

THREE ESSAYS ON ALL-PAY AUCTIONS

A Dissertation
submitted to the Faculty of the
Graduate School of Arts and Sciences
of Georgetown University
in partial fulfillment of the requirements for the
degree of
Doctor of Philosophy
in Economics

By

Minbo Xu, M.A.

Washington, DC
April 19, 2013

Copyright 2013 by Minbo Xu
All Rights Reserved

THREE ESSAYS ON ALL-PAY AUCTIONS

Minbo Xu, M.A.

Thesis Advisor: Ian Gale, Ph.D..

ABSTRACT

The dissertation includes three research papers on all-pay auctions. The first paper (Chapter 1) considers an all-pay auction for a product in which there is an option for bidders to guarantee purchases at a seller specified posted price P at any time. We find the symmetric pure-strategy equilibria in the first- and second-price all-pay auctions (also called *war of attrition*) with a buy-price option. Under these equilibria the buy-price option will affect high-value bidders' behavior, and improve their welfare. At the same time, the seller can select the optimal posted price to collect more revenue, and the Revenue Equivalence Theorem holds as well. The second paper (Chapter 2) conducts empirical analysis on online penny auctions, which 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 a website to study if penny auctions can sustain excessive profits over time. The overwhelming majority of new bidders lose money, but they quit quickly. 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 who will lose money. The third paper (Chapter 3) proposes a nonparametric estimation approach to empirical analysis of the war of attrition. In order to construct a tractable model, we consider the uncertain competition and derive a structural model with a stochastic number of bidders. We admit the contamination from observables and introduce a deconvolution problem with

heteroscedastic errors into the nonparametric approach. By a two-step nonparametric procedure, we can attain a consistent estimator of the distribution of bidders' private values from the observables. Finally, we apply the estimation procedure to field data from penny auctions.

I would like to thank my advisor, Professor Ian Gale, for his support and encouragement on my research. I am grateful to Professor Guido Kuersteiner for his advice and suggestions. Dr. Zhongmin Wang is a great mentor and friend as we worked together. The research is also dedicated to everyone who helped in every aspect. Of course, all errors are mine.

Finally, I would like to thank my parents, Zhongping Xu and Guizhi Zeng, for their endless support and love.

Many thanks,

Minbo Xu

TABLE OF CONTENTS

Chapter 1 All-pay Auctions with a Buy-price Option	1
1.1 Introduction	1
1.2 Basic Model	3
1.3 Equilibrium in the FPA	5
1.4 Equilibrium in the SPA	10
1.5 Seller's Choice	14
1.6 Conclusion	19
Chapter 2 Selling a Dollar for More Than a Dollar? Evidence from Online Penny Auctions	20
2.1 Introduction	20
2.2 Background, Auction Rules and Data	26
2.3 Theoretical Considerations	32
2.4 Empirical Analysis	35
2.5 Conclusion	59
Chapter 3 Nonparametric Estimation of a War of Attrition with a Stochastic Number of Bidders.....	62
3.1 Introduction	62
3.2 Theoretical Framework	66
3.3 Econometric Model	71
3.4 Monte Carlo Experiments	89

3.5 Empirical Application	92
3.6 Conclusion	98
Bibliography	100

Chapter 1

All-pay Auctions with a Buy-price Option

1.1 Introduction

The Internet and online retailers have altered auction mechanisms in many ways. In this paper, we study the role of *buy-price* options in the all-pay auctions. In addition to the chance of winning a product by making the highest bid, buy-price options allow the consumers to obtain the product at a certain posted price.

Most of the online auctions are English auctions, in which the highest bidder wins and pays the second highest price, and the losing bidders don't pay their bids. As the buy-price options are available, the seller provides bidders the option to end the auction early at the posted price (such prices are called "*Buy-It-Now*" on *eBay*, "*Buy Price*" on *Yahoo!*, and "*Take-It*" on *Amazon*). The Buy-It-Now option at eBay is a "temporary" buy price, that is, the buy-price option disappears once a bidder places a bid. Other buy prices are "permanent" option which are available to bidders through the auction's entire duration (Shunda 2009).

There is a growing literature on auctions studying buy-price options. Budish and Takeyama

(2001) provide a simple discrete value example to show that the buy-price option can increase the expected revenue when bidders are risk-averse. Mathews (2004) shows that eBay's temporary buy-price can be explained by impatience on both sides. Durham et. al. (2004) and Anderson et. al. (2008) provide empirical evidence of the effect of buy-price options on eBay auctions. Hidvégi et. al. (2006) analyze the English auction when the buy-price option is permanent, and they conclude that the buy-price option increases expected social welfare and expected surplus to both sides when either buyers or seller are risk-averse. Kirkegaard and Overgaard (2008) extend the model to multi-unit auctions, in which the buy-price option increases the first auction's revenue, but reduces the second auction's revenue. Bose and Daripa (2009) examine the optimal selling mechanism under the environment that a seller owns two venues: the first venue is a store using a posted price to sell the object, and the second venue is an (online) auction site. In these settings, they argue that the eBay auction with a temporary buy-price is an optimal format in the second venue.

In this paper, we study the effect of buy-price options on all-pay auctions. All-pay auctions are widely used in economics because they capture the essential elements of contests, such as rent seeking, R&D races, political contests, and job promotion tournaments. We haven't seen standard all-pay auctions with buy-price options in the real world. A recent internet auction format, called *penny auction* or *pay-to-bid auction*, has the most similar features to all-pay auctions. See Wang and Xu (2011) for more background.

The most related research is Anderson and Ødegaard (2011). Motivated by penny auctions, they study a setting with two sales channels: a fixed posted price and an all-pay auction with a buy-price option, and they propose a symmetric equilibrium bidding strategy under the condition that the buy-price in the auction is much higher than the fixed posted price. Although

our paper is related and complements their insights, there are some fundamental differences. First, we focus on a single sales channel, in which consumers only have access to the auction market. Second, in addition to the standard all-pay auction (i.e., *first-price all-pay auction*), we also consider the war-of-attrition (i.e., *second-price all-pay auction*).

The main contributions of the paper are as follows. First, we provide the symmetric bidding equilibria for the first- and second-price all-pay auctions as the seller offers the buy-price options. Second, we analyze the setting of optimal posted prices by the seller. An interesting result is that the seller should choose different posted prices in the first- and second-price all-pay auctions, but both formats can attain the same expected revenue. Finally, we discuss the welfare effect of the buy-price options. Under the private valuation framework, the buy-price option can change the high-value bidders' behavior and improve their surplus, but has no effect on low-value bidders. Meanwhile, this option has positive effect on seller's expected revenue.

This paper is organized as follows. In section 2, we set up a basic model. In section 3, we discuss the bidders' equilibrium strategy in the first-price all-pay auction with a buy-price option. In section 4, we analyze a parallel development for the second-price all-pay auction. In section 5, we develop revenue comparisons in the view of the seller's choice. We conclude in section 6.

1.2 Basic Model

We extend the framework of Krishna and Morgan (1997) to a buy-price option setting. Consider a seller who wants to sell a product in a market where N potential bidders $i \in \{1, 2, \dots, N\}$ exist. Assume each bidder has one unit of demand, and his valuation is in monetary units, v ,

which is drawn independently from the distribution function $F(v)$ with a continuous density function $f(v)$. Without any loss of generality, we assume $F(\underline{v}) = 0$, $F(\bar{v}) = 1$, $\underline{v} \geq 0$, $\bar{v} \leq \infty$ and bidders are risk-neutral. All of the above is common knowledge among buyers and the seller.

We model the all-pay auction with a buy-price in the following manner. Prior to the start of an auction, the seller posts the description of the product and the buy-price, $P > \underline{v}$. Each bidder knows his private value of the product based on the item description and his own preference. The number of potential bidders and the distribution of valuations are commonly known, but a bidder doesn't know his rivals' values.

As the auction starts, each bidder submits his willingness to pay, b_i . The seller collects all bids, and then reveals all bids and the identity of the winner. In the first-price all-pay (FPA) auction, the highest bidder wins the product and pays his bid. In the second-price all-pay (SPA) auction, the highest bidder wins the product and pays the second-highest bid. If a tie happens, the winner is selected randomly among the highest bidders. In all cases, losing bidders have to pay the amount of their bids as well.

The buy-price option offers each losing bidder an alternative to buy the product at the reduced price $P - b_i$ after the end of auctions. And this option is not available for the winner. This multi-unit assumption is different from the standard auction format, but it is reasonable in the e-commerce market. As the seller receives an order, it takes days to deliver the product and she can request more from the supplier even if the stock is not enough.

1.3 Equilibrium in the FPA

Let's start with the first-price auction. In this section, we discuss the bidders' behavior in the FPA auction with a buy-price option. Since bidders are symmetric, we consider bidder 1's bidding strategy. In what follows, let the random variable $W_{-1} = \max\{v_j\}_{j \neq 1}$ denote the maximum value of bidder 1's rivals. Under the independence assumption, the conditional distribution function of W_{-1} given that $v_1 = v$ is $F_{W_{-1}}(\cdot | v) = F_{W_{-1}}(\cdot) = F^{N-1}(\cdot)$ and the corresponding density function is $f_{W_{-1}}(\cdot | v) = f_{W_{-1}}(\cdot) = (N-1)F^{N-2}(\cdot)f(\cdot)$.

This is a two-stage game. In the first stage, bidders choose the bids to submit; in the second stage, bidders decide whether to execute the buy-price option. We solve the game by backward induction.

In the second stage, the seller reveals all bids and the identity of winner, and the highest bidder wins one unit of product. Since each bidder has single-unit demand, the winner with (v, b) would receive with payoff $v - b$. Under the alternative buy-price option, the posted price P is given, and the losing bidders can choose whether to execute that option. For a non-winner i with submitted bid b_i , his payoff function for the binary choice is

$$\begin{cases} v_i - \max\{P, b_i\} & \text{if } \textit{execute buy-price option} \\ -b_i & \text{if } \textit{otherwise} \end{cases} . \quad (1.1)$$

So the condition to exercise the buy-price option is $v_i - \max\{P, b_i\} \geq -b_i$ and the non-winner

i 's payoff function is

$$\begin{cases} v_i - b_i & \text{if } b_i > P \\ v_i - P & \text{if } P - v_i \leq b_i \leq P \\ -b_i & \text{if } b_i < P - v_i \end{cases} \quad (1.2)$$

In the first stage, all bidders submit their bids simultaneously and the condition for bidder 1 to win is $b_1 > \max_{j \neq 1} \{b_j\}$.¹ Consider bidder 1's payoff function in the second stage, and then bidder 1's payoff function in the first stage is

$$\begin{cases} v_1 - b_1 & \text{if } b_1 > \max_{j \neq 1} \{b_j\} \\ v_1 - b_1 & \text{if } b_1 \leq \max_{j \neq 1} \{b_j\} \text{ and } b_1 > P \\ v_1 - P & \text{if } b_1 < \max_{j \neq 1} \{b_j\} \text{ and } P - v_1 \leq b_1 \leq P \\ v_1 - \frac{b_1 + P(\#\{j:b_j=b_1\}-1)}{\#\{j:b_j=b_1\}} & \text{if } b_1 = \max_{j \neq 1} \{b_j\} \text{ and } P - v_1 \leq b_1 \leq P \\ -b_1 & \text{if } b_1 < \max_{j \neq 1} \{b_j\} \text{ and } b_1 < P - v_1 \\ \frac{v_1}{\#\{j:b_j=b_1\}} - b_1 & \text{if } b_1 = \max_{j \neq 1} \{b_j\} \text{ and } b_1 < P - v_1 \end{cases} \quad (1.3)$$

We begin with a heuristic derivation of the symmetric equilibrium strategy. Suppose that bidders $j \neq 1$ follow the symmetric strategy $\Phi(b_j; v_j)$, which is defined as the probability that a bidder with value v_j bids less than b_j . Then, bidder 1's winning probability is

$$\begin{aligned} \pi(b_1) &= \Pr\left(b_1 > \max_{j \neq 1} \{b_j\}\right) \\ &= \Pr(b_2 < b_1) \Pr(b_3 < b_1) \dots \Pr(b_N < b_1) \\ &= \left[\int \Phi(b_1; v) dF(v) \right]^{N-1}. \end{aligned} \quad (1.4)$$

¹If $b_1 = \max_{j \neq 1} \{b_j\}$, the object goes to each winning bidder with equal probability.

Bidder 1's expected payoff depends on his own bid, value, and the buy-price as follows:

$$\begin{aligned}
 \Pi(b_1; v_1, P) &= (v_1 - b_1) \pi(b_1) + (v_1 - b_1) [1 - \pi(b_1)] \cdot 1(b_1 > P) \\
 &\quad + (v_1 - P) [1 - \pi(b_1)] \cdot 1(P - v_1 \leq b_1 \leq P) \\
 &\quad - b_1 [1 - \pi(b_1)] \cdot 1(b_1 < P - v_1),
 \end{aligned} \tag{1.5}$$

where $1(\cdot)$ is the indicator function defined as follows,

$$1(E) = \begin{cases} 1 & \text{if } E \text{ is true} \\ 0 & \text{otherwise} \end{cases}.$$

From the payoff function, there is a trivial result as follows:

LEMMA 1: *In the FPA auction with a buy-price option, bidders never bid (strictly) more than the buy-price.*

The lemma is straightforward since every bidder can obtain a unit of the product through the buy-price option, and every bidder has a chance to win the product at a cost less than the buy-price. So bidding over the buy-price is a dominated strategy. We will see that this result holds for the SPA auction as well.

Since bidder 1 never bids more than the buy-price, his bidding strategy maximizes the

expected payoff in the following way:

$$\begin{aligned}
& \max_{b_1} \left\{ \begin{array}{l} \max_{b_1 \leq P-v_1} (v_1 - b_1) \pi(b_1) - b_1 [1 - \pi(b_1)] \\ \max_{P-v_1 \leq b_1 \leq P} (v_1 - b_1) \pi(b_1) + (v_1 - P) [1 - \pi(b_1)] \end{array} \right\} \\
&= \max_{b_1} \left\{ \begin{array}{l} \max_{b_1 \leq P-v_1} v_1 \pi(b_1) - b_1 \\ \max_{P-v_1 \leq b_1 \leq P} v_1 - P + (P - b_1) \pi(b_1) \end{array} \right\}. \tag{1.6}
\end{aligned}$$

In this paper, we focus on pure-strategy equilibria throughout. Suppose that bidder $j \neq 1$ with value v_j uses the symmetric increasing bidding strategy $\beta(v_j)$. We have $\pi(b_1) = F^{N-1}(\beta^{-1}(b_1)) = F_{W_{-1}}(\beta^{-1}(b_1))$.

We denote $g(b_1) = v_1 \pi(b_1) - b_1$ and $h(b_1) = v_1 - P + (P - b_1) \pi(b_1)$. Define the policy functions $\alpha(v_1) = \arg \max_{b_1 \in [0, P-v_1]} g(b_1)$ and $\gamma(v_1) = \arg \max_{b_1 \in [P-v_1, P]} h(b_1)$. Thus

$$\beta(v_1) = \begin{cases} \alpha(v_1) & \text{if } g(\alpha(v_1)) \geq h(\gamma(v_1)) \\ \gamma(v_1) & \text{if } g(\alpha(v_1)) \leq h(\gamma(v_1)) \end{cases}.$$

Case 1: $v_1 = \underline{v}$. The trivial result is $\alpha(v_1) = 0$ and $\beta(v_1) = 0$.

Case 2: $v_1 > \underline{v}$. The first-order condition from $g(b_1)$ implies

$$\alpha(v_1) = \begin{cases} \int_{\underline{v}}^{v_1} t f_{W_{-1}}(\cdot) dt & \text{for } v_1 \leq v^* \\ P - v_1 & \text{for } v_1 > v^* \end{cases} \quad \text{where } \int_{\underline{v}}^{v^*} t f_{W_{-1}}(t) dt = P - v^*.$$

Since $\gamma(v_1) = \arg \max_{b_1 \in [P-v_1, P]} h(b_1) = \arg \max_{b_1 \in [P-v_1, P]} (P - b_1) \pi(b_1)$, the policy function does not depend on v_1 , and we can use v_1 as a randomizing device to define an increasing function $\gamma(v_1)$ with the following conditions:

- (i) $\gamma(v^*) = P - v^*$, and $\gamma(v_1) \leq P$;

(ii) $\gamma'(v_1) > 0$ for $\forall v_1 \in [v^*, \bar{v}]$;

(iii) $(P - \gamma(v_1))\pi(\gamma(v_1))$ is constant for $\forall v_1 \in [v^*, \bar{v}]$.

Combine (i) and (iii), we have $(P - \gamma(v_1))\pi(\gamma(v_1)) = v^*F_{W_{-1}}(v^*)$. Since the bidding strategy is an increasing function and $\gamma(v_1) \leq P$, it implies that $\pi(\gamma(v_1)) = F_{W_{-1}}(v_1)$, and we have $\gamma(v_1) = P - \frac{v^*F_{W_{-1}}(v^*)}{F_{W_{-1}}(v_1)}$.

PROPOSITION 1: *In the FPA auction with a buy-price option, there exists a symmetric equilibrium bidding strategy $\beta(\cdot)$ as follows*

$$\beta(v) = \begin{cases} \int_{\underline{v}}^v t f_{W_{-1}}(t) dt & \text{for } v \leq v^* \\ P - \frac{v^* F_{W_{-1}}(v^*)}{F_{W_{-1}}(v)} & \text{for } v > v^* \end{cases}, \quad (1.7)$$

where the threshold point v^* satisfies the following condition

$$\int_{\underline{v}}^{v^*} t f_{W_{-1}}(t) dt + v^* = P. \quad (1.8)$$

Proof. If all bidders' valuations satisfy $v \leq v^*$, we have the pure-strategy bidding strategy $\beta(v) = \int_{\underline{v}}^v t f_{W_{-1}}(t) dt$ (as in Krishna and Morgan 1997). Since $\pi(\cdot)$ is an increasing function under the all-pay auction, if $v_1 > v^*$, it is a dominated strategy to bid $\beta(v_1) < \beta(v^*) = P - v^*$, and her bid is at least $P - v^*$. We prove the result by showing that $g(\alpha(v_1)) \geq h(\gamma(v_1))$ for $v_1 \leq v^*$ and $g(\alpha(v_1)) \leq h(\gamma(v_1))$ for $v_1 > v^*$.

If $v_1 \leq v^*$, we have $g(\alpha(v_1)) \geq g(P - v_1) = v_1\pi(P - v_1) - (P - v_1) = h(P - v_1) = h(\gamma(v_1))$. The last equality is implied from the result that the policy function does not depend on v_1 . Likewise, if $v_1 > v^*$, we have $g(\alpha(v_1)) = g(P - v_1) = v_1\pi(P - v_1) - (P - v_1) = h(P - v_1) = h(\gamma(v_1))$. ■

COROLLARY 1: *The buy-price option will increase bidders' surplus in the FPA auction.*

Proof. For the standard FPA auction without a buy-price option, the symmetric equilibrium bidding strategy is $\beta^0(v) = \int_{\underline{v}}^v t f_{W_{-1}}(t) dt$ for all $v \in [\underline{v}, \bar{v}]$, and the bidders' expected surplus is defined as

$$\Phi^0(v) = vF_{W_{-1}}(v) - \beta^0(v) = vF_{W_{-1}}(v) - \int_{\underline{v}}^v t f_{W_{-1}}(t) dt.$$

On the other hand, with a buy-price option, the bidding strategy is $\beta(v)$, and the bidders' expected surplus is

$$\Phi(v) = \begin{cases} vF_{W_{-1}}(v) - \int_{\underline{v}}^v t f_{W_{-1}}(t) dt & \text{for } v \leq v^* \\ v - P + v^*F_{W_{-1}}(v^*) & \text{for } v > v^* \end{cases}.$$

For $v \leq v^*$, we can see that $\Phi^0(v) = \Phi(v)$. At the same time, we have $\Phi'(v) > \Phi^{0'}(v)$ for all $v > v^*$, which implies $\Phi(v) > \Phi^0(v)$. ■

1.4 Equilibrium in the SPA

In a standard SPA auction, each bidder submits a sealed bid of b_i , and the highest bidder wins the product while paying the second-highest bid. All other bidders lose and pay their bids exactly. Now we consider that the buy-price option is available for non-winners, and then we extend the SPA auction to a two-stage game.

For a non-winner i with bid b_i , his payoff is $v_i - \max\{0, P - b_i\} - b_i = v_i - \max\{P, b_i\}$ if he uses the buy-price option at the retail price P , otherwise his payoff is $-b_i$. So the condition to

exercise the buy-price option is $v_i - \max\{P, b_i\} \geq -b_i$, and we have the same payoff function for non-winners as (1.2):

$$\begin{cases} v_i - b_i & \text{if } b_i > P \\ v_i - P & \text{if } P - v_i \leq b_i \leq P \\ -b_i & \text{if } b_i < P - v_i \end{cases} .$$

In the first stage, all bidders submit their bids and the condition for bidder 1 to win is $b_1 > \max_{j \neq 1} \{b_j\}$ (in case of tying, a lottery is used), and the winner's payment is the second-highest bid. Thus, bidder 1's payoff function is

$$\begin{cases} v_1 - \max_{j \neq 1} \{b_j\} & \text{if } b_1 > \max_{j \neq 1} \{b_j\} \\ v_1 - b_1 & \text{if } b_1 \leq \max_{j \neq 1} \{b_j\} \text{ and } b_1 > P \\ v_1 - P & \text{if } b_1 < \max_{j \neq 1} \{b_j\} \text{ and } P - v_1 \leq b_1 \leq P \\ v_1 - \frac{b_1 + P(\#\{j:b_j=b_1\}-1)}{\#\{j:b_j=b_1\}} & \text{if } b_1 = \max_{j \neq 1} \{b_j\} \text{ and } P - v_1 \leq b_1 \leq P \\ -b_1 & \text{if } b_1 < \max_{j \neq 1} \{b_j\} \text{ and } b_1 < P - v_1 \\ \frac{v_1}{\#\{j:b_j=b_1\}} - b_1 & \text{if } b_1 = \max_{j \neq 1} \{b_j\} \text{ and } b_1 < P - v_1 \end{cases} . \quad (1.9)$$

Once again we begin with a heuristic derivation of the symmetric equilibrium strategy for the reduced game.

Suppose that bidders $j \neq 1$ follow the symmetric and increasing equilibrium strategies $\lambda(v_j)$. Then bidder 1's expected payoff when he bids b_1 is:

$$\begin{aligned}
\Pi(b_1; v_1, P) &= \int_{\underline{v}}^{\lambda^{-1}(b_1)} (v_1 - \lambda(t)) f_{W_{-1}}(t) dt \\
&+ (v_1 - b_1) [1 - F_{W_{-1}}(\lambda^{-1}(b_1))] \cdot 1(b_1 > P) \\
&+ (v_1 - P) [1 - F_{W_{-1}}(\lambda^{-1}(b_1))] \cdot 1(P - v_1 \leq b_1 \leq P) \\
&+ (-b_1) [1 - F_{W_{-1}}(\lambda^{-1}(b_1))] \cdot 1(b_1 < P - v_1).
\end{aligned} \tag{1.10}$$

From the payoff function, we have the same result as in lemma 1, that bidders never bid more than the buy-price. Thus bidder 1's bidding strategy maximizes the following expected payoff:

$$\max_{b_1} \left\{ \begin{array}{l} \max_{b_1 \leq P - v_1} \int_{\underline{v}}^{\lambda^{-1}(b_1)} (v_1 - \lambda(t)) f_{W_{-1}}(t) dt - b_1 [1 - F_{W_{-1}}(\lambda^{-1}(b_1))] \\ \max_{P - v_1 \leq b_1 \leq P} \int_{\underline{v}}^{\lambda^{-1}(b_1)} (v_1 - \lambda(t)) f_{W_{-1}}(t) dt + (v_1 - P) [1 - F_{W_{-1}}(\lambda^{-1}(b_1))] \end{array} \right\}. \tag{1.11}$$

In the case of $b_1 \leq P - v_1$, the first-order condition is $v_1 f_{W_{-1}}(\lambda^{-1}(b_1)) \frac{1}{\lambda'(\lambda^{-1}(b_1))} - [1 - F_{W_{-1}}(\lambda^{-1}(b_1))] = 0$. At a symmetric equilibrium, we have $b_1 = \lambda(v_1)$ as well, and it yields $\lambda'(v_1) = v_1 \frac{f_{W_{-1}}(v_1)}{1 - F_{W_{-1}}(v_1)}$. Thus we can solve the symmetric equilibrium as

$$\lambda(v) = \int_{\underline{v}}^v t \frac{f_{W_{-1}}(t)}{1 - F_{W_{-1}}(t)} dt \tag{1.12}$$

if $\lambda(v) \leq P - v$.

In the case of $P - v_1 \leq b_1 \leq P$, the first-order condition is $(P - b_1) f_{W_{-1}}(\lambda^{-1}(b_1)) \cdot \frac{1}{\lambda'(\lambda^{-1}(b_1))} > 0$, which means the expected payoff is a strictly increasing function, and the bidder 1 will bid

the maximum bid. Thus we can conclude that $\lambda(v) = P$ as $P - v \leq \lambda(v) \leq P$.

PROPOSITION 2: *In the SPA auction with a buy-price option, there exists a symmetric equilibrium bidding strategy $\lambda(\cdot)$ as follows*

$$\lambda(v) = \begin{cases} \int_{\underline{v}}^v t \frac{f_{W-1}(t)}{1-F_{W-1}(t)} dt & \text{for } v \leq v^{**} \\ P & \text{for } v > v^{**} \end{cases}, \quad (1.13)$$

where the threshold point v^{**} will satisfy the indifference condition:

$$\lambda(v^{**}) = P - v^{**},$$

that is

$$\int_{\underline{v}}^{v^{**}} t \frac{f_{W-1}(t)}{1-F_{W-1}(t)} dt + v^{**} = P. \quad (1.14)$$

As we compare (1.8) and (1.14), the following result is straightforward. We need this result in the next section.

COROLLARY 2: *Under the same retail price P , the threshold point v^* in the FPA auction with a buy-price option is greater than the threshold point v^{**} in the SPA auction.*

Similarly to the previous section, we have the following result on bidders' surplus.

COROLLARY 3: *The buy-price option will increase bidders' surplus in the SPA auction.*

Proof. For a SPA auction without a buy-price option, the symmetric equilibrium bidding

strategy is $\lambda^0(v) = \int_{\underline{v}}^v t \frac{f_{W_{-1}}(t)}{1-F_{W_{-1}}(t)} dt$ for all $v \in [\underline{v}, \bar{v}]$. And the bidders' expected surplus is

$$\Psi^0(v) = \int_{\underline{v}}^v (v - \lambda^0(t)) f_{W_{-1}}(t) dt - \lambda^0(v) [1 - F_{W_{-1}}(v)].$$

On the other hand, with a buy-price option, the bidders' expected surplus is

$$\Psi(v) = \begin{cases} \int_{\underline{v}}^v (v - \lambda(t)) f_{W_{-1}}(t) dt - \lambda(v) [1 - F_{W_{-1}}(v)] & \text{for } v \leq v^{**} \\ \int_{\underline{v}}^v (v - \lambda(t)) f_{W_{-1}}(t) dt + (v - P) [1 - F_{W_{-1}}(v^{**})] & \text{for } v > v^{**} \end{cases}.$$

If $v \leq v^{**}$, the bidders' expected surplus is the same; if $v > v^{**}$, we can imply $\Psi(v) - \Psi^0(v) = \int_{v^{**}}^v (1 - F_{W_{-1}}(t)) dt > 0$. ■

1.5 Seller's Choice

In this section we discuss the expected revenue to the seller from the all-pay auctions with a buy-price option. As a benchmark, recall from Myerson (1981) and Riley and Samuelson (1981), the *Revenue Equivalence Theorem (RET)* holds if there is no buy-price option, and the expected payment by a bidder with value v is

$$e^0(v) = \int_{\underline{v}}^v t f_{W_{-1}}(t) dt.$$

Thus the total expected revenue is $N \cdot E(e^0(v))$ and

$$E(e^0(v)) = \int_{\underline{v}}^{\bar{v}} \left[\int_{\underline{v}}^v t f_{W_{-1}}(t) dt \right] dF(v), \quad (1.15)$$

which is the same for all standard all-pay auctions without the buy-price option.

Now consider that the buy-price option is available. Given the distribution of each bidder's valuation, F , the demand in equilibrium for the all-pay auctions with a buy-price option is expected to be

$$D(F) = F(\tilde{v})^n + N[1 - F(\tilde{v})], \quad (1.16)$$

where \tilde{v} is the threshold point, i.e., $\tilde{v} = v^*$ for the FPA, and $\tilde{v} = v^{**}$ for the SPA.

When the retail price P is the same in both formats, we have $v^* \geq v^{**}$, thus we can conclude that $1 \leq D^1(F) \leq D^2(F)$, i.e., the equilibrium demand in the FPA is less than the equilibrium demand in the SPA when the buy-price option exists, and the equilibrium demand is more than 1 unit for each format.

LEMMA 2: *In both FPA and SPA auctions, the buy-price option increases the total social surplus.*

Note that in the all-pay auction with a buy-price option, the seller will choose the posted price P . For each format, the seller's objective function is to maximize her expected revenue.

In the FPA auction with a buy-price option, the expected payment in equilibrium by a bidder with valuation v is

$$\begin{aligned} e^1(v) &= \begin{cases} \beta(v) & \text{for } v \leq v^* \\ \beta(v)\pi(\beta(v)) + P[1 - \pi(\beta(v))] & \text{for } v > v^* \end{cases} \\ &= \begin{cases} \int_{\underline{v}}^v t f_{W_{-1}}(t) dt & \text{for } v \leq v^* \\ P - v^* F_{W_{-1}}(v^*) & \text{for } v > v^* \end{cases}. \end{aligned}$$

The equation (1.8) implies that $P - v^* F_{W_{-1}}(v^*) = \int_{\underline{v}}^{v^*} t f_{W_{-1}}(t) dt + v^* [1 - F_{W_{-1}}(v^*)]$. We

can examine the total expected revenue collected from these bidders $i \in \{1, 2, \dots, N\}$. Given the assumption of symmetric valuations, the total expected revenue is $N \cdot E(e^1(v))$.

$$E(e^1(v)) = \int_{\underline{v}}^{v^*} \left[\int_{\underline{v}}^v t f_{W_{-1}}(t) dt \right] dF(v) + \left[\int_{\underline{v}}^{v^*} t f_{W_{-1}}(t) dt + v^* [1 - F_{W_{-1}}(v^*)] \right] [1 - F(v^*)]. \quad (1.17)$$

We can see that the effect of the buy-price option on seller's expected revenue is $N \cdot E(e^1(v)) - N \cdot E(e^0(v))$.

$$E(e^1(v)) - E(e^0(v)) = \int_{v^*}^{\bar{v}} \left[v^* [1 - F_{W_{-1}}(v^*)] - \int_{v^*}^v t f_{W_{-1}}(t) dt \right] dF(v). \quad (1.18)$$

Analogously, in the SPA auction with a buy-price option, the expected payment in equilibrium by a bidder with valuation v is

$$e^2(v) = \begin{cases} \int_{\underline{v}}^v \lambda(t) f_{W_{-1}}(t) dt + \lambda(v) [1 - F_{W_{-1}}(v)] & \text{if } v \leq v^{**} \\ \int_{\underline{v}}^{v^{**}} \lambda(t) f_{W_{-1}}(t) dt + P [1 - F_{W_{-1}}(v^{**})] & \text{if } v > v^{**} \end{cases}$$

$$= \begin{cases} \int_{\underline{v}}^v t f_{W_{-1}}(t) dt & \text{for } v \leq v^{**} \\ P - v^{**} F_{W_{-1}}(v^{**}) - \int_{\underline{v}}^{v^{**}} F_{W_{-1}}(t) t \frac{f_{W_{-1}}(t)}{1 - F_{W_{-1}}(t)} dt & \text{for } v > v^{**} \end{cases}$$

and the expected revenue for each bidder is $E(e^2(v))$. At the same time, the equation (1.14) implies that $P - v^{**} F_{W_{-1}}(v^{**}) - \int_{\underline{v}}^{v^{**}} F_{W_{-1}}(t) t \frac{f_{W_{-1}}(t)}{1 - F_{W_{-1}}(t)} dt = \int_{\underline{v}}^{v^{**}} t f_{W_{-1}}(t) dt + v^{**} [1 - F_{W_{-1}}(v^{**})]$.

We can derive that

$$\begin{aligned}
E(e^2(v)) &= \int_{\underline{v}}^{v^{**}} \left[\int_{\underline{v}}^v t f_{W_{-1}}(t) dt \right] dF(v) \\
&\quad + \left[\int_{\underline{v}}^{v^{**}} t f_{W_{-1}}(t) dt + v^{**} [1 - F_{W_{-1}}(v^{**})] \right] [1 - F(v^{**})].
\end{aligned} \tag{1.19}$$

Similarly, we can see that the effect of the buy-price option on seller's expected revenue is $N \cdot E(e^2(v)) - N \cdot E(e^0(v))$.

$$E(e^2(v)) - E(e^0(v)) = \int_{v^{**}}^{\bar{v}} \left[v^{**} [1 - F_{W_{-1}}(v^{**})] - \int_{v^{**}}^v t f_{W_{-1}}(t) dt \right] dF(v). \tag{1.20}$$

From the equations of (1.8) and (1.14), as the distribution of bidders' valuation is given, the posted price P will be the unique determinate variable on the threshold points v^* and v^{**} .

Suppose that the seller knows the optimal threshold point \hat{v} , i.e., $v^* = v^{**} = \hat{v}$,

$$\hat{v} \in \arg \max_{v^*} \int_{v^*}^{\bar{v}} \left[v^* [1 - F_{W_{-1}}(v^*)] - \int_{v^*}^v t f_{W_{-1}}(t) dt \right] dF(v), \tag{1.21}$$

then she will choose the optimal posted price,

$$\int_{\underline{v}}^{\hat{v}} t f_{W_{-1}}(t) dt + \hat{v} = P^1 \tag{1.22}$$

$$\int_{\underline{v}}^{\hat{v}} t \frac{f_{W_{-1}}(t)}{1 - F_{W_{-1}}(t)} dt + \hat{v} = P^2, \tag{1.23}$$

and $P^1 < P^2$, i.e., the optimal posted price in FPA is lower than that in SPA.

PROPOSITION 3: *When the buy-price option is available, the seller will prefer a higher*

posted price in SPA than in FPA. Under the optimal posted prices, the Revenue Equivalence Theorem holds.

Intuitively, under the settings without the buy-price option, the FPA auction and the SPA auction satisfy the *RET* and the expected payment from each bidder is equivalent. As we introduce the option of buy-price into the standard all-pay auctions, the behavior of low-value bidders is unchanged, and the high-value bidders adjust their bidding behavior to attain higher expected surplus. Given that the total social surplus increases at the same time, the seller can attain higher expected revenue as well.

Let's define $\phi(z) = \int_z^{\bar{v}} [z [1 - F_{W-1}(z)] - \int_z^v t f_{W-1}(t) dt] dF(v)$, $z \in [\underline{v}, \bar{v}]$, and then we can show the following properties for $\phi(z)$:

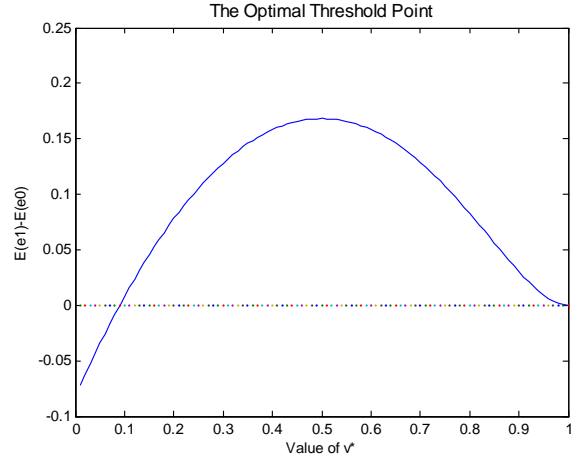
- (a) $\phi(z)$ is a continuous and differentiable function in $z \in [\underline{v}, \bar{v}]$;
- (b) $\phi(\underline{v}) < 0$;
- (c) $\phi(\bar{v}) = 0$ and $\phi'(\bar{v}) = 0$;
- (d) If $\bar{v} \geq \frac{1}{12(N-1)}$, there exists an interval $I \subset [\underline{v}, \bar{v}]$, such that $\phi(z) > 0$, $z \in I$.

Based on these results, we can conclude as follows:

PROPOSITION 4: *In general, there exists an optimal posted price, such that, the buy-price option increases the seller's expected revenue.*

We provide a numerical example to illuminate the result:

Example: Uniform distribution. Suppose the distribution function of bidders' valuation is $F(v) = v$, $v \in [0, 1]$, and $N = 10$ potential bidders. The following graph shows the value of function $\phi(z)$.



1.6 Conclusion

The buy-price option is more and more common in the internet auction market. A hybrid of the posted price and the all-pay auction is relatively new in practice. In this paper, we have studied the effects of buy-price options on the all-pay auctions. Using the similar basic settings as Krishna and Morgan (1997), we analyze two formats of all-pay auctions, the standard all-pay auction, so called first-price all-pay (FPA), and the war-of-attrition, so called second-price all-pay (SPA). First of all, there exist symmetric pure-strategy equilibria in the FPA and SPA with buy-price options. The presence of buy-price options is known to affect the bidders' behavior, specifically, on high-value bidders, and it has no influence on low-value bidders. Second, the option for losing bidders to attain the product at reduced-price makes more transactions happen, and it increases buyers' welfare and the total social welfare as well. Finally, the seller can choose the optimal price to post and then collect more expected revenue. From all of these perspectives, the option of buy-price is a successful tool in the internet market.

Chapter 2

Selling a Dollar for More Than a Dollar? Evidence from Online Penny Auctions¹

2.1 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.² 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

¹This chapter is a joint paper with Zhongmin Wang, who is a research fellow at the Resources for the Future.

²Richard H. Thaler, “Paying a Price for the Thrill of the Hunt,” *New York Times*, November 15, 2009.

irrational impulse buyers have.”³ 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 overwhelming majority of new bidders lose money to the website, but they quit quickly. 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.⁴ This conclusion 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 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.⁵

³Mark Gimein, “The Big Money: The Pennies Add Up at Swoopo.com,” *Washington Post*, July 12, 2009.

⁴This feature is shared somewhat by Ponzi schemes. We are not claiming that penny auctions are Ponzi schemes or necessarily scams.

⁵A 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

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.

We find that 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. This 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. We also find that the performance of experienced bidders depends on their strategic sophistication. Sophisticated experienced bidders start to earn positive profits from their first few auctions and learn to play better in subsequent auctions, but unsophisticated experienced bidders lose money in their first few auctions and do not learn to play better over time.

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

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.

(2009) for reviews of the literature.⁶ See Malmendier and Lee (2011) and the references therein for empirical studies of overbidding in auctions. Our finding 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.⁷

Our paper also relates 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 highlight the importance of learning across games and 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 which she places a bid in the middle of the timer. Bids in the middle of the timer, we shall argue, indicate that a bidder is not mindful of her competition.

Four papers on penny auctions (Augenblick 2011; Platt et al. 2010; Hinnosaar 2010; and

⁶DellaVigna and Malmendier (2006) is an excellent example of empirical behavioral industrial organization study.

⁷See List (2003) for evidence that market experiences may eliminate some forms of market anomalies.

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. Only Augenblick (2011) looks into bidder behavior across auctions, and his conclusion is fundamentally different from ours. He writes (p. 2) 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. ... To explain this ..., I show that consumer learning occurs but is extremely slow, allowing the auctioneer to profit during the learning process” (emphasis in original). Our findings suggest that penny auctions are not a sustainable business strategy. We contend that the finding of extremely slow learning is questionable. First, each of Augenblick’s learning regressions considers all bidders in his sample together, presuming that all bidders, sophisticated or not, 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 when the experience variable takes small values: 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 sophisticated and experienced bidders.

Second, even though his learning regressions consider all bidders together, his results still indicate that bidders start to make positive profits after placing a large number of bids. Logic

suggests that such bidders would continue to play and make positive profits in the periods following Augenblick’s bid-level sample (which covers a period of about four months). This raises the possibility of the existence of such bidders even at the beginning of his sample.

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.

Two papers on penny auctions (Caldara 2012; Goodman 2012) appeared after our paper. Goodman (2012) focuses on the role of reputation in penny auctions. Caldara (2012) conducts lab experiments to study penny auctions. He writes (p. 6) that his evidence supports “the [findings] of Wang and Xu (2011) that pay-to-bid auction websites profit from a ‘resolving door of new bidders’”, and he concludes (p. 32) that “excessive revenues will only last as long as pay-to-bid auction websites can attract new, inexperienced bidders”. His lab findings also strongly support our measure of bidder strategic sophistication.

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, emphasizing bidder learning across auctions and bidder heterogeneity in strategic sophistication. Section 4 presents our empirical results. Section 5 concludes.

2.2 Background, Auction Rules, and Data

2.2.1 The Penny Auction Industry

Penny auctions are also known as *pay-to-bid* or *bidding fee* auctions. 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.⁸ In fact, it named penny auctions one of the top 10 scams of 2011.⁹ 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

⁸ “Online Penny Auctions: Friend or Foe?” <http://www.bbb.org/blog/2010/10/online-penny-auctions-friend-or-foe/>.

⁹ <http://www.bbb.org/us/article/bbb-names-top-ten-scams-of-2011-31711>.

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.

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.”¹⁰

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 auctions 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.

¹⁰Brad Stone, “Penny Auction Sites Hurt by Glut of Competitors”, *Bloomberg Businessweek*, August 12, 2010.

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.

2.2.2 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.¹¹ 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,

¹¹Brad Stone, "BigDeal Puts a New Spin on 'Entertainment Shopping'," *New York Times* Bits Blog, December 19, 2009.

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.¹²

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 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 automati-

¹²BigDeal 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.

cally 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.

2.2.3 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 shows the number of regular (non-beginner) auctions ended each day for the entire

Figure 1: Daily Number of Non-beginner Auctions



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.

2.2.4 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.

2.3 Theoretical Considerations

To address the question of whether penny auctions can sustain excessive profits over time, we focus on bidder learning 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.¹³ Such models may generate the result of selling a dollar for more than a dollar, if bidders suffer from behavioral biases¹⁴ 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

¹³ Augenblick (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 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.

¹⁴ Byers 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.

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).¹⁵ 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.

If a bidder decides to play in another auction, she needs to make two more decisions: which auction to participate in and how to bid in the chosen auction. These two decisions are much more complicated in that they involve strategic thinking, and bidders are not given any direct feedback on how to play better. Therefore, we hypothesize that some sophisticated bidders may learn to play better, but unsophisticated ones may lack the strategic ability to do so. Not everyone can learn to play chess or poker at a high level. Indeed, a major finding of the behavioral game theory literature is that subjects in experimental games exhibit heterogeneity in strategic sophistication. We expect this lab finding to extend to the field setting of penny auctions. That

¹⁵Chance 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.

is, we hypothesize that players in penny auctions differ in their strategic sophistication, and strategically sophisticated players are more likely to win. These hypotheses, if proved to be true, point to the existence of experienced bidders who are strategically sophisticated and may earn positive profits from the auctioneer. The existence of such players would make it harder for inexperienced bidders to win penny auctions, and provides an explanation for why penny auction websites impose win limits.

Testing these hypotheses raises the challenge of measuring bidders' strategic sophistication. Similar to much of the behavioral game theory literature, our measure is based on players' behavior in the game. 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. Our basic argument is that placing bids frequently and deliberately at the beginning or at the end of the time clock, but not in the middle of the time clock, reflects strategic thinking. To understand the 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." Our measure is based on the idea that unsophisticated bidders, those that do not analyze the strategic environment, may place bids randomly during a time period, but sophisticated bidders should not place a bid randomly during a time period.

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. Last-second bids reflect strategic thinking in that they allow a bidder to learn about her competitors. 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. A sophisticated bidder may not always bid this way. It may be optimal for a bidder to bid aggressively (i.e., place a bid immediately after a competing bid) for some periods when she thinks that she is competing with a small number of bidders who are not sophisticated. Indeed, many bidders often place a bid immediately after a competing bid and do so repeatedly for some periods. Since a bid in the middle of the period indicates that a player is not mindful of her competition, a large number of middle bids thus suggests a lack of strategic sophistication.

We note that aggressive bids, by themselves, are not a good indicator of a player's strategic sophistication because the effectiveness of aggressive bids depends critically on the competitive environment in which they are used.¹⁶ Augenblick (2011) studies the effect of aggressive bids, irrespective of the bidder who places the bids.

2.4 Empirical Analysis

2.4.1 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

¹⁶Imagine a set of time periods during an auction when a large number of sophisticated bidders are actively competing with bidder A, and a second set of time periods during the same auction when only a small number of unsophisticated bidders are competing with bidder A. Aggressive bids by bidder A are, presumably, less likely to be effective during the first set of time periods than during the second set of periods.

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 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 the number of auctions participated is 8, the 75th percentile of the number of bids is 55, and the 75th percentile of bidders’ duration is only 4 days. A small percentage of bidders were persistent participants. Only 5.2% of the bidders played in 25 auctions or more, 5.1% of the bidders placed 300 bids or more, and 10.1% of the bidders lasted 29 days or more.

It is illuminating to consider the dynamics of bidder participation over time. The finding that most bidders were fleeting participants holds true for essentially all weeks. 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.¹⁷

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



¹⁷On 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(b): Weekly Average of Daily Percentage of New Bidders Whose Duration Is No More Than 7, 14, or 28 Days




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

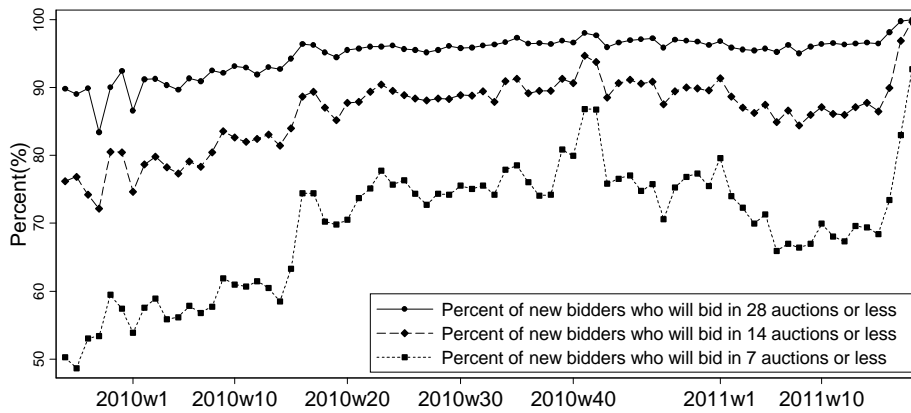
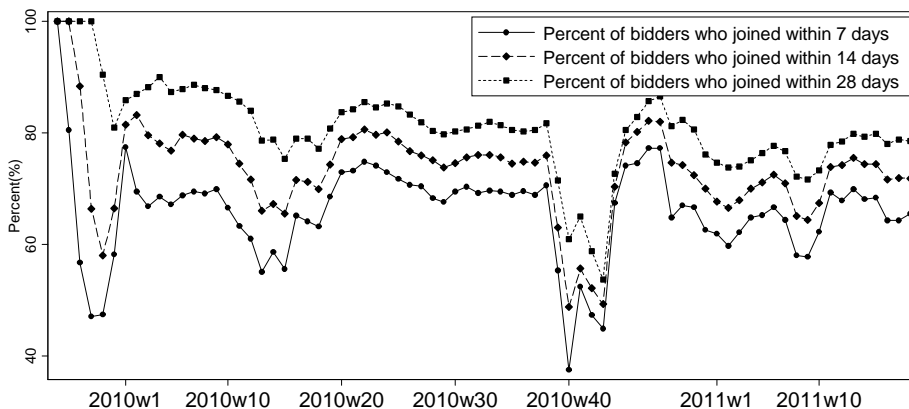


Figure 2(d): Weekly Average of Daily Percentage of Bidders Who Joined
the Site No More Than 7, 14, or 28 Days Ago



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. 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

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.

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.

2.4.2 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 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.

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. 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.¹⁸ We find 61.7% of the non-token BigDeal auctions involved products sold at Amazon.¹⁹ 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.

¹⁸We 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.

¹⁹Non-token auctions refer to any auctions that do not feature packs of bid tokens.

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 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

$$\pi_{i1} = v_1 - p_1 - 0.75b_{i1}. \quad (2.1)$$

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

$$\pi_{i1} = -0.75b_{i1}. \quad (2.2)$$

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

$$\pi_{i1} = \begin{cases} -0.75b_{i1} & \text{if } r_1 - 0.75b_{i1} > v_1 \\ -(r_1 - v_1) & \text{if } r_1 - 0.75b_{i1} \leq v_1 \end{cases}. \quad (2.3)$$

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 (2.3) 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 (2.1). 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 (2.2). If she lost this auction but

the BIN option is available, her loss can be computed as in equation (2.3). 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:²⁰ such tokens have reduced values toward exercising the BIN option in a subsequent auction.²¹ 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 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

²⁰Recognize 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.

²¹Suppose 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).

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

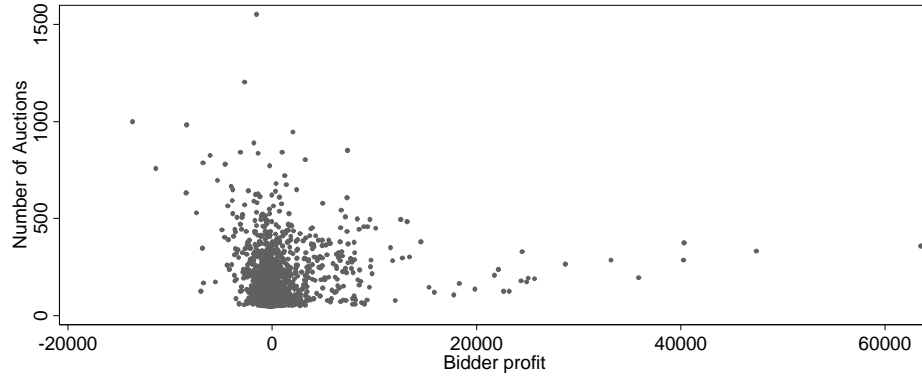
Note: The three numbers in parentheses (0.9, 0.8, and 0.7) are the assumed possible discount rates for bid tokens bought through the BIN option.

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.

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

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



total value of the products sold through the BIN option (\$16.6 million).

Some of the persistent bidders lost a considerable amount of money while others earned a significant amount: 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? Figure 3 shows the relationship between persistent bidders' profit and the number of auctions they participated in. 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


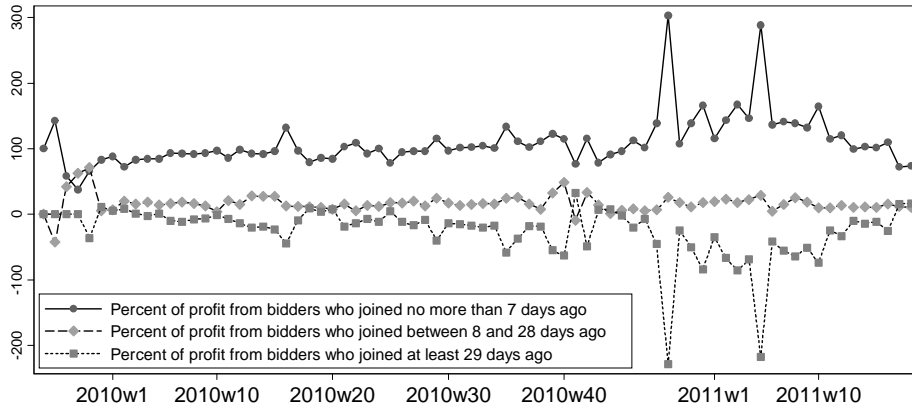


Figure 4(b): Weekly Average of Daily Percentage of Profit Generated from Three Group of Bidders



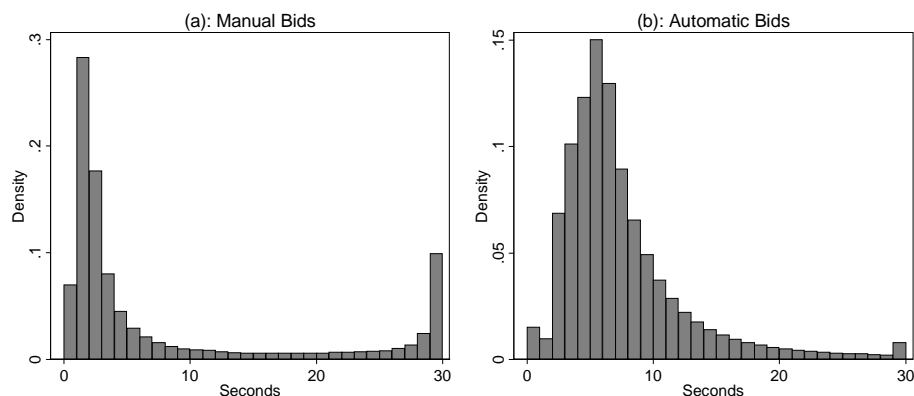
2.4.3 Strategic Sophistication and Persistent Bidders' Performance

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.²² 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 penny auction websites impose win limits.

²²The results for moderate bidders, not presented here, are qualitatively similar, but our measure of strategic sophistication does not characterize fleeting bidders well.

The existence of persistent but unsophisticated bidders, on the other hand, suggests that a small number of bidders may have gamblers’ preferences. This subsection also provides some evidence that higher proportions of aggressive bids are not associated with better performance.

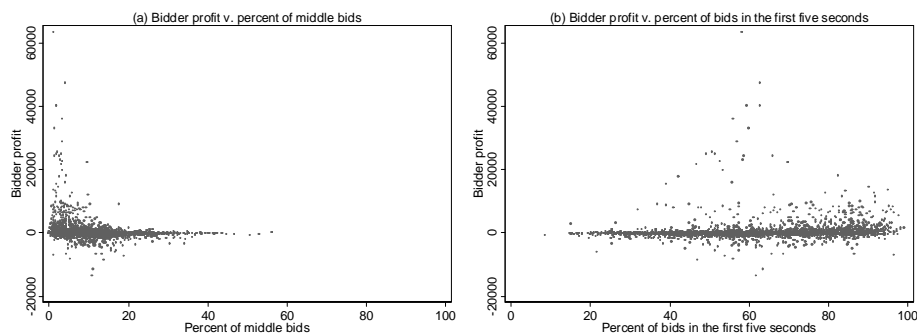
Figure 5: Histogram of the Timing of Manual or Automatic Bids



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 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.

Persistent bidders differ in their degree of strategic sophistication. While 986 of the 3,746 persistent bidders placed less than 5% of their bids in the middle, 374 placed more than 20%

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



of their bids in the middle. Figure 6(a) shows the relationship between strategic sophistication and bidder profit. Smaller proportions of middle bids are associated with higher bidder profits. The 986 persistent bidders with 5% or less middle bids together earned a profit of \$1,149,395. In contrast, the 374 persistent bidders with more than 20% middle bids together lost \$120,458.

Figure 6(b) shows the relationship between persistent bidders' profits and their proportions of bids placed in the first 5 seconds. For the 14 most successful bidders, who each earned at least \$22,000, the percent of bids in the first 5 seconds is between 47% and 66%, and the percentage of bids in the last 4 seconds is between 27% and 45%. None of these 14 bidders placed more than 4.2% of their bids in the middle of the time clock. The most successful bidders thus tend to place their bids at both the beginning and the end of the 30-second timer period, but not in the middle of the time clock. Hence, more aggressive bidding does not necessarily imply more bidder profits. This finding is consistent with our idea that aggressive bids, by themselves, do not reflect a bidder's strategic sophistication.

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

$$\pi_i = c + \beta_1 Middle_i + \beta_2 N_i + \beta_3 Middle_i \cdot N_i + \epsilon_i, \quad (2.4)$$

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. 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 (2.4). 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 (2.4) 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 (2.4) 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 (2.4) 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.

2.4.4 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. Sophisticated bidders start to make positive profits from their first few

auctions, and they learn to play better. Unsophisticated bidders, on the other hand, lose money in their first few auctions and do not learn to play better; these 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:

$$Middle_{ig} = c + \alpha Exp_{ig} + \theta_i + \epsilon_{ig}, \quad (2.5)$$

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 (2.5). 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 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.

Table 6: The Effect of Experience on Strategic Sophistication

	All persistent bidders (1)	Persistent bidders with a positive profit (2)	Persistent bidders with a negative profit (3)
Experience	0.000093 (1.82)	-0.000079 (-1.14)	0.00023*** (3.19)
Constant	0.10*** (89.73)	0.085*** (50.2)	0.11*** (76.48)
Num. of bidders	3,738	1,199	2,539
Num. of observations	77,579	30,711	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. *** $p < 0.01$.

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:

$$\pi_{in} = c + \delta_1 Exp_{in} + \delta_2 Exp_{in} \cdot Middle_i + \delta_3 Exp_{in}^2 + \varphi_i + \epsilon_{in}, \quad (2.6)$$

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 (2.6) 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 (2.6).

Table 7 reports the estimated results for equation (2.6). 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 or as experience increases. 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 unsophisticated bidders are in stark contrast to those for the 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 sophisticated bidders learn to play better but unsophisticated bidders do not. The 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

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 × experience	-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. *** $p < 0.01$, ** $p < 0.05$

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.

2.4.5 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 bids²³ 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 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

²³The 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



the availability of the BIN option.²⁴ 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 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.

2.5 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

²⁴ As 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. *** $p < 0.01$.

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.

The key to understanding penny auctions as a selling mechanism is to focus on bidder learning across auctions and bidder heterogeneity in strategic sophistication 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.

Chapter 3

Nonparametric Estimation of a War of Attrition with a Stochastic Number of Bidders

3.1 Introduction

This paper proposes a nonparametric estimation approach to empirical analysis of a typical all-pay auction: the *war of attrition*, in which bidders place cumulative bids until only one bidder remains. In contrast to other ascending auctions, all bidders in the war of attrition must pay regardless of whether they win, and the optimal bidding strategies are changed accordingly.

Pioneered by Guerre, Perrigne, and Vuong (2000), the nonparametric estimation approach to auctions has attracted a lot of attention (see Athey and Haile (2007) for a recent survey). If we can observe bids in a sample of auctions, we can estimate the distribution of bidders' private values without parametric assumptions or the computations of the Bayesian Nash equilibrium strategies. Recent research has proved that most of standard auctions can be identified from the observed bids (see, e.g., Perrigne and Vuong 1999; Guerre, Perrigne, and Vuong 2000; Athey and Haile 2002, 2007), and all these studies presume that the structure of these auctions is captured

well in a theoretical model and the value of the observed bids is an exact value without errors. In practice, however, the true data-generating process is not perfectly characterized by the existing models; that is, errors exist, and the implications from these abstractions may be misleading.

Haile and Tamer (2003) provide an approach to address this concern for the English auction. Instead of using a complete parametric model, they use an incomplete model consisting of two simple assumptions: bidders neither bid more than their willingness-to-pay nor allow an opponent to win at a price less than their willingness-to-pay. While they construct the informative bounds on the distribution function characterizing bidder demand and information, they cannot identify the distribution of bidders' values. But knowledge of these distributions is so essential¹ that we cannot avoid the issue of identification.

In this paper, we introduce measurement errors into the observables. Based on a reliable structural model, it links the distribution of bidders' private values with *expected* observables. The data-generating in practice provides some observables which are not the true values of the expected observables, and differences are the measurement errors in our estimation process. If the distribution of error terms is known, the main challenge of the identification issue is a deconvolution problem.

In particular, we focus on the war of attrition game with a stochastic number of bidders. To our knowledge, this paper is the first to extend the nonparametric approach to all-pay auctions. In general, especially in internet markets, the number of opponents is unknown, and buyers can generate beliefs about the competition from past history. As we assume that all participants make decisions on their bidding plans at the initial point, we can setup the

¹For example, the distribution of bidders' private values is required to design the optimal format in selling mechanisms.

structural model using the classic framework of *second-price sealed-bid all-pay* auction. Indeed, the theoretical model is a mapping from the distribution of bidders' private values into the expected observables.

Based on the symmetric independent private value (IPV) paradigm, the standard nonparametric estimation approach to auctions can identify the distribution of bidders' private values when the distribution of the expected observables is revealed. Here we use the recently developed method for deconvolution problems with heteroscedastic errors.

The main contribution of this paper is to present a solution for the nonparametric estimation of a war of attrition in which the observed bids are not perfect measures of the expected observables. We develop a nonparametric procedure for recovering the distribution and density of expected observables conditional on some requirements for the error terms. Our procedure applies results from the recent literature on nonparametric deconvolution problems, e.g. De-laigle and Meister (2008) and Wang, Fan, and Wang (2010). The deconvolution method can be extended into the nonparametric estimation approach to auctions in general.

This paper also contributes to the literature on the issue of unobserved auction heterogeneity. Krasnokutskaya (2011) studies the first-price auction environment with private information and unobserved auction heterogeneity. She assumes an environment with unobserved auction heterogeneity which is characterized by two components. First, a common component represents information attributes that are available to all bidders. Second, an individual component reflects information attributes privately observed by each bidder. Based on the multifactor measurement error model, she proposes a nonparametric estimation method to recover distributions of both components from submitted bids. In her paper, the deconvolution method is applied as well. However, unlike the assumption of known density for error terms in my paper,

the multifactor measurement error model requires the common factor across bidders.

Our approach is motivated by a new online selling mechanism, i.e., *penny auction*. Penny auctions, also known as *pay-to-bid* auctions, emerged recently on the internet to sell some popular items, such as Apple's products. Unlike the common English auction, the penny auction is an all-pay mechanism, which runs as follows. Prior to participating in the competition, all participants must buy a stock of bid-tokens. Each bid-token costs a small fee (e.g., \$0.75). Under the internet environment, bidders cannot observe opponents, and their beliefs on the competition are *ex ante* determined. As the game starts, participants place bid-tokens in turn. The placed bid-token is non-refundable, and it offers the bidder an opportunity to buy the product at the current auction price. The auction price for any product is initially 0 and increases by a fixed amount whenever a bid-token is placed. The increment is typically one penny, so they are called penny auctions. The auction winner is the last bidder whose bid-token is not followed by any others. Note that while the winner receives the product and pays the auction price, all losing bidders pay the costs of bid-tokens as sunk costs.

We model the penny auction as a war of attrition with a stochastic number of bidders. Suppose that all bidders apply the automatic bid agents² to set amounts of bid-tokens planned to bid simultaneously at the beginning, and then the bid agents place the bid-tokens one-by-one until only one bidder remains. Let us ignore the auction price since the increment is tiny. In this case, the data-generating process is exactly the same as a classic war of attrition with a stochastic number of bidders. In practice, however, bidding behavior in penny auctions can be asymmetric and bidders' sophistication may matter; therefore, bidders may not follow the

²In most of penny auction websites, bidder robots are available. The robot follows a bidder's setting to place bids automatically.

theoretical predictions and it is difficult to provide an exact interpretation of bidding data (see e.g., Augenblick 2011; Wang and Xu 2012). Here we start from a model with strong abstractions for tractability and then modify the inference of interests accordingly. We will apply the nonparametric estimation procedure to the field data from penny auctions, and analyze the inference of interest.

The remainder of the paper is organized as follows. We first describe our basic structural model and the equilibrium bidding strategies. In Section 3 we consider identification and the nonparametric estimation procedure, and then analyze the asymptotic properties of estimators. Section 4 provides Monte Carlo evidence. For empirical application, we discuss penny auctions in Section 5, and apply the econometric model to field data from penny auctions. We conclude in Section 6.

3.2 Theoretical Framework

In this section, we present the model of a war of attrition following Krishna and Morgan (1997) and Bos (2012). Here we assume the war of attrition to be the IPV second-price sealed-bid all-pay auction with a stochastic number of bidders, and then provide a simple theoretical analysis of the model.

3.2.1 Basic Setup

A single indivisible object is sold through a second-price all-pay auction (*i.e.*, *war of attrition*) with an uncertain number of bidders. There exist at most m potential risk-neutral bidders, with $1 < m < \infty$. Bidders enter the game sequentially, and the arrival process is exogenously given. Assume the actual number of bidders participating in the game is unknown to bidders. Let γ_n

denote the *ex ante* probability that n bidders are present, $n = 1, \dots, m$, such that $\sum_{n=1}^m \gamma_n = 1$. The prior beliefs are exogenously given and commonly known to bidders. Once a bidder k decides to participate in the competition, he updates his belief about the number of opponents. Let p_n^k denote bidder k 's updated probability that there are n bidders present conditional upon the event that he is an active bidder. With presumption of symmetric priors among active bidders, we have

$$p_n = p_n^k = \frac{n\gamma_n}{\sum_{l=1}^m l\gamma_l} \quad (3.1)$$

for all k and n (see McAfee and McMillan 1987).

Throughout we discuss the scenario within the symmetric independent private value (IPV) paradigm. Each potential bidder $k = 1, \dots, m$ is assumed to have a private value v_k for the auctioned object. A bidder does not know other bidders' private values but believes that all private values have been drawn independently from a common distribution function $F(\cdot)$, with associated density function $f(\cdot)$ and support $[v, \bar{v}] \subset \mathbb{R}_+$. All above information is common knowledge, and all bidders are identical *ex ante*. The main idea of this paper is to use the observable information in a sample of games along with the presumption of optimal bidding strategies to learn about the latent private information $F(\cdot)$.

3.2.2 Equilibrium Bidding Strategies

Note that in a classic war of attrition model, all participants submit bids simultaneously, and the highest bidder wins the object with the cost equal to the second-highest bid while all losing bidders pay their bids exactly. As usual, in case of tying, a lottery is used. Let us assume the bidding strategy for bidder k with value v_k is b_k , and we derive the symmetric equilibrium

strategy as follows.

The payoff of the bidder k is

$$\begin{cases} v_k - \max_{j \neq k} b_j & \text{if } b_k > \max_{j \neq k} b_j \\ -b_k & \text{if } b_k < \max_{j \neq k} b_j \\ \frac{v_k}{\#\{j: b_j = b_k\}} - b_k & \text{if } b_k = \max_{j \neq k} b_j \end{cases} . \quad (3.2)$$

We now derive a symmetric bidding equilibrium. Suppose that bidders $j \neq k$ follow the symmetric and increasing equilibrium bidding strategies $\beta(\cdot)$. In what follows, let the random variable $V_{-k}^n = \max\{v_j\}_{j \neq k}$ denote the maximum value of bidder k 's rivals conditional on n bidders participating. Under the independence assumption, the conditional distribution function of V_{-k}^n given that $v_k = v$ is $F_{V_{-k}^n}(\cdot|v) = F_{V_{-k}^n}(\cdot) = F^{n-1}(\cdot)$ and the corresponding density function is $f_{V_{-k}^n}(\cdot|v) = f_{V_{-k}^n}(\cdot) = (n-1)F^{n-2}(\cdot)f(\cdot)$. Since bidders do not know the number of rivals, the expected payoff depends on the beliefs on the competitive environment. We have bidder k 's expected payoff as follows:

$$\Pi(b_k; v_k) = \sum_{n=1}^m p_n \left[\int_{\underline{v}}^{\beta^{-1}(b_k)} (v_k - \beta(t)) f_{V_{-k}^n}(t) dt - b_k [1 - F_{V_{-k}^n}(\beta^{-1}(b_k))] \right]. \quad (3.3)$$

The maximization with respect to b_k leads to:

$$\sum_{n=1}^m p_n \left[v_k f_{V_{-k}^n}(\beta^{-1}(b_k)) \frac{1}{\beta'(\beta^{-1}(b_k))} - [1 - F_{V_{-k}^n}(\beta^{-1}(b_k))] \right] = 0. \quad (3.4)$$

At the symmetric equilibrium, $b_k = \beta(v_k)$, thus equation (3.4) yields

$$\begin{aligned}\beta'(v_k) &= \frac{\sum_{n=1}^m p_n v_k f_{V_{-k}^n}(v_k)}{1 - \sum_{l=1}^m p_l F_{V_{-k}^n}(v_k)} \\ &= \frac{\sum_{n=1}^m p_n (n-1) v_k F^{n-2}(v_k) f(v_k)}{1 - \sum_{l=1}^m p_l F^{l-1}(v_k)},\end{aligned}\tag{3.5}$$

and we state the following result on the symmetric bidding equilibrium.

Proposition 1 *Under the symmetric IPV paradigm, there exists a unique symmetric Bayesian Nash equilibrium in a war of attrition with a stochastic number of bidders with bid function*

$$\beta(v) = \sum_{n=1}^m p_n (n-1) \int_{\underline{v}}^v \frac{t F^{n-2}(t) f(t)}{1 - \sum_{l=1}^m p_l F^{l-1}(t)} dt.\tag{3.6}$$

When $m \geq 2$, $\beta(\cdot)$ is a continuous and strictly increasing function.

Proof. First of all, we know that $\beta(\cdot)$ is a continuous and differentiable function. And then we need to verify the optimality of $\beta(u)$ when bidder k 's value is v . Using equation (3.4) we have

$$\frac{\partial \Pi(\beta(u); v)}{\partial b} = \sum_{n=1}^m p_n v f_{V_{-k}^n}(u) \frac{1}{\beta'(u)} - \left[1 - \sum_{n=1}^m p_n F_{V_{-k}^n}(u) \right]$$

and

$$\beta'(u) = \frac{\sum_{n=1}^m p_n u f_{V_{-k}^n}(u)}{1 - \sum_{l=1}^m p_l F_{V_{-k}^n}(u)}.$$

Plug the latter function into the former, thus we find that

$$\frac{\partial \Pi(\beta(u); v)}{\partial b} = \frac{(v-u)}{\beta'(u)} \sum_{n=1}^m p_n f_{V_{-i}^n}(u).$$

When $v > u$, it follows $\frac{\partial \Pi(\beta(u);v)}{\partial b} > 0$, and $\frac{\partial \Pi(\beta(u);v)}{\partial b} < 0$ if $v < u$. Since $\Pi(\beta(u);v)$ reaches the optimal point at $u = v$, that means bidder k with value v should choose the bidding strategy $\beta(v)$. ■

Though we don't use the closed form of bidding strategies in the estimation procedure, it is important to note that the unique symmetric Bayesian Nash equilibrium exists and the bidding function is continuous and strictly increasing. For Monte Carlo experiments, we need the closed form in the bidding data-generating process.

3.2.3 Unobserved Factors

The bidding strategy is a mapping from a bidder's value to his bid, which is his *expected* "willingness-to-pay". However, in practice, there are a lot of unobserved factors changing bidders' behavior such that bidders would bid more or less. For example, the generalized war of attrition model proposed by Bulow and Klemperer (1999) characterizes the feature that bidders would drop out immediately and pay less than expected. Augenblick (2011) has explored the sunk cost fallacy that bidders would overbid over time in all-pay competition. Bidders' individual heterogeneity also plays important roles on the bids' data-generating process. A sophisticated bidder may understand the optimal bidding strategy better than an unsophisticated bidder.

To address this issue, a potential approach is to provide a more complicated model to identify all of these effects. But even if more information is available, we can never eliminate all unobserved factors entirely. In fact, the perfect structural model does not exist.

The traditional method in econometric analysis is to introduce an error term which contains unobserved factors. In this paper, we use the additive measurement error model to specify the

relationship of observables and unobserved factors, and apply the deconvolution method to identify the distribution of the latent variable. In particular, we distinguish the *planned bid* and the *observed bid*, the former refers to the expected optimal bidding strategy following the structural model, and the latter refers to the bid observed in practice. All unobservable contamination is contained in the error term. Note the unobserved factors characterize some individual heterogeneity, which is different from the unobserved auction heterogeneity examined in literature (see, e.g., Krasnokutskaya 2011).

3.3 Econometric Model

In this section, we develop a nonparametric methodology to identify and estimate the war of attrition with a stochastic number of bidders. More specifically, from observables in a sample of war-of-attrition games, we can identify the distribution of bidders' values. To identify the model, we consider bidders would not follow their planned bids perfectly and the observables are measurements of those planned values with errors. As we have some information about the error term, we can apply the recent developed methods in statistics and econometrics to identify the distribution of the latent primitives.

3.3.1 Identification

The model of the war of attrition is identified if, given the implications of equilibrium behavior in the model, the distribution of bidders' values is uniquely determined by the distribution of observables. Given the symmetric Bayesian Nash equilibrium as the function (3.6), let B_k denote the random variable for the planned bid made by player k , b_k is the realization of the variable, and naturally the ideal observables should be these bids. Let's introduce the distri-

bution $G(\cdot)$ of b_k and its density $g(\cdot)$, and our nonparametric estimation relies upon the fact that the first derivative $\beta'(\cdot)$ and the distribution $F(\cdot)$ with its density $f(\cdot)$ can be eliminated simultaneously from the differential equation (3.5) by using the distribution $G(\cdot)$ and density $g(\cdot)$. The symmetric monotonic equilibrium bidding strategies provide the corresponding relationship between bids and values. Specifically, for every $b \in [\underline{b}, \bar{b}] = [\beta(v), \beta(\bar{v})]$ we have $G(b) = \Pr(\tilde{b} \leq b) = \Pr(\tilde{v} \leq \beta^{-1}(b)) = F(\beta^{-1}(b)) = F(v)$ since $b = \beta(v)$, and it follows that $g(b)\beta'(v) = f(v)$. Therefore the differential equation (3.5) becomes

$$v_k = \xi(b_k) \equiv \frac{1 - \sum_{l=1}^m p_l G^{l-1}(b_k)}{g(b_k) \cdot \sum_{n=1}^m (n-1) p_n G^{n-2}(b_k)}. \quad (3.7)$$

Equation (3.7) gives the individual private value v_k as a function of the bidders' beliefs about competition p_n , the individual's equilibrium bid b_k , its distribution $G(\cdot)$ and density $g(\cdot)$.

In most literature on nonparametric estimation of auctions, the distribution $G(\cdot)$ and density $g(\cdot)$ can be estimated nonparametrically by the empirical distribution and the kernel density estimator using the observations $\{b_1, b_2, \dots\}$ (e.g., Guerre, Perrigne, and Vuong 2000). As is widely known, this is no longer the case when measurement error arises and we need more observables and information to identify the distribution of primitive interests.

Before proceeding, we introduce some notation for our econometric model. Suppose there are T *ex ante* identical auctions, $t = 1, \dots, T$. For auction t , we observe n_t bidders with non-zero bids, therefore the number of observations is $N = \sum_{t=1}^T n_t$. From now on, we use j to denote the observation index, and i is defined as the imaginary number, i.e., $i = \sqrt{-1}$. Note that $j = 1, 2, \dots, N$ refers to observations rather than any bidder's identity k .

Consider the nonparametric estimation of a density from a sample contaminated by random

error. This problem, which is called a *deconvolution problem*, arises very frequently in fields of data application (see, e.g., Li and Vuong 1998; An and Hu 2011; Krasnokutskaya 2011). In particular, the observations are a sample of independent and identically distributed (i.i.d.) variables (Y_j, Z_j) generated by the additive measurement error model

$$Y_j = B_j + \varepsilon_j(Z_j) \tag{3.8}$$

where B_j is the planned bid, which has an unknown density $g(\cdot)$, and the error term is $\varepsilon_j(Z_j)$, follows the density $f_{Z_j}(\cdot)$ which is determined by a variable Z_j . Note that this model allows heteroscedastic contamination, and each error term has its own density function, which may depend on the observation Z_j . Further, we assume that B_j and ε_j are real-valued and independent. The independence of error terms is a standard condition in the deconvolution problem.³

In addition to the existing variable B_j , we require two auxiliary variables:

1. A proxy Y_j , which is a mismeasured version of B_j , e.g., the variable of observed bids, where each Y_j has the common density function $f_Y(\cdot)$.
2. An observation Z_j to characterize the heteroscedastic error term, e.g., the standard deviation of the error term may depend on the duration of participating in competition for each bidder.

Under the independence assumption, the density function $f_Y(\cdot)$ is the convolution of $g(\cdot)$

³The independence assumption is a strong condition. It may be extended to Conditional Independence (CI), i.e., conditional on Z , the random variable B is independent of error term ε . For example, a bidder with a higher bid is more likely to deviate from his planned bid. As bidders change the bidding plans over time, the variance of the error term is positively related to the duration. For bids with same duration, the error term is uncorrelated with B .

and $f_Z(\cdot)$,

$$f_Y(y) = \int g(b) f_Z(y - b) db. \quad (3.9)$$

When both $f_Y(\cdot)$ and $f_Z(\cdot)$ are known, $g(\cdot)$ is recovered by Fourier inversion, thus we refer to the problem of estimating $g(\cdot)$ in the absence of parametric assumptions as *deconvolution*. In the case of homoscedastic errors, Carroll and Hall (1988) and Stefanski and Carroll (1990) proposed the deconvolution kernel density estimator. In recent literature, Delaigle and Meister (2008) and Wang, Fan, and Wang (2010) extended the deconvolution kernel density estimator to heteroscedastic errors. Here we briefly introduce the identification results, and the estimation results are in the following subsection.

For each random variable X , let $\phi_X(t)$ denote its characteristic function, i.e., $\phi_X(t) = Ee^{itX} = \int e^{itx} f_X(x) dx$. From equation (3.8) we have $\phi_{Y_j}(t) = \phi_B(t) \cdot \phi_{\varepsilon(z_j)}(t)$, which can be used to derive $\phi_B(t)$ as a function of series of $\phi_{Y_j}(t)$ and $\phi_{\varepsilon(z_j)}(t)$, for $j = 1, 2, \dots, N$. And then we apply the following Fourier inversion to $\phi_B(t)$,

$$g(b) = \frac{1}{2\pi} \int e^{-itb} \phi_B(t) dt. \quad (3.10)$$

Therefore, under the conditions that the density function $f_Z(\cdot)$ is known and we observe a sample of (Y_j, Z_j) , we have nonparametric estimators for the distribution and density functions of B .

This result is important in the sense that not only does it state the density $g(\cdot)$ is identified by the sample of (Y_j, Z_j) , but it also gives an explicit formula for the nonparametric estimation of the density $g(\cdot)$. Based on the result, we have the following conclusion naturally.

Proposition 2 *Under the symmetric IPV paradigm, suppose that $\xi(b)$ is strictly increasing in b and error density $f_Z(\cdot)$ is known, then the distribution of bidders' values is nonparametrically identified from the sample of (Y, Z) .*

Proposition 2 gives us an opportunity to attenuate the estimation bias, and the knowledge of the distribution of the error term is a crucial challenge in practice. In particular, we may assume density $f_Z(\cdot)$ is a normal distribution with standard deviation that depends on Z . Note that the monotonicity for $\xi(b)$ is corresponding to the decreasing hazard rate of the highest bids distribution. We discuss it as follows.

Equation (3.7) implies

$$\xi(\cdot) \equiv \frac{\sum_{l=1}^m p_l (1 - G^{l-1}(\cdot))}{\sum_{n=1}^m p_n (n-1) G^{n-2}(\cdot) g(\cdot)} = \frac{\sum_{l=1}^m p_l (1 - \Lambda(\cdot))}{\sum_{n=1}^m p_n \Lambda'(\cdot)}$$

where $\Lambda(\cdot) = G^{n-1}(\cdot)$ and $\Lambda'(\cdot) = (n-1)G^{n-2}(\cdot)g(\cdot)$ corresponds to the distribution and density function of the highest bid. Thus, if the number of bids is fixed and known, we have $\xi(\cdot) = \frac{(1-\Lambda(\cdot))}{\Lambda'(\cdot)}$, and the monotonicity for $\xi(\cdot)$ requires that the hazard rate function of distribution $\Lambda(\cdot)$ is decreasing.

3.3.2 Nonparametric Estimation

The procedure of our nonparametric estimation is straightforward: if we know the distribution of Y , given the distribution function for the error term, then we could estimate nonparametrically the distribution $G(\cdot)$ and its density $g(\cdot)$. And then we could use the inverse of bidding strategy (3.7) to estimate the distribution $F(\cdot)$. Hence we propose a constructive two-step procedure: in the first step, we consider a nonparametric deconvolution problem with heteroscedastic errors;

in the second step, we apply the established methodology for nonparametric estimation to auctions. For most of this paper, we focus on the first step of the procedure, since the second step is a well-established technique.

As we observe the numbers of bidders for T auctions, let n_{\max} be the maximum number of bidders present in T auctions. We defined m as the maximum number of potential bidders, thus $n_{\max} \leq m$, and the natural estimator for γ_n is defined by the empirical probability that n bidders are present,

$$\hat{\gamma}_n = \frac{\#\{t : n_t = n\}}{T} \quad \text{for } n = 1, \dots, \infty, \quad (3.11)$$

which is a consistent estimator by the strong law of large numbers, $\hat{\gamma}_n \rightarrow \gamma_n$ for every n , as $T \rightarrow \infty$, and then the empirical estimator for p_n is defined simply from (3.1) as,

$$\hat{p}_n = \frac{n\hat{\gamma}_n}{\sum_{l=1}^m l\hat{\gamma}_l} \quad \text{for } n = 1, \dots, n_{\max}. \quad (3.12)$$

It is well defined and a consistent estimator of p_n provided that $\sum_{l=1}^m l\hat{\gamma}_l \neq 0$ by the continuous mapping theorem.

Now, given the error term density $f_Z(\cdot)$ is known and we observe a sample of (Y_j, Z_j) , using (3.10) and $\phi_Y(t) = \phi_B(t) \cdot \phi_{\varepsilon(z)}(t)$, Delaigle and Meister (2008) propose a deconvolution estimator for the density with heteroscedastic errors, which can be written as a form of a kernel-type density estimator,

$$\hat{g}_N(b) = \frac{1}{Nh_N} \sum_{j=1}^N \tilde{K}_j \left(\frac{b - Y_j}{h_N} \right), \quad (3.13)$$

where

$$\tilde{K}_j(x) = \frac{1}{2\pi} \int e^{-itx} \phi_K(t) \psi_{\varepsilon(z_j)}(t/h_N) dt, \quad \psi_{\varepsilon(z_j)}(t) = \frac{\phi_{\varepsilon(z_j)}(-t)}{\frac{1}{N} \sum_{l=1}^N |\phi_{\varepsilon(z_l)}(t)|^2}, \quad (3.14)$$

and ϕ_K is the characteristic function of a symmetric probability kernel, $K(\cdot)$, with a finite variance, h_N is a bandwidth with $h_N \rightarrow 0$ and $Nh_N \rightarrow \infty$ as $N \rightarrow \infty$. In this context, we use the second-order kernel⁴

$$K(x) = \frac{48 \cos(x)}{\pi x^4} \left(1 - \frac{15}{x^2}\right) - \frac{144 \sin(x)}{\pi x^5} \left(2 - \frac{5}{x^2}\right). \quad (3.15)$$

Delaigle and Hall (2006) recommend to use this kernel based on the performance of numerical simulations; it is a commonly used kernel in the deconvolution problems (see Delaigle and Gijbels 2006b; Delaigle and Hall 2006; Bonhomme and Robin 2010), and it corresponds to the characteristic function

$$\phi_K(t) = (1 - t^2)^3 \cdot 1_{\{t \in [-1, 1]\}}. \quad (3.16)$$

If we assume the following conditions, the estimator $\hat{g}_N(b)$ is well defined. These conditions are standard in deconvolution problems.

Condition A1 There exists some j such that $|\phi_{\varepsilon(z_j)}(t)| \neq 0$ for all $t \in \mathbb{R}$;

Condition A2 $\phi_K(t)$ is bounded, continuous at $t = 0$ and $\phi_K(0) = 1$;

Condition A3 $\frac{\phi_{\varepsilon(z_j)}(-t)\phi_K(th_N)}{\sum_{l=1}^N |\phi_{\varepsilon(z_l)}(t)|^2} \in L_2(\mathbb{R})$, that is, quadratically integrable functions.

⁴A second-order kernel satisfies $\int K(t) dt = 1$, $\int tK(t) dt = 0$, and $\int t^2K(t) dt = 0$.

Meanwhile the distribution estimator $\widehat{G}_N(b)$ is defined as simply the integral of $\widehat{g}_N(\cdot)$ over $(-\infty, b]$. Following Wang, Fan, and Wang (2010), we have the distribution estimator as a form of the kernel-type,

$$\widehat{G}_N(b) = \int_{-\infty}^b \widehat{g}_N(t) dt = \frac{1}{N} \sum_{j=1}^N \widetilde{L}_j(b - Y_j), \quad (3.17)$$

where

$$\widetilde{L}_j(x) = \frac{1}{2} + \frac{1}{2\pi} \int \frac{\sin(tx) \phi_K(th_N) \psi_{\varepsilon(z_j)}(t)}{t} dt, \quad \text{for } j = 1, \dots, N. \quad (3.18)$$

To estimate the unknown support of bids $[\underline{b}, \bar{b}] = [\beta(\underline{v}), \beta(\bar{v})]$, we first recall that Proposition 1 implies the left endpoint $\underline{b} = \beta(\underline{v}) = 0$. For the right endpoint \bar{b} , the largest observation $\sup\{B_j\}$ is a consistent estimator if the data of bids $\{B_j\}$ are observed without error. In the case of measurement error, however, the simple estimator $\widehat{Y}_N = \sup\{Y_j : j = 1, \dots, N\}$ is not consistent for \bar{b} . Indeed, under the assumption of mean zero for the error term, the simple estimator converges to some value beyond the support of $[\underline{b}, \bar{b}]$, i.e., $\lim \widehat{Y}_N \geq \bar{b}$. Delaigle and Gijbels (2006a, 2006b) studied the boundary estimation in deconvolution problems, and the performance of the consistent estimators in practice *strongly* depends on the choice of the bandwidth. For simplicity, we consider the support of bids in the compact set of interval $[0, \widehat{Y}_N]$.

Under quite general conditions, $\widehat{G}_N(b)$ and $\widehat{g}_N(b)$ are L_2 -uniformly consistent estimators of the distribution and density of planned bids B (Delaigle and Meister 2008, Wang, Fan, and Wang 2010). Due to the boundary effects of kernel estimator $\widehat{g}_N(\cdot)$, $\frac{1}{\widehat{g}_N(\cdot)}$ is an asymptotically biased estimate at the boundaries. To overcome this problem, we introduce the trimming method proposed by Guerre, Perrigne, and Vuong (2000), and thus, using (3.7) leads to define the pseudo private value \widehat{V}_j corresponding to any value b_j in the inner compact subset of the

support,

$$\widehat{V}_j = \widehat{\xi}_N(b_j) = \begin{cases} \frac{1 - \sum_{l=1}^m \widehat{p}_l \widehat{G}_N^{l-1}(b_j)}{\widehat{g}_N(b_j) \cdot \sum_{n=1}^m (n-1) \widehat{p}_n \widehat{G}_N^{n-2}(b_j)} & \text{if } \frac{\rho h_N}{2} \leq b_j \leq \widehat{Y}_N - \frac{\rho h_N}{2} \\ +\infty & \text{otherwise} \end{cases}, \quad (3.19)$$

for $j = 1, \dots, N$, with $\rho < \infty$ is the length of the support of $K(\cdot)$. While this inverse of the bidding strategy is a strictly increasing function, the value distribution $F(\cdot)$ can be estimated by

$$\widehat{F}_N(\cdot) = \widehat{G}_N(\widehat{\xi}_N^{-1}(\cdot)) \quad (3.20)$$

where the support of $F(\cdot)$ is estimated as $[\widehat{v}, \widehat{v}] = [\widehat{\xi}_N(\frac{\rho h_N}{2}), \widehat{\xi}_N(\widehat{Y}_N - \frac{\rho h_N}{2})]$, respectively⁵.

As we know, $\xi(b)$ is the inverse bidding function, thus $\xi^{-1}(v) = \beta(v)$ is the (quasi-) bidding function. Here $\widehat{\xi}_N^{-1}(\cdot)$ is the estimator of the (quasi-) bidding function, and it is defined as an extremum estimator

$$\widehat{\xi}_N^{-1}(v) = \arg \min_b \left\{ \left(\widehat{\xi}_N(b) - v \right)^2 \right\} \quad (3.21)$$

for any value $v \in [\widehat{v}, \widehat{v}]$ suggested by Newey and McFadden (1994). When N is big enough, $\widehat{\xi}_N(\cdot)$ is a continuous function in $[\frac{\rho h_N}{2}, \widehat{Y}_N - \frac{\rho h_N}{2}]$, thus, there exists a point $\widehat{b} = \widehat{\xi}_N^{-1}(v)$ such that $\widehat{\xi}_N(\widehat{b}) = v$.

3.3.3 Asymptotic Properties of Estimators

This subsection summarizes properties of the proposed estimators, and then provides the details of the proofs. Given the estimators in the previous subsection, we need to show the uniform

⁵Note the estimators of the support of private values are not consistent. We focus on the property of inner compact subset of the support. Since we cannot observe the value of B_j , we estimate the value distribution by (3.20) rather than the empirical CDF.

consistency of estimators of the bid distribution $\widehat{G}_N(b)$, the bid density $\widehat{g}_N(b)$, the pseudo-value estimator function $\widehat{\xi}_N(\cdot)$, and the estimator for value distribution $\widehat{F}_N(\cdot)$.

In this paper, we consider auctioned objects that are homogenous and bidders' private values that are independent within or across auctions. Throughout, we work under the assumptions that $g(\cdot)$ is continuous and bounded and hence square integrable. To evaluate the quality of the estimator more precisely, we study the convergence rates and uniform consistency. Here we need some regularity assumptions to obtain the upper bound and lower bound of the rates of convergence of the estimator. Some definitions are provided as follows.

Definition 1 Let $\mathcal{F}_{\lambda,C}$ denote the class of densities uniformly bounded relative to their Sobolev (λ -) norm. That is, the density of B , $g(\cdot) \in \mathcal{F}_{\lambda,C}$, such that

$$\int |\phi_B(t)|^2 (1+t^2)^\lambda dt \leq C.$$

Definition 2 (Ordinary smooth density) The error term $\varepsilon(Z_j)$ follows an *ordinary smooth* error density $f_{Z_j}(\cdot)$, that is

$$C_1 (1+|t|)^{-\alpha} \leq \left| \phi_{\varepsilon(z_j)}(t) \right| \leq C_2 (1+|t|)^{-\alpha}, \text{ for all } t \in \mathbb{R},$$

for some $C_2 > C_1 > 0$ and $\alpha > 0$.

Definition 3 (Supersmooth density) The error term $\varepsilon(Z_j)$ follows a *supersmooth* error density $f_{Z_j}(\cdot)$, that is

$$C_1 \exp(-d_1 |t|^\gamma) \leq \left| \phi_{\varepsilon(z_j)}(t) \right| \leq C_2 \exp(-d_2 |t|^\gamma), \text{ for all } t \in \mathbb{R},$$

for some $C_2 > C_1 > 0, 0 < d_2 < d_1, \gamma > 0$.

Remark: Examples of ordinary smooth distributions are uniform, gamma, symmetric gamma, double exponential, etc. Examples of supersmooth distributions are normal, Cauchy, mixture normal, etc.

Condition A4 $|\phi_K(t)| \leq 1$ for all t , ϕ_K is supported on $[-1, 1]$ and $|\phi_K(t) - 1| = o(|t|^\lambda)$ with order λ .

Proposition 3 *Under conditions A1-A4, if the error term density $f_{Z_j}(\cdot)$ is ordinary smooth or supersmooth, then $\hat{g}_N(b)$ is a uniformly consistent estimator of $g(b)$.*

Proof. From the Fourier inversion formula,

$$g(b) = \frac{1}{2\pi} \int e^{-itb} \phi_B(t) dt,$$

and the kernel-type density estimator (3.13) can be written as

$$\hat{g}_N(b) = \frac{1}{2\pi} \int e^{-itb} \hat{\phi}_B(t) dt,$$

where

$$\hat{\phi}_B(t) = \frac{1}{N} \sum_{j=1}^N e^{itY_j} \phi_K(th_N) \psi_{\varepsilon(z_j)}(t).$$

The (3.8) implies $\phi_{Y_j}(t) = \phi_B(t) \cdot \phi_{\varepsilon(z_j)}(t)$, and we can have

$$\phi_B(t) = \sum_{j=1}^N \phi_{Y_j}(t) \frac{\phi_{\varepsilon(z_j)}(-t)}{\sum_{l=1}^N |\phi_{\varepsilon(z_l)}(t)|^2} = \frac{1}{N} \sum_{j=1}^N \phi_{Y_j}(t) \psi_{\varepsilon(z_j)}(t).$$

For each positive integer M

$$\begin{aligned}
E \left(\frac{1}{2\pi} \int_{-M}^M \left| \widehat{\phi}_B(t) - \phi_B(t) \right| dt \right) &= \frac{1}{2\pi} \int_{-M}^M E \left| \widehat{\phi}_B(t) - \phi_B(t) \right| dt \\
&\leq \frac{1}{2\pi} \int_{-M}^M |\phi_K(th_N)| \left(E \left| \frac{1}{N} \sum_{j=1}^N (e^{itY_j} - \phi_{Y_j}(t)) \psi_{\varepsilon(z_j)}(t) \right|^2 \right)^{1/2} dt \\
&\quad + \frac{1}{2\pi} \int_{-M}^M |\phi_B(t)| |\phi_K(th_N) - 1| dt \\
&\leq \frac{1}{2\pi} \int_{-M}^M |\phi_K(th_N)| \frac{\sqrt{2}}{\sqrt{N}} \max_j |\psi_{\varepsilon(z_j)}(t)| \\
&\quad + \frac{1}{2\pi} \int_{-M}^M |\phi_B(t)| |\phi_K(th_N) - 1| dt \\
&\leq \frac{M}{\pi} \frac{\sqrt{2}}{\sqrt{N}} \max_{j, |t| \leq M} |\psi_{\varepsilon(z_j)}(t)| + \frac{1}{2\pi} \int_{-M}^M |\phi_B(t)| |\phi_K(th_N) - 1| dt.
\end{aligned}$$

From conditions A1-A4 and $\int |\phi_B(t)| dt < \infty$, it follows that a sequence $M (= M_N)$ can be chosen so that

$$E \left(\frac{1}{2\pi} \int_{-M}^M \left| \widehat{\phi}_B(t) - \phi_B(t) \right| dt \right) \rightarrow 0.$$

Hence

$$\begin{aligned}
|\widehat{g}_N(b) - g(b)| &= \left| \frac{1}{2\pi} \int e^{-itb} \widehat{\phi}_B(t) dt - \frac{1}{2\pi} \int e^{-itb} \phi_B(t) dt \right| \\
&\leq \frac{1}{2\pi} \int_{-M}^M |e^{-itb}| \left| \widehat{\phi}_B(t) - \phi_B(t) \right| dt + \frac{1}{2\pi} \int_{|t|>M} |e^{-itb}| |\phi_B(t)| dt \\
&= \frac{1}{2\pi} \int_{-M}^M \left| \widehat{\phi}_B(t) - \phi_B(t) \right| dt + \frac{1}{2\pi} \int_{|t|>M} |\phi_B(t)| dt
\end{aligned}$$

is independent of $b \in \mathbb{R}$. Since $\int_{|t|>M} |\phi_B(t)| dt \rightarrow 0$ as $M \rightarrow \infty$, it follows that

$$E \left(\sup_b |\widehat{g}_N(b) - g(b)| \right) \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

i.e.

$$\sup_b |\widehat{g}_N(b) - g(b)| = O(d_N) \quad \text{a.s.}$$

where d_N denote the rate of convergence⁶. Thus the estimate $\widehat{g}_N(b)$ is a uniformly consistent estimator of $g(b)$ (see Theorem 3.1 in Liu and Taylor 1989). ■

Note that $G(b) = \int_{-\infty}^b g(t) dt$ and $\widehat{G}_N(b) = \int_{-\infty}^b \widehat{g}_N(t) dt$ imply that $|\widehat{G}_N(b) - G(b)| \leq \int_{-\infty}^b |\widehat{g}_N(t) - g(t)| dt \leq \sup_t |\widehat{g}_N(t) - g(t)| \cdot (b - \underline{b})$. The support of $g(\cdot)$ is always finite, indeed $0 = \underline{b} \leq \bar{b} < \infty$, and then the uniform consistency of estimator $\widehat{G}_N(b)$ holds under the same conditions.

Proposition 4 *Suppose the conditions in Proposition 3 hold. Then, $\widehat{G}_N(b)$ defined by (3.17) is a uniformly consistent estimator of $G(b)$ provided $N \rightarrow \infty$, $h_N \rightarrow 0$ and $Nh_N \rightarrow \infty$.*

Now we proceed to analyze the conditions of the boundary estimators. Since the left end point $\underline{b} = 0$ is well-known, the asymptotic property for the estimator of the right end point is summarized in the following proposition.

Proposition 5 *The consistent estimator of the right end point \bar{b} for $g(\cdot)$, \widehat{b}_N , is in the interval of $[0, \widehat{Y}_N]$. That is*

$$P\left(0 \leq \lim \widehat{b}_N \leq \lim \widehat{Y}_N\right) = 1.$$

Remark: The result suggests that the estimate of the right end point is not well-constructed. In this situation, our consistency for all estimators is limited in a small subset of the interior compact set. In the following content, we denote $\Delta(B)$ as the interior compact subset, and

⁶Li and Vuong (1998) established the rate of convergence for four different cases.

$\Delta(B) \subset [0, \bar{b}] \subset [0, \widehat{Y}_N]$. In particular, the left end point of $\Delta(B)$ is 0, and the right end point of $\Delta(B)$ is bounded by \widehat{Y}_N .

In the next step, we analyze the uniform consistency of the pseudo-value estimator function $\widehat{\xi}_N(\cdot)$. In order to recover private values, it is convenient that $g(\cdot)$ be bounded away from zero in its support. The proofs of Proposition 6 and 7 follow the procedures in Krasnokutskaya (2011, p. 323-324).

Proposition 6 *Suppose the conditions A1-A4 hold and $g(b) \geq \bar{c}_g > 0$ for all bids b . Then*

$\widehat{\xi}_N(\cdot)$ defined by (3.19) is a uniformly consistent estimator of $\xi(\cdot)$ provided $N \rightarrow \infty$, $h_N \rightarrow 0$ and $Nh_N \rightarrow \infty$.

Proof. Since we have proved the uniform convergence, we denote the rate of convergence for $\widehat{g}_N(b)$ as d_N , and then $\widehat{G}_N(b)$ converges to $G(b)$ at the same rate of d_N . Then we prove the uniform consistency of the estimator for the individual inverse bid function $\widehat{\xi}_N(\cdot)$. By Equation (7), we know

$$v = \xi(b) = \frac{1 - \sum_{l=1}^m p_l G^{l-1}(b)}{g(b) \cdot \sum_{n=1}^m (n-1) p_n G^{n-2}(b)},$$

which is the inverse bidding function. The proposed estimator $\widehat{\xi}_N(\cdot)$ is constructed as follows

$$\widehat{\xi}_N(b) = \frac{1 - \sum_{l=1}^m \widehat{p}_l \widehat{G}_N^{l-1}(b)}{\widehat{g}_N(b) \cdot \sum_{n=1}^m (n-1) \widehat{p}_n \widehat{G}_N^{n-2}(b)}$$

in an interior compact subset $\Delta(B) \subset [0, \bar{b}] \subset [0, \widehat{Y}_N]$. Where \widehat{Y}_N is well-defined as $\widehat{Y}_N = \sup\{Y_j : j = 1, \dots, N\}$. Up to now, we have proved $\sup_{b \in \Delta(B)} |\widehat{g}_N(b) - g(b)| = O(d_N)$ *a.s.* from Proposition 3 and $\sup_{b \in \Delta(B)} |\widehat{G}_N(b) - G(b)| = O(d_N)$ *a.s.* from Proposition 4. And we know \widehat{p}_n is a consistent estimator of p_n for every n . Here we assume $g(b) \geq \bar{c}_g > 0$ for every

$b \in \Delta(B)$. In the interior compact subset $\Delta(B)$, we have $G(b) \geq \bar{c}_G > 0$ for some \bar{c}_G . Note that $\hat{g}_N(b) \geq c_g > 0$ and $\hat{G}_N(b) \geq c_G > 0$ for some $c_g \in (0, \bar{c}_g]$ and $c_G \in (0, \bar{c}_G]$ since \hat{g}_N and \hat{G}_N uniformly converge to g and G , respectively. We define the following notation:

$$\chi(b) = 1 - \sum_{l=1}^m p_l G^{l-1}(b),$$

$$\hat{\chi}_N(b) = 1 - \sum_{l=1}^m \hat{p}_l \hat{G}_N^{l-1}(b),$$

$$\kappa(b) = g(b) \cdot \sum_{n=1}^m (n-1) p_n G^{n-2}(b),$$

$$\hat{\kappa}_N(b) = \hat{g}_N(b) \cdot \sum_{n=1}^m (n-1) \hat{p}_n \hat{G}_N^{n-2}(b).$$

So we have $\xi(b) = \frac{\chi(b)}{\kappa(b)}$ and $\hat{\xi}_N(b) = \frac{\hat{\chi}_N(b)}{\hat{\kappa}_N(b)}$. Then $\hat{\xi}_N(b) - \xi(b) = \frac{\hat{\chi}_N(b)\kappa(b)}{\hat{\kappa}_N(b)\kappa(b)} - \frac{\hat{\kappa}_N(b)\chi(b)}{\hat{\kappa}_N(b)\kappa(b)} = \frac{\hat{\chi}_N(b)\kappa(b) - \hat{\kappa}_N(b)\chi(b)}{\hat{\kappa}_N(b)\kappa(b)}$ implies

$$\left| \hat{\xi}_N(b) - \xi(b) \right| = \frac{|\hat{\chi}_N(b)\kappa(b) - \hat{\kappa}_N(b)\chi(b)|}{\hat{\kappa}_N(b)\kappa(b)}.$$

Note $\kappa(b) = g(b) \cdot \sum_{n=1}^m (n-1) p_n G^{n-2}(b) \geq \bar{c}_g \cdot \sum_{n=1}^m (n-1) p_n \bar{c}_G^{n-2} > 0$, $\hat{\kappa}_N(b) = \hat{g}_N(b) \cdot \sum_{n=1}^m (n-1) \hat{p}_n \hat{G}_N^{n-2}(b) \geq c_g \cdot \sum_{n=1}^m (n-1) \hat{p}_n c_G^{n-2} > 0$. Since the lower bounds $\bar{c}_g \cdot \sum_{n=1}^m (n-1) p_n \bar{c}_G^{n-2}$ and $c_g \cdot \sum_{n=1}^m (n-1) \hat{p}_n c_G^{n-2}$ do not depend on the variable b , we can find a constant $C_1 > 0$,

such that

$$\hat{\kappa}_N(b)\kappa(b) \geq C_1.$$

Then,

$$\begin{aligned} \left| \widehat{\xi}_N(b) - \xi(b) \right| &\leq \frac{1}{C_1} |\widehat{\chi}_N(b) \kappa(b) - \widehat{\kappa}_N(b) \chi(b)| \\ &\leq \frac{1}{C_1} (|\widehat{\chi}_N(b) - \chi(b)| |\kappa(b)| + |\chi(b)| |\kappa(b) - \widehat{\kappa}_N(b)|). \end{aligned}$$

Let's define $\widetilde{c}_g = \max_{b \in \Delta(B)} \{g(b)\}$. It is well defined because g is continuous function and the interior subset is a compact set. Then $\kappa(b) \leq \widetilde{c}_g \cdot \sum_{n=1}^m (n-1) p_n \leq m \widetilde{c}_g$ and $\chi(b) \leq 1$. Thus we have

$$\left| \widehat{\xi}_N(b) - \xi(b) \right| \leq \frac{m \widetilde{c}_g}{C_1} |\widehat{\chi}_N(b) - \chi(b)| + \frac{1}{C_1} |\kappa(b) - \widehat{\kappa}_N(b)|.$$

Pointwise application of the delta method and uniform convergence of \widehat{g}_N and \widehat{G}_N to g and G respectively allows us to conclude that

$$\begin{aligned} \sup_{b \in \Delta(B)} |\widehat{\chi}_N(b) - \chi(b)| &= O(d_n), \text{ a.s.} \\ \sup_{b \in \Delta(B)} |\kappa(b) - \widehat{\kappa}_N(b)| &= O(d_n), \text{ a.s.} \\ \sup_{b \in \Delta(B)} \left| \widehat{\xi}_N(b) - \xi(b) \right| &= O(d_n), \text{ a.s..} \end{aligned}$$

■

Proposition 7 *Suppose the conditions A1-3 hold and $g(\cdot)$ is continuous and bounded. Then,*

$\widehat{F}_N(v)$ defined by (3.20) is a uniformly consistent estimator of $F(v)$ for all $v \in \Delta(V)$ provided $N \rightarrow \infty$, $h_N \rightarrow 0$ and $Nh_N \rightarrow \infty$.

Proof.

$$\begin{aligned}
\left| \widehat{F}_N(v) - F(v) \right| &= \left| \widehat{G}_N \left(\widehat{\xi}_N^{-1}(v) \right) - G \left(\xi^{-1}(v) \right) \right| \\
&\leq \left| \widehat{G}_N \left(\widehat{\xi}_N^{-1}(v) \right) - G \left(\widehat{\xi}_N^{-1}(v) \right) \right| + \left| G \left(\widehat{\xi}_N^{-1}(v) \right) - G \left(\xi^{-1}(v) \right) \right| \\
&= \left| \widehat{G}_N \left(\widehat{\xi}_N^{-1}(v) \right) - G \left(\widehat{\xi}_N^{-1}(v) \right) \right| + |g(\theta)| \cdot \left| \widehat{\xi}_N^{-1}(v) - \xi^{-1}(v) \right|.
\end{aligned}$$

where $\theta \in (0, \bar{b})$. As we know, the estimator of the (quasi-) bidding function $\widehat{\xi}_N^{-1}(v)$ is defined as the solution to $\widehat{\xi}_N(b) = v$ which can be obtained by solving

$$\widehat{\xi}_N^{-1}(v) = \arg \min_b \left\{ \left(\widehat{\xi}_N(b) - v \right)^2 \right\}$$

for any value $v \in [\widehat{v}, \bar{v}]$. Thus, in the interior subset $\Delta(V) \subset [\widehat{v}, \bar{v}]$, we need to establish the uniform convergence of the (quasi-) bidding function estimation. For a given $v \in \Delta(V)$, let $b_0 = \xi^{-1}(v)$ and $b_N = \widehat{\xi}_N^{-1}(v)$. Here b_0 is some number from $(0, \bar{b})$ and b_N is a random variable with realizations in $(0, \bar{b})$ for large N . Note that $\widehat{\xi}_N(\cdot)$ is a continuous function, and for every point $v \in [\widehat{v}, \bar{v}]$ there exists a point $b \in (0, \bar{b})$ such that $\widehat{\xi}_N(b) = v$. For every realization of b_N , there is a number b_N^* such that

$$\xi(b_0) - \xi(b_N) = \xi'(b_N^*)(b_0 - b_N), \quad b_N^* \in [b_0, b_N],$$

since $\xi(\cdot)$ is continuously differentiable on the interior compact set. Let us also denote by b_N^* a random variable with realizations as above. Note that we allow $b_0 \leq b_N$ or $b_0 \geq b_N$. If b_0, b_N always belong to the interior of $\Delta(B)$, then b_N^* also always belongs to the interior of $\Delta(B)$. Since the inverse bidding function is strictly increasing on the compact set, then

$\|\xi'(b_N^*)\| \geq c_\xi > 0$ for some constant number c_ξ , and therefore,

$$\|b_0 - b_N\| \leq \frac{1}{c_\xi} \|\xi(b_0) - \xi(b_N)\|.$$

On the other hand, $v = \xi(b_0) = \widehat{\xi}_N(b_N)$ implies

$$\xi(b_0) - \xi(b_N) = \widehat{\xi}_N(b_N) - \xi(b_N).$$

Since, as we have shown above, $\widehat{\xi}_N(\cdot)$ converges uniformly to $\xi(\cdot)$, then

$$\|\xi(b_0) - \xi(b_N)\| = \left\| \widehat{\xi}_N(b_N) - \xi(b_N) \right\| = O(d_N) \text{ a.s.}$$

and thus we have

$$\|b_0 - b_N\| = \left\| \widehat{\xi}_N^{-1}(v) - \xi^{-1}(v) \right\| = O(d_N) \text{ a.s.}$$

by approach for optimization estimators (see Newey and McFadden 1994). Note here $\|\cdot\| = \sup|\cdot|$. That means we have proved the uniform consistency for estimator $\widehat{\xi}_N^{-1}(\cdot)$. Therefore, the uniform convergence of $\widehat{G}_N(\cdot)$ and the bounded property of $g(\cdot)$ obtain

$$\sup_{v \in \Delta(V)} \left| \widehat{F}_N(v) - F(v) \right| = O(d_N) \text{ a.s.}$$

■

3.4 Monte Carlo Experiments

In this section, we report the results from some Monte Carlo experiments for the two-step nonparametric procedure. Suppose $m = 5$ potential bidders participate in $T = 200$ homogenous war-of-attrition games. The number of active bidders in each individual game is equally likely to be 1 through 5; that is, $\gamma_n = \frac{1}{5}$ for $n = 1, \dots, 5$, thus the expected number of observed bids for each game is 3, and the expected number of observations is 500 in total. As in Guerre, Perrigne, and Vuong (2000), we simulate for 1000 replications, and the true distribution of values $F(\cdot)$ is log-normal with parameters zero and one, truncated in $[0.055, 2.5]$.

For each replication, we first generate randomly a series of numbers of bidders $\{n_1, \dots, n_T\}$, and then draw $\sum_{t=1}^T n_t = N$ private values from the truncated log-normal distribution. Following the symmetric Bayesian Nash equilibrium, we compute numerically the corresponding (planned) bids B_j using (3.6) for each private value. Second, since B_j and Z_j are highly correlated, and the correlation has no effect on estimation based on the assumption of error term independence, we choose the duration Z_j as a function of B_j , for simplicity, as $Z_j = B_j$. Here we assume the error term follows normal distribution $\varepsilon(Z_j) \sim N\left(0, [0.5 + Z_j/\max(Z_j)]^2\right)$. Third, we compute numerically the corresponding observed bids Y_j using (3.8). Note the value of Y_j may be negative in this data-generating process.

Given the sample of (Y, Z) , we apply our estimation procedure for each replication. First, we estimate the distribution and density function of planned bids using (3.13) and (3.17). Second, we compute the pseudo-value function (3.19) and then estimate the distribution function of private values by (3.20). Delaigle and Gijbels (2004) compared several plug-in bandwidth selectors with the cross-validation (CV) bandwidth selector and the bootstrap bandwidth selec-

tor, and Wang and Wang (2011) generalized the plug-in and the bootstrap bandwidth selection methods to the case of heteroscedastic errors and they also provided an R package *decon* on the bandwidth selection for practical use. In this paper, we choose the bootstrap method with real resampling generalized by Wang and Wang (2011).

To evaluate performance, we report the Monte Carlo results in Figure 1 and 2. Figure 1(a) displays the estimator of bid distribution and Figure 1(b) for the estimator of density. From these Figures, the estimators of distribution and density tend to be biased to one-side before we correct the measurement errors. As we correct the impact of errors, the results converge toward the true values just as we have expected. Figure 2 displays the true value distribution and its estimator. In the interior compact subset on the left-hand side, the estimated distribution converges uniformly to the true value, and the poor convergence on the right end point is due to the boundary condition and the inconsistency of the estimate in Proposition 5.

Figure 1(a): True and estimated bids distribution




Figure 1(b): True and estimated bids density

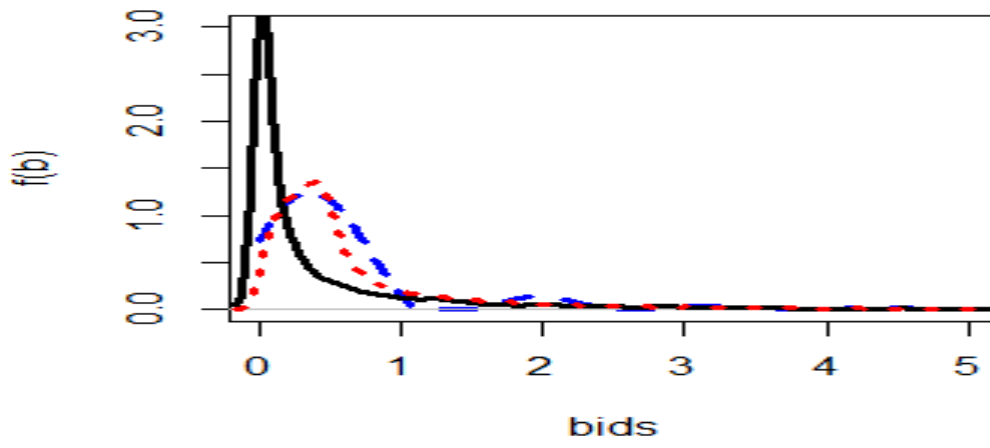



Figure 2: True and estimated private value distribution



3.5 Empirical Application

In this section, we apply the proposed two-step estimation procedure to the data from penny auctions, and get some empirical inferences from field data source.

3.5.1 Penny auctions

The penny auction is an online retail selling mechanism that has emerged very recently. *Swoopo*, founded in 2005, was the pioneer in this industry. By November 2010, at least 125 penny auction websites were running in the U.S. Some recent papers examine the bidders' behavior and the auctioneer's revenue in penny auctions. For example, Augenblick (2011) and Platt, Price, and Tappen (2010) are impressed by the extremely high observed revenue from Swoopo's data, and they provide some explanations on the irrational bidding behavior. Anderson and Ødegaard (2011) examine the penny auction under a setting with two sales channels: a fixed posted price and a standard all-pay auction. In contrast to eBay auctions, the penny auction is a contest game with elements of all-pay auctions and lotteries. In this paper, we use the data from BigDeal.com, which was a typical penny auction website operating from November 2009 to August 2011. More details on penny auctions and BigDeal are in Wang and Xu (2012). The rules in most of penny auction websites are similar to the following.

Prior to participating in any auction, bidders must buy packs of bid-tokens. Each bid-token costs a fixed price (e.g., \$0.75). The website typically releases an auction with an initial countdown clock that lasts for 1-2 days. Potential buyers can access the website to view the description of the product. The posted retail price is also available, which is usually close to the market price. And the auction price for any product starts at \$0.

First, a bidder must place a bid-token to bid. A bid is an offer to buy the product at the current auction price. The auction price is initially 0 and is increased by a fixed tiny amount whenever a bid is placed. The tiny increment is typically one penny, thus the name of penny auction.

Second, as a bid is placed, a countdown clock starts and the timer lasts for a fixed length of time (e.g., 30 seconds).

Third, the other bidders decide whether to follow before the clock expires. And the timer is reset whenever a new bid is placed.

The winner is the *last* bidder, the person whose bid is not followed by any other bid by the time the countdown clock expires. The auction winner receives the product and pays the auction price. Since each bid-token cost is not refundable, every bidder actually pays the bidding cost.

3.5.2 Planned bids *vs.* observed bids

To analyze the mechanism of the penny auction, Augenblick (2011) provides a basic analysis framework within the common value paradigm. All other existing literature follows his perspective. A penny auction is not a standard auction because of the rule that the last rather than the highest bidder wins. Under this framework, most of papers focus on bidders' sophistication and bidding behavior. Now we propose a new view under the independent private value paradigm, in which the penny auction is a war of attrition with a stochastic number of bidders.

We propose the following mechanism in the data generating process for penny auctions:

Bid-agent game: Most penny auction websites offer automatic bid-agents to help bidders to place bids. A bidder can set up with a maximum number of bids to place by a bid-agent. The bid-agent will start to bid as soon as the bidder click the “Activate” button. When the

bid-agent places a bid, the countdown clock resets to 30 seconds. If another bid follows during the time period, the bid-agent will place another bid until the bids run out. While more than two bid-agents are activated, they will bid in turn. Suppose that all bidders use bid-agents, and they set their total bids at the beginning and activate the bid-agents simultaneously. No new bidder enters in the midst, and bidders cannot modify the bidding plans. As the game starts, bid-agents automatically place bids one-by-one, and the last bid wins. In each round, each bidder gets the chance to place a bid. Thus the winner, the *last* bidder, is also the *highest* bidder. This framework is a second-price sealed-bid all-pay auction.

Since the increment is tiny, assume we can ignore the effect of the auction price. By the bid-agent game, each bidder’s bidding plan is predicted as the planned bid B , which is unobserved. The observed bid Y is a proxy of the planned bid, and the relation is the following, i.e., as equation (3.8)

$$Y = B + \varepsilon,$$

with $\varepsilon \sim N(0, \sigma_Z^2)$, where σ_Z denotes the standard deviation of the error term, Z is bidders’ duration in the game, and we assume $\sigma_Z = \log(1 + Z)$ in the empirical analysis.

3.5.3 Data Structure

In this paper, we use the data from a penny auction site, BigDeal.com. There are 1,686 different products auctioned in BigDeal, and we choose 4 pairs from the top 20 products to analyze and compare: 10 BigDeal Bid Tokens *vs.* 50 BigDeal Bid Tokens; \$25 Gift Card *vs.* \$100 Gift Card; iPod Nano 16GB *vs.* iPad 3G 64GB; Amazon Kindle *vs.* Kindle DX. Table 1 summarizes the statistics of them. The number of bids for each product is reasonable to achieve the convergence

results.

Table 3.1: Data Description

Product	Retail price	Num. of auctions	Num. of bids
10 BigDeal Bid Tokens	7.5	20,462	57,166
50 BigDeal Bid Tokens	37.5	7,839	176,824
\$25 Gift Card	25	2,130	12,793
\$100 Gift Card	100	509	13,355
iPod Nano 16GB	179	909	20,861
iPad 3G 64GB	900	410	82,897
Amazon Kindle	259	549	16,731
Kindle DX	379	620	46,350

Note that the penny auction is a second-price all-pay auction. In a standard second-price auction, the highest bid is unobservable. This is a potential threat to our analysis. However, a typical penny auction is not a standard auction, because the winner is not the highest bidder but the last; therefore, we can observe the highest bids in most cases, with the cost of the last bid biased downward.

3.5.4 Estimate Results

In this subsection, we report the estimate results based on the nonparametric structural model. Suppose that the density function of error term is known, and then our results depend on the specifications of $f_{\varepsilon(Z)}$. Here we assume the error term follows the Normal distribution, i.e., $\varepsilon \sim N(0, \sigma_Z^2)$. In particular, we set $\sigma_Z = \log(1 + Z)$, and the variable Z denotes the duration of bidders in each auction. Figures 3a-3d display the estimated distributions of private values for these four pairs of products. We summarize two important findings. First, we find the estimated private values are distributed around the retail prices. For example, in bid-token auctions, about 70% of bidders hold values below the retail prices. The result is reasonable since the retail prices can be seen as benchmarks of market values. Bidders are likely to refer

the retail price when they generate their own values. Second, bidders' values are highly related to the popularity of the products. For example, in Apple product auctions, more than 50% of bidders are willing to pay the retail price for iPad while few of bidders treat iPod with the same attitude. These findings are consistent with common intuition, and they provide some justifications for our estimation approach.

Figure 3(a): 10 Bid Tokens (\$7.5) vs. 50 Bid Tokens (\$37.5)




Figure 3(b): \$25 Gift Card vs. \$100 Gift Card




Figure 3(c): iPod Nano 16GB (\$179) vs. iPad 3G 64GB (\$899.99)





Figure 3(d): Amazon Kindle (\$259) vs. Kindle DX (\$379)



3.6 Conclusion

As we apply the nonparametric estimation approach to auctions, our observables may be different from the expected values which are predicted by the structural model. This paper provides a two-step nonparametric procedure to estimate the war of attrition with a stochastic number of bidders. In particular, we allow the observables are associated with heteroscedastic error. We apply this approach to penny auctions and estimate the distribution of bidders' private values.

In this paper, we assume that the distribution of error terms is known, otherwise, the unknown distribution of bidder's private values is not identifiable. However, the distribution of error terms is not specified in many practical situations. This classical condition is relaxed in some recent papers. As a trade-off, those models require either the availability of additional direct data from the error distribution or replicated measurements or more restrictive conditions

on the unknown density (see Li and Vuong 1998; Schennach 2004; Delaigle and Meister 2008; Krasnokutskaya 2011).

This paper is the first to estimate all-pay auctions using the nonparametric approach. Penny auctions provide the ideal data to conduct empirical analysis on all-pay auctions. The deconvolution method can also be applied in the nonparametric estimation approach for other standard auctions.

Bibliography

- [1] An, Yonghong, and Yingyao Hu. 2011. “Well-posedness of Measurement Error Models for Self-Reported Data.” Working paper.
- [2] Anderson, Chris K., and Fredrik Ødegaard. 2011. “Retail Selling with All-Pay Auctions.” Working paper.
- [3] Anderson, Steven, Daniel Friedman, Garrett Milam, and Nirvikar Singh. 2008. “Buy It Now: A Hybrid Internet Market Institution.” *Journal of Electronic Commerce Research*, 9(2): 137-153.
- [4] Athey, Susan, and Philip A. Haile. 2002. “Identification of Standard Auction Models.” *Econometrica*, 70(6): 2107-2140.
- [5] Athey, Susan, and Philip A. Haile. 2007. “Nonparametric Approaches to Auctions.” In J. Heckman & E. Leamer (Eds.). *Handbook of Econometrics*, Vol. 6A, Chapter 60.
- [6] Augenblick, Ned. 2011. “Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis.” Working paper.
- [7] Bajari, Patrick, and Ali Hortaçsu. 2004. “Economic Insights from Internet Auctions.” *Journal of Economic Literature*, 42(2): 457–486.

- [8] Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 1996. "The All-Pay Auction with Complete Information." *Economic Theory*, 8: 291-305.
- [9] Bonhomme, Stéphane, and Jean-Marc Robin. 2010. "Generalized Non-Parametric Deconvolution with an Application to Earnings Dynamics." *Review of Economic Studies*, 77: 491-533.
- [10] Bos, Olivier. 2012. "Wars of Attrition and All-pay Auctions with Stochastic Competition." *Journal of Mathematical Economics*, 48(2): 83-91.
- [11] Bose, Subir, and Arup Daripa. 2009. "Optimal Sale Across Venues and Auctions with a Buy-now Option." *Economic Theory*, 38: 137-168.
- [12] 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.
- [13] Budish, Eric B., and Lisa N. Takeyama. 2001. "Buy prices in Online Auctions: irrationality on the internet?" *Economics Letters*, 72: 325-333.
- [14] Bulow, Jeremy, and Paul Klemperer. 1999. "The Generalized War of Attrition." *American Economic Review*, 89(1): 175-189.
- [15] Bulow, Jeremy, and Paul Klemperer. 2009. "Why Do Sellers (Usually) Prefer Auctions?" *American Economic Review*, 99(4): 1544-1575.
- [16] 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.

- [17] Caldara, Michael. 2012. "Bidding Behavior in Pay-to-Bid Auctions: An Experimental Study." Working paper.
- [18] Camerer, Colin F. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- [19] 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.
- [20] Canals-Cerdá, José J., and Jason Percy. 2010. "Arriving in Time: Estimation of English Auctions with a Stochastic Number of Bidders." Working paper.
- [21] Carroll, Raymond J., and Peter Hall. 1988. "Optimal Rates of Convergence for Deconvolving a Density." *Journal of the American Statistical Association*, 83, 1184-1186.
- [22] Carroll, Raymond J., David Ruppert, Leonard A. Stefanski, and Ciprian M. Crainiceanu. "Measurement Error in Nonlinear Models: A Modern Perspective." Second Edition. Chapman & Hall/CRC. 2006.
- [23] Che, Yeon-Koo, and Ian Gale. 1996. "Financial Constraints in Auctions: Effects and Antidotes." *Advances in Applied Microeconomics*, Vol 6, 97-120.
- [24] Che, Yeon-Koo, and Ian Gale. 1998. "Caps on Political Lobbying." *American Economic Review*, Vol 88, No.3, 643-651.
- [25] 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.

- [26] 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.
- [27] Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. 2010. "Strategic Thinking." Working paper.
- [28] Delaigle, Aurore, and I. Gijbels. 2004. "Practical Bandwidth Selection in Deconvolution Kernel Density Estimation." *Computational Statistics & Data Analysis*, 45, 249-267.
- [29] Delaigle, Aurore, and I. Gijbels. 2006a. "Estimation of Boundary and Discontinuity Points in Deconvolution Problems." *Statistica Sinica*, 16, 773-788.
- [30] Delaigle, Aurore, and I. Gijbels. 2006b. "Data-driven Boundary Estimation in Deconvolution Problems." *Computational Statistics & Data Analysis*, 50, 1965-1994.
- [31] Delaigle, Aurore, and Peter Hall. 2006. "On Optimal Kernel Choice for Deconvolution." *Statistics & Probability Letters*, 76 (15), 1594-1602.
- [32] Delaigle, Aurore, and I. Gijbels. 2007. "Frequent Problems in Calculating Integrals and Optimizing Objective Functions: A Case Study in Density Deconvolution." *Statistics and Computing*, 17 (4), 349-355.
- [33] Delaigle, Aurore, and Alexander Meister. 2008. "Density Estimation with Heteroscedastic Error." *Bernoulli*, 14 (2), 562-579.
- [34] DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature*, 47(2): 315-372.

- [35] DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review*, 96(3): 694-719.
- [36] Durham, Yvonne, Matthew R. Roelofs, and Stephen S. Standifird. 2004. "eBay's Buy-It-Now Function: Who, When, and How." *Topics in Economic Analysis and Policy*, 4(1): Article 28.
- [37] Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan. 2011. "Learning from Seller Experiments in Online Markets." Working paper.
- [38] 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.
- [39] Friedman, Milton, and L. J. Savage. 1948. "The Utility Analysis of Choices Involving Risk." *Journal of Political Economy*, 56(4): 279-304.
- [40] Gneezy, Uri., and Ran Smorodinsky. 2006. "All-Pay Auctions—An Experimental Study." *Journal of Economic Behavior and Organization*, 61: 255–275.
- [41] Golec, Joseph, and Maury Tamarkin. 1998. "Bettors Love Skewness, Not Risk, at the Horse Track." *Journal of Political Economy*, 106: 205-225.
- [42] Goldfarb, Avi, and Botao Yang. 2009. "Are All Managers Created Equal?" *Journal of Marketing Research*, 46(5): 612-622.

- [43] 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.
- [44] Goodman, Joseph. 2012. "Reputations in Bidding Fee Auctions." SSRN working paper.
- [45] Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-price Auctions." *Econometrica*, 68(3): 525-574.
- [46] Haile, Philip A., and Elie Tamer. 2003. "Inference with an Incomplete Model of English Auctions." *Journal of Political Economy*, 111(1): 1-51.
- [47] Hammond, Robert G. 2010. "Comparing Revenue from Auctions and Posted Prices." *International Journal of Industrial Organization*, 28: 1-9.
- [48] Hidvégi, Zoltán, Wenli Wang, and Andrew B. Whinston. 2006. "Buy-price English Auction." *Journal of Economic Theory*, 129: 31-56.
- [49] Hinnoosaar, Toomas. 2010. "Penny Auctions". Working paper.
- [50] Houba, Harold, Dinard van der Laan, and Dirk Veldhuizen. 2011. "Endogenous Entry in Lowest-Unique Sealed-Bid Auctions." *Theory and Decision*, 71: 269-295.
- [51] 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.
- [52] 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.

- [53] Kirkegaard, René, and Per Baltzer Overgaard. 2008. "Buy-out Prices in Auctions: Seller Competition and Multi-unit Demands." *RAND Journal of Economics*, 39(3): 770-789.
- [54] Krasnokutskaya, Elena. 2011. "Identification and Estimation of Auction Models with Unobserved Heterogeneity." *Review of Economic Studies*, 78: 293-327.
- [55] Krishna, Vijay. 2002. *Auction Theory*. San Diego, CA: Academic Press.
- [56] Krishna, Vijay, and John Morgan. 1997. "An Analysis of the War of Attrition and the All-pay Auction." *Journal of Economic Theory*, 72(2): 343-362.
- [57] Kultti, Klaus. 1999. "Equivalence of Auctions and Posted Prices." *Games and Economic Behavior*, 27: 106-113.
- [58] Laffont, Jean-Jacques, and Quang Vuong. 1996. "Structural Analysis of Auction Data." *American Economic Review, Papers and Proceedings*, 86(2): 414-420.
- [59] Levin, Jonathan D. 2011. "The Economics of Internet Markets." NBER Working Paper No. 16852.
- [60] Li, Tong, and Quang Vuong. 1998. "Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators." *Journal of Multivariate Analysis*, 65: 139-165.
- [61] List, John A. 2003. "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics*, 118(1): 41-71.
- [62] List, John A., and David Lucking-Reiley. 2000. "Demand Reduction in Multiunit Auctions: Evidence from a Sportscard Field Experiment." *American Economic Review*, 90(4): 961-972.

- [63] Liu, MingChung, and Robert L. Taylor. 1989. "A Consistent Nonparametric Density Estimator for the Deconvolution Problem." *The Canadian Journal of Statistics*, 17(4): 427-438.
- [64] Malmendier, Ulrike, and Young Han Lee. 2011. "The Bidder's Curse." *American Economic Review*, 101(2): 749-87.
- [65] Matthews, Timothy. 2004. "The Impact of Discounting on an Auction with a Buyout Option: a Theoretical Analysis Motivated by eBay's Buy-It-Now Feature." *Journal of Economics*, 81(1): 25-52.
- [66] McAfee, R. Preston, and John McMillan. 1987. "Auctions with a Stochastic Number of Bidders." *Journal of Economic Theory*, 43(1): 1-19.
- [67] Meister, Alexander. 2009. *Deconvolution Problems in Nonparametric Statistics*. Springer Series in Statistics.
- [68] Myerson, Roger B. 1981. "Optimal Auction Design." *Mathematics of Operations Research*, 6(1): 58-73.
- [69] Nakajima, Daisuke. 2008. "First-Price Auctions, Dutch Auctions, and Buy-It-Now-Prices with Allais Paradox Bidders." Working paper.
- [70] Newey, Whitney, and Danial McFadden. 1994. "Large Sample Estimation and Hypothesis Testing." *Handbook of Econometrics*, 4, 2111-2245.
- [71] Ostling, Robert, Joseph Tao-yi Wang, Eileen Y. Chou, and Colin F. Camerer. 2011. "Testing Game Theory in the Field: Swedish LUPU Lottery Games." *American Economic Journal: Microeconomics*, 3(3): 1-33.

- [72] Perrigne, Isabelle, and Quang Vuong. 1999. "Structural Econometrics of First-price Auctions: A Survey of Methods." *Canadian Journal of Agricultural Economics*, 47: 203-223.
- [73] Platt, Brennan C., Joseph Price, and Henry Tappen. 2010. "Pay-to-Bid Auctions." Working paper.
- [74] Raviv, Yaron, and Gabor Virag. 2009. "Auctions by Gambling." *International Journal of Industrial Organization*, 27: 369-378.
- [75] Reynolds, Stanley S., and John Wooders. 2009. "Auctions with a Buy Price." *Economic Theory*, 38: 9-39.
- [76] Riley, John G., and William F. Samuelson. 1981. "Optimal Auctions." *American Economic Review*, 71(3): 381-392.
- [77] Schennach, Susanne. 2004. "Nonparametric Regression in the Presence of Measurement Error." *Econometric Theory*, 20: 1046-1093.
- [78] Schwarz, Maik, and Sébastien Van Bellegem. 2010. "Consistent Density Deconvolution Under Partially Known Error Distribution." *Statistics and Probability Letters*, 80: 236-241.
- [79] Shubik, Martin. 1971. "The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation." *Journal of Conflict Resolution*, 15(1): 109-111.
- [80] Shunda, Nicholas. 2009. "Auctions with a Buy Price: The Case of Reference-Dependent Preferences." *Games and Economic Behavior*, 67: 645-664.

- [81] Stefanski, Leonard A., and Raymond J. Carroll. 1990. "Deconvoluting Kernel Density Estimators." *Statistics*, 21: 169-184.
- [82] Tsybakov, Alexandre B. 2009. *Introduction to Nonparametric Estimation*. Springer Series in Statistics.
- [83] Tversky, Amos, and Daniel Kahneman. 1986. "Rational Choice and the Framing of Decisions." *Journal of Business*, 59(4), part 2: S251-S275.
- [84] Wang, Ruqu. 1993. "Auctions versus Posted-Price Selling." *American Economic Review*, 83(4): 838-851.
- [85] Wang, Xiaofeng, Zhaozhi Fan, and Bin Wang. 2010. "Estimating Smooth Distribution Function in the Presence of Heteroscedastic Measurement Errors." *Computational Statistics and Data Analysis*, 54: 25-36.
- [86] Wang, Xiaofeng, and Bin Wang. 2011. "Deconvolution Estimation in Measurement Error Models: The R Package decon." *Journal of Statistical Software*, 39(10): 1-24.
- [87] Wang, Zhongmin, and Minbo Xu. 2011. "Learning and Strategic Sophistication in Games: The Case of Penny Auctions on the Internet." Working paper.