extended by 30 seconds. A bidder won only if her bid was not followed by any other bid when the 30-second timer expired. It is not surprising that nearly all bids were placed after the 30-second timer started. Once the 30-second timer started, the timer was set to last 30 seconds *ex ante*, but whenever a bid was placed within this period, this period ended immediately and a new period started. Hence, the length of a time period *ex post* could range from 0 to 30 seconds.

In addition to her bidding cost, the winner also paid the auction price to attain the product. BigDeal offered losing bidders the BIN option in all auctions except for some bid pack and iPad auctions. BigDeal offered bidders a bid agent (called BidBuddy) that placed bids automatically on their behalf. The bid agent did not bid strategically. A bidder could impose three restrictions on her bid agent: the maximum number of bids, at what auction price to start to bid, and at what auction price to stop. A bidder could also deactivate a bid agent at any time. BigDeal auctioned several categories of products, including packs of bid tokens, video games and consoles, Apple products, non-Apple electronics such as computers, TVs, phones, cameras, and GPS, housewares, gift cards, handbags, jewelry, and movies.

2.3 Data

Our dataset, downloaded from BigDeal.com, covers the general information and the bidding history of all auctions released by BigDeal from November 19, 2009, the first day of the website's operation, through August 6, 2011, two days

before the website was closed. Auction-level information includes the auction price increment, the posted retail price, product name and description, the final auction price, the winner, and whether the BIN option was available. We do not observe which losing bidder(s) exercised the BIN option. The BIN option was not available for bid pack auctions until late November 2010, and it was also not available for iPad auctions for some periods “due to inventory restrictions.”

Another auction-level variable is whether an auction was a beginner auction that only accepted bids from new members. Most beginner auctions featured 10-token or 20-token bid packs. Beginner auctions were not offered until November 30, 2010.

The bid history for each auction includes every single bid: the exact second when a bid was placed, the screen name of the bidder, and whether the bid was placed manually or by a bid agent.

Figure 1: Daily Number of Non-beginner Auctions




Figure 1 shows the number of regular (non-beginner) auctions ended each day for the entire sample period. There was a dramatic decline in the number of auctions per day in late April 2011, which was a sign that BigDeal was preparing to shut down. Because the operation of BigDeal was no longer normal after that, we do not consider the auctions ended on or after May 1, 2011. For the sample period of November 19, 2009, through April 30, 2011, BigDeal offered a total of 110,703 auctions, including 78,634 regular auctions and 32,069 beginner auctions. Among these auctions, 61 regular auctions and

3,423 beginner auctions failed to attract a single bidder. A total of 207,069 bidders placed at least one bid during our sample period, and together they placed a total of 22,598,036 bids.

2.4 The Bidder with the Most Bids Often Does Not Win

Since the winner of a penny auction is the bidder who bids last, the bidder with the most bids in a penny auction often does not win the auction. The winner’s total number of bids is strictly smaller than that of at least one losing bidder in 40.9% of the 77,944 regular auctions with two bidders or more, and is equal to the maximum number of bids by any losing bidder in 12.9% of the auctions. Hence, the winner has the (strictly) largest number of bids in less than half of the regular auctions. In fact, in 3,302 auctions, the total number of bids placed by the last bidder is less than 10% of that by another bidder. In 154 auctions, the total number of bids placed by the last bidder is less than 1% of that by another bidder. The winners of such auctions often are “jumpers” in that they used the strategy of jumping in: starting to bid in an auction only after a large number of bids had already been placed in the auction.

3 Theoretical Considerations

To address the question of whether penny auctions can sustain excessive profits over time, we focus on bidder behavior across auctions instead of bidder behavior within individual auctions. Models focusing on an individual auction presumably predict that a dollar can only sell for a dollar, if all bidders are fully informed, rational, and risk-neutral or risk-averse.¹² Such models may

¹²Augenblick (2011) and Platt et al. (2010) present an equilibrium model of a single auction that predicts the zero-profit result. By assuming all bidders are homogeneous, fully informed, and rational; the number of bidders is fixed and known; the BIN option is not present; and the timing of placing a bid within a period can be ignored, the model can be solved by backward induction and is characterized by a mixed strategy equilibrium in which bidders’ expected value of placing a bid equals the cost of the bid so that they are indifferent between bidding and not bidding. If there are two bidders or more, the expected revenue

generate the result of selling a dollar for more than a dollar, if bidders suffer from behavioral biases¹³ or are risk loving, but bidder learning across auctions constrains the auctioneer’s ability to exploit bidders’ behavioral biases.

After playing in at least one auction, a bidder needs to decide whether to participate in another auction. This is a simple binary choice, and bidders are given accurate and immediate feedback on their gains or losses in the auctions in which they have played. According to Tversky and Kahneman (1986, p. S274), “accurate and immediate feedback about the relation between the situational conditions and the appropriate response” is conducive to effective learning. We then expect the principle of individual rationality to hold for all bidders with regard to the decision of whether to bid in another auction.

Suppose bidders are risk-neutral or risk-averse. Under this assumption, bidders quit the website if they lose enough to form a negative expected gain.

Suppose some bidders’ preferences are similar to those of lottery players (Platt et al. 2010).¹⁴ Under this assumption, an auctioneer may obtain excessive profits from experienced bidders who continue to play even if they lose money.

Therefore, sustained excessive profits may come from inexperienced bidders who have not learned the consequences of playing penny auctions or experienced bidders with gamblers’ preferences. Subsections 4.1 and 4.2 present evidence that BigDeal is characterized by a revolving door of new bidders and that the auctioneer profits from the revolving door of new bidders but loses

for the auctioneer is the value of the product since all bidders’ expected gain from bidding is zero in equilibrium. The BIN option complicates any attempt to build equilibrium models of an individual penny auction, but it does not affect our argument on bidder learning across auctions.

¹³Byers et al. (2010) present a model in which overbidding occurs if bidders underestimate the true number of bidders in the auction, and Augenblick (2011) sketches a model in which the sunk cost fallacy leads to overbidding.

¹⁴Chance plays an important role in determining the outcome of penny auctions. Penny auction bidders, however, are unlikely to have the Friedman and Savage (1948) utility function that is concave at the current wealth level and convex above it. The maximum return in penny auctions is relatively small; no product auctioned at BigDeal had a retail price over \$3,000. However, Golec and Tamarkin (1998) present evidence that horse track bettors seek skewness in return, not risk. It is also possible that some bidders may derive intrinsic utility from the mere act of bidding in penny auctions.

money to experienced bidders as a group.

Do experienced bidders exhibit heterogeneity in strategic sophistication? A major finding of the behavioral game theory literature is that subjects in experimental games exhibit heterogeneity in strategic sophistication. Penny auctions are much more complicated than experimental games, and its all-pay nature provides all players with strong incentives to perform. Hence, we expect the lab finding of bidder heterogeneity in strategic sophistication to extend to the field setting of penny auctions. That is, we hypothesize that experienced bidders in penny auctions differ in strategic sophistication and that bidders' performance and learning function depend on strategic sophistication.

Testing these two hypotheses raises the challenge of measuring bidders' strategic sophistication. The ideal measure of strategic sophistication should be based on bidders' exogenous characteristics, but unfortunately, we do not observe bidders' true identity. Similar to much of the behavioral game theory literature, we are thus forced to use players' behavior in the game as our measure. We measure a bidder's lack of strategic sophistication by the frequency with which she places a bid in the middle of the 30-second time clock. Call such bids "middle bids." To understand our justifications, recall that strategic sophistication, according to Crawford (1997, p. 209), "refers to the extent to which a player's beliefs and behavior reflect his analysis of the environment as a game rather than a decision problem, taking other players' incentives and the structure into account." A basic requirement for being strategically sophisticated is to analyze the strategic environment and respond accordingly. We argue that bids at the beginning or at the end of the time clock satisfy this requirement, but middle bids do not.

Last-second bids may be a best response in certain strategic environment in that they allow a bidder to learn about her competitors. A sophisticated bidder analyzes the bidding environment to learn who are competing with her and what strategies her competitors are using so that she may respond optimally. If a player bids in the middle of the time clock, she loses the chance to observe if any other bidder may place a bid between her bid and the end of the time period. If she waits for the last second to bid, she can observe

if someone else bids before then and she can always plan to bid at the last second of the following period. By bidding this way, she saves bids, keeps the auction alive, and obtains more information about who are competing with her and what strategies her competitors are using. We are making an informal argument that middle bids are inferior to last-second bids. To prove formally that last-second bids weakly dominate middle bids, one needs to impose strong assumptions because penny auctions are a complicated dynamic game in which players' actual strategy space is large and difficult to specify. However, is there a reasonable strategic situation in which middle bids are clearly a better response than last-second bids? We fail to find one, presuming that the opportunity cost of waiting for a few more seconds is negligible.

Bids in the first few seconds, commonly called aggressive bids, may be a best response in certain strategic environment as well. Indeed, many bidders often place a bid immediately after a competing bid and do so repeatedly for some periods (e.g., Augenblick 2011 and Goldman 2012). Aggressive bids may intimidate some unsophisticated bidders. Whether effective or not, aggressive bids are at least an act with a strategic intention. Middle bids are not.

We thus measure a bidder's lack of strategic sophistication by her proportion of middle bids. A smaller proportion of middle bids reflects a stronger degree of strategic sophistication. This is equivalent to stating that a larger proportion of aggressive bids and last-second bids together reflects a stronger degree of strategic sophistication. A larger proportion of aggressive bids alone (or last-second bids alone) may not reflect a stronger degree of strategic sophistication. For example, it is not always a good strategy to bid aggressively; there are reasonable strategic situations in which a bidder needs to use last-second bids to learn her competitors.

As mentioned in the introduction, Caldara (2012) adopts our measure of strategic sophistication in his lab experiments and finds that our measure is strongly correlated with bidder performance in the lab.

4 Empirical Analysis

Subsections 4.1 and 4.2 present our primary findings that suggest penny auctions cannot sustain excessive profits. Subsections 4.3 and 4.4 present our findings on bidder heterogeneity in strategic sophistication. Subsection 4.5 studies the impact of the BIN option on auctioneer profit.

4.1 A Revolving Door of New Bidders

In this subsection, we present compelling evidence that BigDeal was characterized by a revolving door of new bidders. A vast majority of new bidders who joined BigDeal on a given day played in only a few auctions, placed a small number of bids, and then quit the site within a week or so without winning any regular (i.e., non-beginner) auctions. This finding is consistent with our hypothesis that learning across auctions constrains the auctioneer’s ability to exploit bidders. A very small percentage of bidders were persistent participants, but they won most of the regular auctions. These findings confirm that bidders are heterogeneous and suggest that the revolving door of new bidders is a major source of profit for the auctioneer.

Table 2: Distribution of Three Measures of Bidder Participation Intensity

	Percentiles							
	50%	75%	90%	95%	99%	99.5%	99.75%	99.95%
Number of auctions	3	8	16	25	76	128	201	422
Number of bids	22	55	150	300	1,350	2,622	4,954	16,928
Duration	1	4	29	84	258	319	364	430

Table 2 shows the distribution of three measures of bidder participation: the number of auctions a bidder participated in, the number of bids submitted, and importantly, the duration of a bidder. We define the duration of a bidder by the number of days from the date she placed her first bid through the date she placed her last bid in our sample. All three measures of participation indicate that the vast majority of the bidders at BigDeal were fleeting participants. The 75th percentile of bidders’ duration is only 4 days, the 75th

percentile of the number of auctions participated is 8, and the 75th percentile of the number of bids is 55. A small percentage of bidders were persistent participants. Only 10.1% of the bidders lasted 29 days or more, 5.2% of the bidders played in 25 auctions or more, and 5.1% of the bidders placed 300 bids or more.

Figure 2(a): Weekly Sum of New Bidders Each Day






Figure 2(b): Weekly Average of Daily Percentage of New Bidders Whose Duration Is No More Than 7, 14, or 28 Days



To directly see the phenomenon of a revolving door of new bidders, consider Figures 2(a) to 2(d), which present the dynamics of bidder participation over time. Figure 2(a) shows the weekly sum of each day's new bidders at BigDeal. Figure 2(b) shows the weekly average of the daily percentage of new bidders whose duration was no more than 7, 14, or 28 days. Figure 2(c) shows the

weekly average of the daily percentage of new bidders whose total number of auctions was no more than 7, 14, or 28. Note that bidders who joined BigDeal toward the end of our sample naturally have lower participation intensity. Figure 2(d) shows the weekly average of the daily percentage of bidders who appeared on the website for less than 7, 14 or 28 days. Most bidders on a given day were relatively new to the website. Note that the weekly averages here are all weighted by the number of bidders on each weekday. Note also that the sudden drop in the number of new bidders in Figures 2(a) and 2(d) around week 40 of 2010 was related to the sudden drop in the number of non-beginner auctions in Figure 1 around the same time.¹⁵


Figure 2(c): Weekly Average of Daily Percentage of New Bidders Who Bid in No More Than 7, 14, or 28 Auctions



To facilitate exposition, we classify bidders into three mutually exclusive groups: persistent, fleeting, or moderate bidders. Whether a bidder is fleeting or persistent is inherently a matter of degree. We shall use the following working definition. A bidder is persistent if her total number of auctions is at least 50. A bidder is fleeting if her total number of auctions is at most 15. Moderate bidders are those in between, neither persistent nor fleeting. Panel A of Table 3 presents summary statistics of the three groups of bidders.

¹⁵On September 27, 2010, the number of new bidders decreased suddenly and caused a big loss for BigDeal. So BigDeal offered fewer auctions the next day. Though the number of new bidders recovered in October, BigDeal retained the low level of supply until the end of November. BigDeal started to offer beginner auctions on November 30, 2010.

Figure 2(d): Weekly Average of Daily Percentage of Bidders Who Joined the Site No More Than 7, 14, or 28 Days Ago



By our definition, 89.2% of the bidders are fleeting, and only 1.8% persistent. However, the persistent bidders won 64.4% of the regular or non-beginner auctions. Note that 96% of the fleeting bidders and 61% of the moderate bidders never won a regular auction, and only 10.2% of the persistent bidders never won a regular auction. Subsection 4.2 shows that 94% of the fleeting bidders lost money (after considering the effect of beginner auctions).

Why do most new bidders lose money and then quit quickly? Our interpretation is that most bidders, before playing, did not know the difficulty of winning penny auctions or the existence of persistent bidders who win most of the auctions. Though we do not have direct evidence, it appears plausible that many bidders may have been enticed by the advertisements of deep discounts and joined the website in the hope of winning some items easily and cheaply. If so, such bidders quickly realized that their expectations were wrong.

4.2 Bidder or Auctioneer Profit

In this subsection, we estimate the auctioneer’s profit and each bidder’s profit or loss. Our results show that BigDeal made considerable profit from the fleeting and moderate bidders, but lost money to the persistent bidders as a group. This finding confirms that the main source of auctioneer profit is the revolving

Table 3: Descriptive Statistics of Three Groups of Bidders

	Fleeting	Moderate	Persistent	All bidders
Panel A:				
Number of bidders	184,689	18,634	3,746	207,069
(% of all bidders)	(89.2)	(9.0)	(1.8)	(100)
Number of bids	7,132,908	4,902,971	10,562,156	22,598,036
(% of all bids)	(31.6)	(21.7)	(46.7)	(100)
Number of regular auction wins	9,175	18,789	50,609	78,573
(% of all regular auction wins)	(11.7)	(23.9)	(64.4)	(100)
% of bidders who never won a regular auction	96.1	60.9	10.2	91.3
Panel B:				
Bidder profit in token auctions (0.9)	-474,007	-378,930	-384,452	-1,237,389
Bidder profit in token auctions (0.8)	-540,186	-445,885	-494,364	-1,480,435
Bidder profit in token auctions (0.7)	-575,081	-485,693	-570,833	-1,631,608
Bidder profit in all auctions (0.9)	-3,493,993	-1,176,934	924,342	-3,746,585
Bidder profit in all auctions (0.8)	-3,560,172	-1,243,889	814,430	-3,989,631
Bidder profit in all auctions (0.7)	-3,595,067	-1,283,697	737,961	-4,140,803
% of bidders who lost money (0.9)	94.3	86.1	66.7	93.0
% of bidders who lost money (0.8)	94.4	86.6	67.9	93.3
% of bidders who lost money (0.7)	94.5	86.9	68.8	93.4

Notes: Regular auctions refer to non-beginner auctions. The three numbers in parentheses (0.9, 0.8, and 0.7) are the assumed possible discount rates for bid tokens bought through the BIN option. See subsection 4.2 for explanations.

door of new bidders, suggesting that penny auctions cannot sustain excessive profits without attracting new bidders who will lose money. The persistent bidders differ greatly in their performance; while most persistent bidders lost money, a small percentage of persistent bidders made significant amounts of positive profits, confirming that the experienced bidders are heterogeneous.

4.2.1 Profit Definition and Computation

We define a bidder's profit as the total value of the products she won or bought minus her total cost. We define the auctioneer's profit as its revenue minus the total value of the products auctioned or sold through the BIN option. These two definitions suit the purpose of studying whether penny auctions generate

revenues that are above the values of the products sold, and if so, which types of bidders are the sources of the excessive profit. We are not concerned with the auctioneer’s profit over its cost, which we do not observe. Given these definitions, our conclusion that penny auctions cannot sustain excessive profits does not mean that they cannot sustain normal economic profits. Since the auctioneer’s revenue equals bidders’ total cost, one dollar lost by a bidder is one dollar of additional profit earned by the auctioneer. We describe below how to compute profit from the bidders’ perspective.

Following the literature on penny auctions, we approximate the value of a product by the retail price of the same product at Amazon.com.¹⁶ We find 61.7% of the non-token BigDeal auctions involved products sold at Amazon.¹⁷ For these auctions, the Amazon prices were, on average, 78.0% of the retail prices posted by BigDeal. In 97.6% of these auctions, the Amazon price was smaller than the BigDeal retail price. We assume that the value of a non-token product that does not have a matched Amazon product was 78% of the retail price posted by BigDeal. We will discuss the value of bid tokens below.

A bidder’s profit depends on the number of auctions she won and lost and the dollar amount she made in each of the auctions she played. Consider bidder i who participated in $n = 1, 2, \dots, N$ auctions. Let π_{in} denote bidder i ’s profit (or loss) from her n th auction. Her total profit, π_i , is then $\pi_i = \pi_{i1} + \pi_{i2} + \dots + \pi_{iN}$. It is straightforward to calculate her profit in any auction that she won. It is a bit involved to calculate her loss in an auction that she did not win because of the need to estimate whether she exercised the BIN option. We use the following two observations to estimate whether a bidder exercised the BIN option. Suppose bidder i lost an auction after placing b bids, and the posted retail price for the product is r . To exercise the BIN option, bidder i needs to pay $r - bc$ to purchase the product, where c is the cost per

¹⁶We searched Amazon.com in mid-June 2011, and found an exact match for 601 of the 1,687 unique non-token products auctioned by BigDeal. The vast majority of these matched products were sold by multiple sellers on Amazon, often at different prices. We recorded the price posted by the main or featured seller, which is the manufacturing firm of the product or Amazon itself or a large seller. For iPads, we use Apple’s official prices.

¹⁷Non-token auctions refer to any auctions that do not feature packs of bid tokens.

bid.

If the BIN option is available, then (a) the inequality $bc \leq r$ must hold; (b) bidder i exercises the BIN option if and only if $r - bc \leq v$.

Part (a) says that bidder i 's cost of total bids should not exceed the posted retail price of the product if the BIN option is available. Once a bidder's cost has reached the posted retail price, she can exercise the BIN option and obtain the product for free. We present some evidence for this observation in subsection 4.5. Part (b) says that bidder i exercises the BIN option if and only if her additional cost of bids, $r - bc$, is no more than v , the value of the product.

Assume her first auction is for a non-token product and the second auction is for bid tokens. We demonstrate here how to compute her profits in these two auctions. Her profits for the other $N - 2$ auctions can be similarly computed.

Suppose the posted retail price for the product in her first auction is r_1 , the value of the product is v_1 , the final auction price is p_1 , and her number of bids is b_{i1} . Then, if she won, her profit is

$$\pi_{i1} = v_1 - p_1 - 0.75b_{i1}. \tag{1}$$

Note that the cost of a bid is always \$0.75. The winner of a bid pack auction may obtain tokens at substantial discounts, but when such tokens are used in subsequent auctions, the opportunity cost of such a token should still be the price of a token, \$0.75.

If bidder i lost, her profit depends on whether the BIN option is available, and if the option is available, whether she exercises it. Suppose the BIN option is not available. Then her profit is simply

$$\pi_{i1} = -0.75b_{i1}. \tag{2}$$

If the BIN option is available, bidder i 's profit depends on whether she

exercises the BIN option:

$$\pi_{i1} = \begin{cases} -0.75b_{i1} & \text{if } r_1 - 0.75b_{i1} > v_1 \\ -(r_1 - v_1) & \text{if } r_1 - 0.75b_{i1} \leq v_1 \end{cases}. \quad (3)$$

If the cost of exercising the option is bigger than the value of the product, $r_1 - 0.75b_{i1} > v_1$, she does not exercise the option and her loss is simply her bidding costs, $0.75b$. If she exercises the option, she uses r_1 to obtain a product of value v_1 , so her loss is $r_1 - v_1$. Equation (3) assumes implicitly that $r_1 > v_1$. In the rare event that $r_1 < v_1$, bidder i exercises the BIN option after losing and obtains a positive profit.

Consider the second auction, which features bid tokens. If she won this auction, her profit can be computed as in equation (1). Since a bid token's price is \$0.75, we presume its value is \$0.75 for any winner of any token auctions. If she lost this auction and the BIN option is not available, then her loss can be computed as in equation (2). If she lost this auction but the BIN option is available, her loss can be computed as in equation (3). However, the value of a bid token is no longer \$0.75 when she is deciding whether to exercise the BIN option for the following reason. When BigDeal made the BIN option available to token auctions in late November 2010, it imposed a restriction upon tokens bought through the BIN option:¹⁸ such tokens have reduced values toward exercising the BIN option in a subsequent auction.¹⁹ The value of a token with this usage restriction should be smaller than \$0.75, but we do not have a way of estimating the reduced value.

Fortunately, our overall estimates of bidder profits are not sensitive to

¹⁸Recognize that some usage restrictions have to be imposed on the BIN option for token auctions. Otherwise, since the value of a token purchased through the BIN option is \$0.75, all losing bidders will exercise the BIN option and fully recover the bids they have lost; no bidder ever loses in such auctions. Since the winner of a token auction may obtain a discount, the auctioneer most likely loses money by conducting such token auctions.

¹⁹Suppose a bidder lost an auction of 100 bid tokens after placing 90 bids. She can exercise the BIN option and obtain 100 bid tokens by paying \$7.50 ($= 75 - 90 \times 0.75$), which is called the BIN price for this bidder. The value of a bid obtained this way toward exercising the BIN option in a subsequent auction is only \$0.075, which equals the bidder's BIN price (\$7.50) divided by the number of bids obtained through the BIN option (100).

how bidders discount tokens bought through the BIN option. This is because the BIN option was available for token auctions for only about 25% of the sample period and the discount rate only affects bidders whose number of bids in a token auction was significant enough to consider exercising the BIN option. Consider three possible reduced values for a BIN-purchased bid token: 0.9×0.75 , 0.8×0.75 , and 0.7×0.75 . Call 0.9, 0.8, and 0.7 the discount rates. Table 4 contains the distribution of bidder profits from all auctions, with bidders' losses in token auctions computed using these three possible discount rates. The difference between any two of the three 10th percentiles is less than a dollar, and so is the difference between any two of the three 90th percentiles. Only the extreme percentiles noticeably differ; a smaller discount rate, which implies bigger loss upper bounds, leads to a slightly smaller extreme percentile. In addition, the Spearman rank order correlation coefficient is above 0.99 between any pair of the three bidder profits.

Table 4: Distribution of Bidder Profit from All Auctions

	0.05%	0.1%	1%	10%	50%	90%	95%	99%	99.9%	99.99%
Bidder profit (0.9)	-1,798	-1,278	-342	-74	-9.0	-.75	6.25	166	2,499	15,433
Bidder profit (0.8)	-1,860	-1,312	-352	-75	-9.0	-.75	5.96	160	2,471	15,395
Bidder profit (0.7)	-1,974	-1,359	-358	-75	-9.8	-.75	5.59	156	2,448	15,358

Note: The three numbers in parentheses (0.9, 0.8, and 0.7) are the assumed possible discount rates for bid tokens bought through the BIN option.

We use the relationship between bidder profit and bidder group to further illustrate that our results are not sensitive to the assumed discount rate for tokens purchased through the BIN option. Consider panel B of Table 3, which contains, by bidder group, bidder profits from token auctions only, bidder profits from all auctions, and proportion of bidders who lost money when considering all auctions. It is apparent that these three statistics are not sensitive to the assumed discount rate (0.9, 0.8, or 0.7) for BIN-purchased tokens. Therefore, we shall report results assuming 0.8 is the discount rate for such tokens.

4.2.2 Sources of Auctioneer Profit

The fleeting bidders together lost \$3.56 million in all auctions, and 94.4% of the fleeting bidders lost money. The moderate bidders together lost \$1.24 million in all auctions, and 86.6% of the moderate bidders lost money. The persistent bidders as a group, however, made a positive profit of \$0.81 million in all auctions, though 67.9% of the persistent bidders still lost money. BigDeal thus generated a total profit of \$3.99 million, 15.1% of the total value of the products it auctioned or sold through the BIN option. The profit margin of 15.1% for BigDeal is much smaller than the profit margin of 150% found by Augenblick (2011) for Swoopo. We present some evidence in subsection 4.5 that the BIN option reduces the profit margins of auctions of the same product. The total value of the products auctioned (\$9.9 million) is smaller than the total value of the products sold through the BIN option (\$16.6 million).

Figure 3: Number of Auctions Versus Bidder Profit (Persistent Bidders)




Figure 3 shows the relationship between persistent bidders' profit and the number of auctions they participated in. Some of the persistent bidders lost a considerable amount of money while others earned a significant amount. For example, 2 bidders lost over \$10,000, while 30 earned over \$10,000; 93 bidders lost at least \$2,000 each, and together they lost \$333,291; 247 bidders earned at least \$2,000 each, and together they earned \$1,700,824. What causes the significant difference in bidders' performance? Somewhat to our surprise, it does not appear that larger numbers of auctions are associated with bigger profits. In subsection 4.3, we present evidence that persistent bidders' performance is

highly correlated with their strategic sophistication.

Figure 4(a) shows the auctioneer's weekly profit. The profit was small in the first few weeks since the number of auctions was small. Figure 4(b) shows the weekly average of the percentage of profit each day generated from three groups of bidders: those who had appeared on the website for 7 days or less, those between 8 and 28 days, and those 29 days or more. The vast majority of the auctioneer's profit in almost all weeks came from those who joined the website less than 7 days earlier, and the auctioneer lost money in most weeks to those bidders who stayed on the website for over 4 weeks.

Figure 4(a): The Auctioneer's Weekly Profit






Figure 4(b): Weekly Average of Daily Percentage of Profit Generated from Three Group of Bidders



4.3 Strategic Sophistication and Persistent Bidders' Performance

In this subsection, we present evidence that (1) persistent bidders differ in their strategic sophistication, and (2) strategic sophistication is predictive of persistent bidders' overall and *future* performance.²⁰ The existence of persistent bidders who make significant positive profits suggests that not all bidders suffer from behavioral biases when bidding in penny auctions. This finding also provides a natural explanation for why it is difficult for inexperienced bidders to win and why penny auction websites impose win limits. The existence of persistent but unsophisticated bidders, on the other hand, suggests that a small number of bidders may have gamblers' preferences.

Figure 5: Histogram of the Timing of Manual or Automatic Bids




When measuring a bidder's strategic sophistication, we only consider manual bids that were placed in the middle of the 30-second timer. To see our definition of "the middle," consider Figure 5(a), which shows the histogram of the timing of all manual bids (21.5 million) that were placed after the 30-second timer started. The vast majority of these manual bids were placed either at the beginning or at the end of a time period; 68.5% were in the first 5 seconds and 13.7% in the last 4 seconds. We consider manual bids only

²⁰The results for moderate bidders, not presented here, are qualitatively similar, but our measure of strategic sophistication does not characterize fleeting bidders well.

because bidders do not have control over the timing of those bids placed by the bid agent. Figure 5(b) shows the histogram of the timing of all the bids (2.1 million) placed by the bid agent. To be conservative, we classify a manual bid to be in the middle of the 30-second time period if it was placed from the 10th second through the 22th second.

Figure 6: Bidder Profit and Percentage of Middle or Aggressive Bids



Persistent bidders differ in their degree of strategic sophistication. Figure 6(a) shows the relationship between strategic sophistication and bidder profit. Smaller proportions of middle bids are associated with higher bidder profits. For example, 986 of the 3,746 persistent bidders placed less than 5% of their bids in the middle, and together they earned a total profit of \$1,149,395. In contrast, 374 placed more than 20% of their bids in the middle, and together they lost \$120,458. Figure 6(b) shows the relationship between persistent bidders' profits and their proportions of bids placed in the first 5 seconds. The most successful bidders are not the ones with the largest proportion of their bids in the first 5 seconds. The most successful bidders tend to place their bids at both the beginning and at the end of the 30-second timer period, but not in the middle of the time clock.

We use the model below to estimate the relationship between strategic sophistication and bidder profit:

$$\pi_i = c + \beta_1 Middle_i + \beta_2 N_i + \beta_3 Middle_i \cdot N_i + \epsilon_i, \quad (4)$$

where π_i is bidder i 's total profit or loss, $Middle_i$ is bidder i 's proportion of middle bids, and N_i is bidder i 's total number of auctions. Note that the value of $Middle_i$ is between 0 and 100, not between 0.01 and 1. The interaction term $Middle_i \cdot N_i$ is meant to capture the idea that the impact of strategic sophistication on a bidder's profit depends on the number of auctions in which she has played. The impact of strategic sophistication is expected to be bigger for bidders who participated in a larger number of auctions.

Table 5: The Effect of Strategic Sophistication on Bidder Profit

	Dependent variable			
	Bidder profit in all auctions		Bidder profit after the first 30 auctions	
	(1)	(2)	(3)	(4)
Proportion of middle bids in all auctions	-67.5*** (-11.00)	36.2*** (3.90)		
Proportion of middle bids in a bidder's first 30 auctions			-36.93*** (-7.00)	-7.75 (-1.14)
Number of auctions		11.9*** (15.51)		
Number of auctions - 30				5.40*** (9.12)
Proportion of middle bids \times Number of auctions		-0.92*** (-14.45)		
Proportion of middle bids in the first 30 auctions \times (Number of auctions - 30)				-0.29*** (-6.56)
Constant	918.2*** (11.82)	-462.5*** (-3.95)	590.27*** (8.71)	68.30 (0.77)
Number of observations	3,746	3,746	3,746	3,746
Adjusted R^2	0.03	0.09	0.01	0.03

Notes: The numbers in parentheses are t -statistics. *** $p < 0.01$.

Table 5 reports the ordinary least square (OLS) estimates for equation (4). In specification (1), the proportion of middle bids is the only explanatory variable, and its coefficient, as expected, is significantly negative. The marginal effect of a 1% increase in proportion of middle bids is estimated to be \$-67.5. In specification (2), we add in the number of auctions and the interaction term.

The estimated marginal effect of the proportion of middle bids is $36.2 - 0.92N_i$, which is negative (since $N_i \geq 50$) and is increasingly negative for bigger N_i . The estimated marginal effect of N_i is $11.9 - 0.92Middle_i$, which is negative for unsophisticated bidders and positive for strategically sophisticated bidders. We note that the variable N_i is endogenous, so we caution that the estimated marginal effect of N_i is only suggestive. We offer more discussions on the relationship between a bidder's profit and her number of auctions in the next subsection.

A concern with equation (4) is that a bidder's proportion of middle bids and her total profit are determined simultaneously. One way to address this endogeneity problem in equation (4) is to see if our measure of strategic sophistication predicts bidders' *future* performance. That is, we can define $Middle_i$ as bidder i 's proportion of middle bids in her, say, first 30 auctions and π_i as her total profit after her first 30 auctions. In this case, N_i should be defined as bidder i 's total number of auctions minus 30. Specifications (3) and (4) in Table 5 report the estimated results for equation (4) using the new measures of the dependent and independent variables. The results remain similar. When the proportion of middle bids in a bidder's first 30 auctions is the only independent variable, its coefficient is again significantly negative. When the interaction terms are added, the estimated marginal effect is again negative and increasingly negative for bigger N , and the estimated marginal effect of N is again negative for unsophisticated bidders and positive for sophisticated bidders.

4.4 Strategic Sophistication and Learning

In this subsection, we first clarify what our measure of strategic sophistication is and is not. We then present evidence that whether a persistent bidder learns to play better depends critically on her strategic sophistication. Our results indicate strongly that not all bidders have the same learning function. Highly sophisticated bidders start to make positive profits from their first few auctions, and they learn to play better. Highly unsophisticated bidders, on

the other hand, lose money in their first few auctions and do not learn to play better; these unsophisticated bidders may be characterized as gamblers in that they continue to play despite consistently losing money.

It turns out that persistent bidders, on average, do not decrease their proportion of middle bids as they gain more experience. This finding suggests that a bidder's proportion of middle bids reflects a relatively stable attribute of a bidder. This attribute, in our opinion, is the degree to which a bidder is mindful of her competition.

We emphasize that a bidder's proportion of middle bids captures only a basic aspect of her bidding behavior and does not fully characterize her strategic ability. That is, a bidder's proportion of middle bids is not a comprehensive measure of her strategic sophistication. Two bidders with the same proportion of middle bids may not play penny auctions the same way; they may differ in making such decisions as which auction to participate in and when to bid aggressively in an auction. Since a high proportion of middle bids is indicative that a bidder is not mindful of her competition, we hypothesize that such bidders are unlikely to learn to play better in more complicated aspects of the game that are not captured by our measure of strategic sophistication. On the other hand, a bidder with a low proportion of middle bids may learn to play better in more complicated aspects of the game as she gains more experience. An analogy might be useful. Proportion of middle bids as an imperfect measure of strategic sophistication is similar to GRE quantitative score as an imperfect measure of research ability in economics. GRE quantitative score is not a comprehensive measure of research ability, but a student with a poor GRE quantitative score is unlikely to do well in economic research.

To see that experienced bidders, on average, did not learn to decrease their proportions of middle bids, consider a simple fixed-effect regression model in which the dependent variable is bidders' proportion of middle bids. A bidder's proportion of middle bids in an auction in which she placed only one or two bids is not a reliable measure of a bidder's strategic sophistication. Since many bidders do submit only one or two bids in some auctions and a bid may be placed before the 30-second countdown clock started, we group consecutive

auctions into groups and consider bidders' proportion of middle bids in such groups of auctions. Consider the following fixed-effect model:

$$Middle_{ig} = c + \alpha Exp_{ig} + \theta_i + \epsilon_{ig}, \quad (5)$$

where $Middle_{ig}$ is bidder i 's proportion of middle bids in auction group g , Exp_{ig} is bidder i 's experience when playing in group g , and θ_i is the bidder fixed effect. To see how we measure $Middle_{ig}$ and Exp_{ig} , consider an example. Suppose bidder i played in a total number of 58 auctions. Order these 58 auctions by time and let every 5 consecutive auctions constitute an auction group; the first 5 auctions are the first group, auctions 6 through 10 the second group, and so on. The experience variable, Exp_{ig} , takes the value of 1 for the first group of auctions, 2 for the second group, and so on. In this example, bidder i 's last group includes three auctions only. The results are not sensitive to the number of auctions included in a group.

Table 6 reports the estimates for equation (5). Specification (1) considers all persistent bidders, and the estimated coefficient for the experience measure is 0.000093 and is not statistically significant at the 5% level. Specification (2) considers persistent bidders who made a positive profit, and the estimated coefficient for the experience measure is -0.000079 but is statistically insignificant. We obtain similar results even if we restrict the sample to the highly successful bidders only; specification (3) considers persistent bidders with a profit of \$2,000 or more. Specification (4) considers persistent bidders with a negative profit, and the estimated coefficient for the experience measure is 0.00023, with a p-value of 0.001. These results suggest that persistent bidders, on average, did not learn to place a smaller percentage of bids in the middle of the 30-second timer.

We use the model below to study whether persistent bidders learn to play better (in other aspects of the game that affect outcome) as they gain more experience:

$$\pi_{in} = c + \delta_1 Exp_{in} + \delta_2 Exp_{in} \cdot Middle_i + \delta_3 Exp_{in}^2 + \varphi_i + \epsilon_{in}, \quad (6)$$

Table 6: The Effect of Experience on Strategic Sophistication

	All persistent bidders (1)	Persistent bidders with a positive profit (2)	Persistent bidders with profits above \$2,000 (3)	Persistent bidders with a negative profit (4)
Experience	0.000093 (1.82)	-0.000079 (-1.14)	-0.000068 (-0.67)	0.00023*** (3.19)
Constant	0.10*** (89.73)	0.085*** (50.2)	0.06*** (18.11)	0.11*** (76.48)
Num. of bidders	3,738	1,199	225	2,539
Num. of obs.	77,579	30,711	10,315	46,868

Note: Dependent variable is a bidder's proportion of middle bids in a group of 5 auctions. Bidder fixed effects are included in all regressions. The reported constant is the average bidder fixed effect. In parentheses are t -statistics based on Huber/White robust standard errors. *** $p < 0.01$.

where the dependent variable π_{in} is bidder i 's profit or loss in her n th auction, Exp_{in} is bidder i 's experience when she plays her n th auction, $Middle_i$ is bidder i 's proportion of middle bids in all of her auctions, and φ_i is the bidder fixed effect. The interaction term is meant to capture the idea that experience improves a bidder's performance only if she is strategically sophisticated enough. In other words, a bidder with too low a strategic ability may not be able to learn to play better at all. Here, $Exp_{in} = n$. The square of experience is added in equation (6) to capture the idea that the marginal effect of experience may diminish as experience increases. The marginal effect of experience is $\delta_1 + \delta_2 Middle_i + 2\delta_3 Exp_{in}$. We expect the estimated coefficients for both δ_2 and δ_3 to be negative. After presenting the estimated results, we discuss two concerns with the interpretation of equation (6).

Table 7 reports the estimated results for equation (6). Bidder fixed effects are included in all specifications. Specification (1) considers all persistent bidders. The estimated marginal effect of experience from this specification is $0.048 - 0.0023 \cdot Middle_i - 0.000035 \cdot n$, confirming that the marginal effect of experience diminishes as the proportion of middle bids increases. Recall that the value of $Middle_i$ is between 0 and 100, not between 0 and 1, so the impact

Table 7: The Effect of Experience and Strategic Sophistication on Bidder Profit per Auction

	Persistent bidders			First 200 auctions	All bidders
	All	% of middle bids < 5%	% of middle bids > 20%		
	(1)	(2)	(3)	(4)	(5)
Experience	0.048*** (4.88)	0.149*** (2.73)	-0.011 (-0.85)	0.087*** (4.40)	0.017*** (3.66)
% of mid. bids	-0.0023*** (-3.97)	-0.011 (-0.83)	0.00006 (0.13)	-0.004*** (-4.28)	
Experience ²	-0.000018*** (-3.60)	-0.00013*** (-4.70)	0.00002*** (2.58)	-0.00007 (-0.82)	-0.000009** (-2.36)
Constant	-0.452 (-1.08)	1.566 (0.95)	-2.599*** (-7.88)	-0.700 (-0.96)	-2.867*** (-21.55)
Num. of bidders	3,746	986	374	521	207,069
Num. of obs.	457,016	115,284	40,175	104,200	1,697,192

Notes: Bidder fixed effects are included in all regressions. Specification (4) considers only the first 200 auctions of the bidders whose number of auctions was greater than 200. The reported constant is the average bidder fixed effect. The numbers in parentheses are t -statistics based on Huber/White robust standard errors. *** $p < 0.01$, ** $p < 0.05$

of strategic sophistication on learning is economically significant.

Specification (2) considers only persistent bidders whose proportion of middle bids was smaller than 5%. The estimates from this specification indicate that bidders whose proportion of middle bids was 5% earned, on average, a small positive profit in their first auctions and about \$9.7 in their 100th auctions. Specification (3) considers only persistent bidders whose proportion of middle bids was more than 20%. The results for these highly unsophisticated bidders are in stark contrast to those for the highly sophisticated bidders. Bidders whose proportion of middle bids was 20%, on average, lost \$2.6 in their first auctions and \$3.4 in their 100th auctions. These results indicate that highly sophisticated bidders learn to play better but highly unsophisticated bidders do not. The highly unsophisticated but persistent bidders may be characterized as gamblers in that they continue to play despite consistently losing money.

One concern with equation (6) is that the estimated learning effect may simply be a selection effect. This alternative interpretation is based on the idea that more sophisticated players self-select to play in more auctions. Bidder selection is an issue of concern, but its effect depends on the sample selected. If we include all of the bidders in our sample, whether they are fleeting, moderate, or persistent, in the learning regression, as in specification (5), the estimates are driven by the selection effect. Similar to the specification (3) estimates for the unsophisticated persistent bidders, the specification (5) estimates for all bidders indicate that bidders on average lose money when the experience variable is not too large. This is expected, because the vast majority of the more than 200,000 bidders lost money, and they dominate the sample when the experience variable takes on small values. Specification (5) estimates indicate that bidders start to break even when the experience variable is large. This is also expected, because the sophisticated and persistent bidders start to dominate when the experience variable is large enough. If we take specification (5) as the learning function for all bidders, we would obtain the result of extremely slow learning, which is clearly misleading. In fact, the sophisticated and persistent bidders start to earn positive profits from their few auctions.

However, if we restrict the sample to persistent bidders only, as we do in this subsection, the successful (persistent) bidders may not choose to play in more auctions than the losing (persistent) bidders do. In fact, as shown in Figure 3, the relationship between a persistent bidder's total number of auctions and her total profit is not clear-cut. The selection effect, if it exists, is attenuated. The argument that more sophisticated bidders self-select to play more auctions has limited applicability to specification (3), which considers only unsophisticated persistent bidders. In fact, the argument does not apply to specification (4), where we restrict the sample to the first 200 auctions of the 521 bidders who played in more than 200 auctions. This specification answers the question of whether the bidders who played in over 200 auctions learned to play better in their first 200 auctions. The estimates, again, indicate a positive learning effect for those bidders with a small proportion of middle bids, but not for those with a large proportion of middle bids.

Another concern is that the estimated learning effect may be a reputation effect. This alternative interpretation is based on the idea that experienced and sophisticated bidders may have reputations that may help them win auctions. However, to be consistent with our results, the reputation argument would require experienced but unsophisticated bidders not to have positive reputations. While acknowledging that the estimated learning effect may partly reflect a reputation effect, we believe the role of reputation is small in our context. First, BigDeal was characterized by a revolving door of new bidders, and most new bidders are unlikely to know which bidders are experienced and sophisticated. It is time-consuming to check the bidding history of previous auctions. Second, sophisticated bidders presumably are the players who may attempt to learn whether their competitors are sophisticated or not. Since sophisticated bidders can learn their competitors' degree of strategic sophistication from their bidding behavior in the *current* auction, we suspect that few bidders try to memorize and recall their competitors' degree of sophistication in the past, especially considering that the number of experienced competitors is large.


4.5 Impact of the BIN Option on Auction Outcome

The BIN option does not affect our arguments on bidder learning across auctions, but it does complicate our estimation of bidder profit or loss. In this subsection, we present evidence that (1) a bidder's total cost of bids²¹ in an auction should not exceed the posted retail price of the product if the BIN option is available, and (2) the BIN option has the effect of reducing profit margin and attracting more bidders.

If a bidder's total cost of bids in an auction reaches the posted retail price of the product being auctioned, she can exercise the BIN option and obtain a product that is the same as the one being auctioned for free. Thus any rational bidder does not want to bid more than the posted retail price when the BIN

²¹The opportunity cost of a bid is always \$0.75, but those bidders who have won token auctions, due to mental accounting, may not consider the cost of a bid to be \$0.75 in a subsequent auction.

Figure 7: The Maximum Number of Bids by Any Bidder for iPad 64GB 3G Auctions



option is available. Consider Figure 7, which shows the maximum number of bids by any bidder for all the auctions featuring the iPad 64GB 3G. This product was auctioned at BigDeal from May 1, 2010, to March 11, 2011, and the BIN option was not available until November 13, 2010. Before the BIN option became available, the maximum number of bids by any bidder exceeded 1,500 in a considerable number of auctions and exceeded 2,000 in five auctions. After the BIN option became available, the maximum number of bids by any bidder exceeded 1,201 in only 6 of the 204 auctions. This is consistent with the fact that the posted retail price for the iPad 64GB 3G is \$899.99, a price that only required 1,200 bids for a bidder to exercise the BIN option for free.

We use the fact that the BIN option was not available in some of the iPad and bid pack auctions to study the effect of the BIN option on auction outcomes. Several values of bid packs (e.g., 30 tokens, 50 tokens, and other values) and two types of iPads experienced a change in the availability of the BIN option.²² In Table 8, we regress four measures of auction outcome on whether the BIN option was available and product fixed effects. The product was a bid pack of a certain value (Panel A) or an iPad of a certain specification (Panel B). Profit per dollar’s worth of product, our measure of profit margin, is defined as the total profit generated by an auction divided by the total value of

²²As we mentioned earlier, the BIN option was not available for bid pack auctions until late November 2010. When the iPad and iPad 2 were released at the beginning, the BIN option was not available “due to inventory restrictions.”

Table 8: The Impact of BIN on Auction Outcome

	Total profit generated by an auction	Profit per dollar's worth of product in an auction	Num. of actual bidders in an auction	Num. of bids in an auction
Panel A: Bid packs				
BIN	-40.89*** (-12.85)	-179.33*** (-48.63)	4.77*** (14.50)	25.37*** (4.71)
Constant	98.60*** (47.57)	183.71*** (76.50)	23.60*** (110.12)	210.72*** (60.03)
Num. of observations	17,726	17,726	17,726	17,726
Panel B: iPads				
BIN	-168.81 (-0.91)	-215.56*** (-11.73)	51.74*** (4.38)	1168.62*** (3.71)
Constant	1743.27*** (14.00)	222.38*** (18.08)	163.11*** (20.64)	3323.49*** (15.76)
Num. of observations	695	695	695	695

Notes: Constant is the average product fixed effects. The numbers in parentheses are t -statistics. The price increment is \$0.01 for all auctions considered in this table. *** $p < 0.01$.

the products this auction sold directly or through the BIN option. The fixed effect estimates indicate that the BIN option reduced the profit margin for both bid pack and iPad auctions, and it reduced the absolute amount of profit for bid pack auctions but not for iPad auctions. The absolute amount of profit for iPad auctions was not significantly reduced because iPad auctions with the BIN option attracted a much larger number of bidders and bids per auction.

5 Conclusion

Can penny auctions sustain excessive profits in the long run? Our evidence suggests it cannot. A key finding of this paper is that BigDeal profited from a revolving door of new bidders, but lost a significant amount of money to experienced bidders as a group. This finding suggests that a penny auction website, to sustain excessive profits, must continuously attract new bidders

who will lose money. We are not claiming that penny auctions cannot sustain normal economic profits.

The key to understanding penny auctions as a selling mechanism is to focus on bidder heterogeneity in strategic sophistication and bidder behavior across auctions instead of possible bidder biases within an auction. Experienced and strategically sophisticated bidders exploit penny auctions. Inexperienced bidders might suffer from various biases when playing, but they receive immediate and clear outcome feedback so that they may learn to quit quickly. Our results thus highlight that behavioral biases are unlikely to persist in markets in which consumers can obtain quick and unambiguous feedback, and that firms' ability to exploit consumer biases is limited by consumer learning.

Our paper also contributes to the large literature that studies what players actually do in games. A central theme of this behavioral game theory literature, based largely on experimental games, is that learning and strategic sophistication are important for understanding subjects' behavior. This naturally raises the question of whether learning and strategic sophistication are important for understanding players' behavior in the field. Our findings in this paper provide strong evidence that the concepts of learning and strategic sophistication are important for understanding players' behavior in a large-scale field game, and that an equilibrium model that presumes all bidders are experienced and fully rational are inadequate to understand this game.

References

- [1] Augenblick, Ned. 2011. “Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis.” Working paper.
- [2] Bajari, Patrick, and Ali Hortaçsu. 2004. “Economic Insights from Internet Auctions.” *Journal of Economic Literature*, 42(2): 457–486.
- [3] Brown, Alexander L., Collin F. Camerer, and Dan Lovallo. 2012. “To Review or Not to Review? Limited Strategic Thinking at the Movie Box Office.” *American Economic Journal: Microeconomics*, 4(2): 1-28.
- [4] Byers, John W., Michael Mitzenmacher, and Georgios Zervas. 2010. “Information Asymmetries in Pay-per-Bid Auctions.” *Proceedings of the 11th ACM Conference on Electronic Commerce*, 1-11.
- [5] Caldara, Michael. 2012. “Bidding Behavior in Pay-to-Bid Auctions: An Experimental Study.” Working paper.
- [6] Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- [7] Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. 2004. “A Cognitive Hierarchy Model of Games.” *Quarterly Journal of Economics*, 119(3): 861-898.
- [8] Costa-Gomes, Miguel, and Vincent Crawford. 2006. “Cognition and Behavior in Two-Person Guessing Games: An Experimental Study.” *American Economic Review*, 96(5): 1737-1768.
- [9] Crawford, Vincent P. 1997. “Theory and Experiment in the Analysis of Strategic Interaction” in David M. Kreps and Kenneth F. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications*, 206-242, Seventh World Congress, New York: Cambridge University Press.

- [10] Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. 2010. "Strategic Thinking." Working paper.
- [11] DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2): 315-372.
- [12] DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review*, 96(3): 694-719.
- [13] Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2011. "Learning from Seller Experiments in Online Markets." Working paper.
- [14] Ellison, Glenn. 2006. "Bounded Rationality in Industrial Organization" in Richard Blundell, Whitney K. Newey, and Torsten Persson (eds.), *Advances in Economics and Econometrics: Theory and Applications, 142-174, Ninth World Congress*, New York: Cambridge University Press.
- [15] Friedman, Milton, and L. J. Savage. 1948. "The Utility Analysis of Choices Involving Risk." *Journal of Political Economy*, 56(4): 279-304.
- [16] Gneezy, Uri., and Ran Smorodinsky. 2006. "All-Pay Auctions—An Experimental Study." *Journal of Economic Behavior and Organization*, 61: 255–275.
- [17] Golec, Joseph, and Maury Tamarkin. 1998. "Bettors Love Skewness, Not Risk, at the Horse Track." *Journal of Political Economy*, 106: 205-225.
- [18] Goldfarb, Avi, and Botao Yang. 2009. "Are All Managers Created Equal?" *Journal of Marketing Research*, 46(5): 612-622.
- [19] Goldfarb, Avi, and Mo Xiao. 2011. "Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets." *American Economic Review*, 101(7): 3130-3161.
- [20] Goodman, Joseph. 2012. "Reputations in Bidding Fee Auctions." SSRN working paper.

- [21] Hinnoosaar, Toomas. 2010. “Penny Auctions”. Working paper.
- [22] Houba, Harold, Dinard van der Laan, and Dirk Veldhuizen. 2011. “Endogenous Entry in Lowest-Unique Sealed-Bid Auctions.” *Theory and Decision*, 71: 269-295.
- [23] Kagel, John H., and Dan Levin. 1986. “The Winner’s Curse and Public Information in Common Value Auctions.” *American Economic Review*, 76(5): 894–920.
- [24] Kagel, John H., Dan Levin, and Ronald M. Harstad. 1995. “Comparative Static Effects of Number of Bidders and Public Information on Behavior in Second-Price Common Value Auctions.” *International Journal of Game Theory*, 24: 293-319.
- [25] Krishna, Vijay. 2002. *Auction Theory*. San Diego, CA: Academic Press.
- [26] List, John A. 2003. “Does Market Experience Eliminate Market Anomalies?” *Quarterly Journal of Economics*, 118(1): 41-71.
- [27] Malmendier, Ulrike, and Young Han Lee. 2011. “The Bidder’s Curse.” *American Economic Review*, 101(2): 749–87.
- [28] Ostling, Robert, Joseph Tao-yi Wang, Eileen Y. Chou, and Colin F. Camerer. 2011. “Testing Game Theory in the Field: Swedish LUPI Lottery Games.” *American Economic Journal: Microeconomics*, 3(3): 1-33.
- [29] Platt, Brennan C., Joseph Price, and Henry Tappen. 2010. “Pay-to-Bid Auctions.” Working paper.
- [30] Raviv, Yaron, and Gabor Virag. 2009. “Auctions by Gambling.” *International Journal of Industrial Organization*, 27: 369-378.
- [31] Shubik, Martin. 1971. “The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation.” *Journal of Conflict Resolution*, 15(1): 109-111.

- [32] Tversky, Amos, and Daniel Kahneman. 1986. “Rational Choice and the Framing of Decisions.” *Journal of Business*, 59(4), part 2: S251-S275.
- [33] Wang, Zhongmin, and Minbo Xu. 2011. “Learning and Strategic Sophistication in Games: Evidence from Online Penny Auctions.” Working paper.