PENNY AUCTIONS

By

Antonia Hao Chen

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy in Economics

The University of Auckland
2018

Introduction

The low setup cost of online auction platforms has facilitated the creation of new varieties of auction formats. A popular new auction variety invented by a German company 'Swoopo' in 2005 requires participants pay a small fee at each time of bidding (Platt et al., 2012). It was named Penny Auction, as each bid increased the current price by a fixed small increment, which was usually one cent. Swoopo's success has brought hundreds of competitors running similar format pay-to-bid auctions around the world (Zimmerman, 2011). Swoopo remained one of the market leaders from its founding until it filed for bankruptcy in early 2011. At its peak, Swoopo operated internationally in 22 regions (including the United States and most EU countries), earned a profit of $28.3 million in 2008 and had 2.5 million users in 2009 (Oswald 2008; Stone 2009). Since then, more than 150 market entrants including BidCactus, BigDeal, and Quibids have tried to capture a piece of the market (Stone 2010), which swelled to 300,000 U.S. visitors daily in 2011.[1] Table 1 shows the numbers of monthly unique visitors of a few of the largest penny auction websites, monitored by Compete.com, a web traffic monitoring company, showing penny auctions website even reached around 25% of eBay's data at the end of 2010 (Wang Xu, 2016).

Website	**Unique visitors**		
	Feb 2010	**Nov 2010**	**Apr 2011**
Bidcactus.com	1,428,316	3,411,705	1,979,846
BigDeal.com	480,230	1,324,947	943,327
Quibids.com	173,142	4,541,783	4,586,523
Swoopo.com	286,142	171,141	Closed
Total number of sites	47	125	158
All sites	4,701,541	16,866,475	12,524,625
eBay.com	64,766,668	67,197,011	69,929,590
% of eBay traffic	7.3%	25.1%	17.9%

Table 1. Monthly traffic of the largest penny auction websites

[1] As recorded by http://www.pennystats.com/ in March 2011.

Different penny auction websites may have slight differences in their auction settings, such as charging different amounts of bidding fees, with some general rules followed by most of the auctioneers. In general, each penny auction starts at a price of zero with a specified closing time displayed in a countdown clock. To place a bid, a bidder is required to pay a small bidding fee (usually between $0.60 and $1), which can either be charged immediately at the time of bidding, or deducted from his/her pre-purchased bidding credits. After a bid is placed, the current price increases by a fixed amount (usually between $0.01 and $0.15), substantially smaller than the bidding fee. After the original closing time is reached, each new bid can extend the auction by a set short amount of time (usually between 10~30 seconds). If no other bid is placed before the time expires, the last bidder wins and pays the current price. It differs from English auctions, where the winner of a penny auction pays not only the winning bid, but also the bidding fees incurred throughout the entire auction, while the auctioneer gains his revenue not only from the winning bid, but also from the aggregate bidding fees paid by all participating bidders. It also differs from ordinary online auction platforms such as eBay and Amazon, where a penny auction website hosts all its auctions and lists selected brand new items of limited variety, such as popular consumer electronics, video game consoles, gift cards and packs of bidding credits, which have a relatively well-defined market value.

Penny auctions have drawn academic attention because of their similarity and differences to some well-known auction mechanisms such as Martin Shubik's Dollar Auction, and War of Attrition and all-pay auctions (Hinnosaar, 2016). A variety of theoretical models based on complete information and risk neutral sellers and bidders predict zero expected profit, but data suggest otherwise. Empirical evidence has shown that bidding fees are primary revenue sources of penny auction operators who earn substantial and consistent profits, robust over time. Although some penny auctions ended with high profit margins, not all auctions were profitable, e.g. Swoopo only made positive profit in half of its auctions (Platt, Price et al. 2010; Byers, Mitzenmacher et al. 2010).

A simple example would help us understand the profitability of the mechanism more intuitively. Suppose in some penny auction, bidders pay $1 to place a bid, which raises the current price by $0.02, then for every $1 increase in the winning bid, the seller collects $50 additional revenue from bidding fees. Thus, as long as the winning bid reaches 1.96% of the seller's valuation, he would break even, and if the winning bid reaches 5% of his valuation, his profit margin would be 200%. Despite its profitability to the auctioneer, a penny auction may still look attractive to many bidders, especially the newcomers, who observe that the current price of an auction item is usually remarkably low and the winning bid of a closed auction is also usually low. Swoopo managed to deliver expensive consumer products at significantly low prices that beat all traditional retailers in most of its auctions, while retaining a high profit margin. Note that although Swoopo gained negative profits in around half of its auctions, the other half were successful enough to result in an overall profit margin of 50% from 166,000 auctions spanning from September 2005 to June 2009. In fact, the median winning bid of Swoopo auctions is only 10% of the retail price of the auction items (Augenblick, 2011). Despite these obvious pros, the tricky con is that costs spent on past bids are technically sunk and have no bearing on one's likelihood of winning once being outbid, which is one of the reasons that newcomers often spend much more than they plan to. Note that placing a bid in a penny auction has a distinct interpretation from bidding in a traditional auction, in which a bid represents a bidder's willingness to pay for the auctioned item and is not payable unless it turns out to be the winning bid, while money spent on bids in a penny auction is unrecoverable.

Structures of most penny auction websites, including relatively small bidding fees and discounted packages of bidding credits, promote irrational bidding behaviours and exploit the sunk cost fallacy[2]; as bidders continue to participate in an auction, and they spend more money on bids leading them to experience a higher psychological cost from leaving the auction (Eyster, 2002; Augenblick, 2011). On the other hand, for a newcomer of some penny auction websites with a bidding fee of $1, it is likely that he may not realise he has already spent $100 after placing 100 bids. Most penny auction websites do not provide full bidding histories to their visitors. For instance, when viewing a live, or completed auction on Swoopo, only the ten most recent bids are visible, and most other auctioneers provide between five to ten recent bids in live

[2] One potential explanation for high auctioneer profits comes from the dollar auction (Shubik, 1971), which shares many characteristics with the penny auction. In the dollar auction, two players sequentially bid slowly escalating amounts to win a dollar bill, but both are required to pay their last bid. The dollar auction is known as a "prototypical example" of the irrational escalation of commitment (also known as the sunk cost fallacy), in which players become less willing to exit a situation as their financial and mental commitments increase, even if these commitments do not increase the probability of success (Camerer and Weber, 1999). This suggests that the sunk-cost effect also could be playing a role in penny auctions, as players make similarly escalate financial commitments (in the form of bid costs) as the auction continues (Augenblick, 2011).

auctions; although BidCactus displayed a full list of all bidders and their total number of bids in completed auctions, full bidding histories were not provided. Furthermore, the bidding history of a bidder is usually not available for him/her to track either, so it is not obvious to a bidder how much he/she has spent so far in a live auction.

From sellers' points of view, unlike traditional auctioneers who earn from the winning bids, bidding fees are the primary source of revenue for penny auction websites. Their mechanisms and auction settings (such as tiny bidding increments) promote the total number of bids placed in each auction. In a traditional online auction, most bids are placed at the beginning and within the last minutes before it ends, and because the bidding increment is customisable, simply placing a bid may not make a bidder the current leader, which may discourage some bidders from participating. However, in a penny auction, an operator would prefer to keep participants bidding steadily throughout the auction, so settings such as zero starting price, fixed tiny price increment per bid, and high valuation of auction items, which maintain the attractiveness of placing a bid throughout the auction. Similarly, a penny auction operator would also want to maximise the auction's length.

In traditional online auctions, bidders tend to wait until the last moment to place their bids (Roth and Ockenfels, 2002). The dynamic countdown clock in a large bold font is a key feature adopted by all penny auction websites. Together with super short extension periods (usually 10~30 seconds) being counted down in seconds, it builds a great atmosphere to attract new bidders to join in as well as encouraging current participants to place bids continuously, resulting in auto-extension periods that can last for hours, and even days. For instance, 120 bids that only increase the current price by $1.20 could take up to one hour in an auction with a 30-second countdown timer.

However, it is not reasonable to expect ordinary bidders to remain active by placing bids manually in an online auction for hours or days, thus automatic bidding tools such as BidButler of Swoopo, are available to all bidders. For example, a bidder can set his BidButler to place a bid on a certain auction every time when the countdown clock reaches the last second, such that he can stay in the auction as long as he has bidding credits remaining (MacDonald, 2011). Besides this, large operators such as Swoopo and BidRivals run their auctions internationally, so that bidders living in different time zones across the globe can compete in every single auction, which contribute to extending the auction's length.

In my study, I collect my original dataset of over half a million auctions over a timeframe of four years from BidCactus.com, one of the market leaders after the shutdown of Swoopo. The data show wide variability in profitability across categories of items and overall high profitability. I first replicate the Maximum Likelihood methodology of Platt et al. (2013) to estimate risk parameters implied by bidding behaviour for my original dataset, and found evidence of bidders' experiences affecting the profitability, which was consistent with evidence from earlier studies using other datasets that reveal differential bidding behaviour based on experience. I therefore extend the Platt et al. (2013) model to allow for multiple types of bidders with different prior experiences and estimate type-specific risk parameters via Maximum Likelihood when there are up to three types of bidders in an auction. Bidders in different experience groups in our dataset are shown to have significantly different risk attitudes. When bidders are separated into two bidding groups of inexperienced bidders with participation experience of less than 20 prior auctions with the rest being experienced bidders, I observe that the more experienced bidders were more risk-seeking and bidding more aggressively, which contributes more to the seller's revenue, regardless of whether they competed with rivals of the same type, or otherwise.

Another interesting fact that was observed, was that 0.2% of bidders participated in 1,000 or more auctions, and placed over 20% of all bids. If we separate this small proportion of bidders into a new group called the super-experienced bidders, and called the rest of the experienced bidders as ordinary-experienced bidders, we observed that the super-experienced bidders whose behaviours are usually affected by type composition of rivals. They play more aggressively than the other two types when playing only against rivals of their own type; while they play more conservatively compared to the ordinary-experienced bidders (those with 20~1000 prior auction experiences) when there exist rivals of other types in an auction. Implication of bidders' experience types on seller's profitability are discussed in more details in Chapter 6 and 7.

Chapter 1. Literature Survey

Theoretical Literature

Past theoretical literature has established a baseline model which features risk neutral participants and complete information, and which implies zero profitability. Data, as we explained in the Introduction, suggests otherwise. Various authors attempt to explain the observed excess revenue by extending the basic model to capture risk-loving preferences (Platt et al., 2013), sunk cost fallacies (Augenblick, 2015; Hinnosaar, 2016), information asymmetry, collusion and shill bidding (Byers et al., 2010), signalling strategies (Augenblick, 2012; Byers et al., 2010), and additional utility from auction participation (Hinnosaar, 2016).

Platt et al. (2013) proposed and tested a model of penny auctions to predict the distribution of ending prices, which suggests that bidders are risk-seeking to some extent. Their symmetric complete information model is set up as follows:[3]

The auctioneer sells an object with a known, objective value of v dollars to N potential bidders who are risk-neutral, and the state of the auction is described by the number of elapsed periods t and the identity of the current winning bidder, $i \in \{1, \ldots, N\}$. The auction starts at price $p_0 = 0$. For every bid that is placed successfully, the bidder must pay a bidding fee of c dollars to the auctioneer, and the current price is raised by s dollars. Thus, the current price at period t is $p_t = st$. During each period $t > 0$, the $N-1$ non-leaders (or N non-leaders at initial period $t = 0$) simultaneously choose whether to place a bid. If no one places a bid, the auction closes and the current leader wins the object and pays the current price p_t. It is assumed that if $K > 1$ bidders place a bid, one of them is randomly selected with probability $1/K$ to become the new leader, and is the only one to pay the bidding fee. In addition to common initial wealth of W, each bidder has the same von Neumann-Morgenstern (vNM) utility function of $u(\bullet)$ and $u(W) = W$.

[3] Slightly different notation to that in Platt et al. (2013) is used in this literature review for consistency throughout the thesis and for easier reading.

In the symmetric Markovian Subgame Perfect Equilibrium, each player's next move is predicted by the last move of the other player, not by earlier history of moves. All non-leaders at period t employ the identical Markov strategy with the same probability $\beta_{t+1} \in [0,1]$ of attempting to place the $(t+1)^{th}$ bid. Thus the aggregate probability of the t^{th} bid occurring $1-\mu_t \equiv (1-\beta_t)^{N-1}$, where μ_t is the probability of the t^{th} bid being placed by any non-leaders in the $(t-1)^{th}$ period, solved by backward induction.

There exists the maximum number of periods possible, as no bidders would be willing to place a bid once the current price plus the bidding fee exceeds the item's value. Let $T \equiv \dfrac{v-c}{s}$ for $s>0$, and assume T is an integer for easier analysis. Even if this assumption does not hold, players would strictly prefer to bid in the $(T-1)^{th}$ period, since no one would bids in the T^{th} period. No bidder in the T^{th} period or beyond is willing to place a bid, i.e. no one is willing to place the $(T+1)^{th}$ bid. Because if a bidder places a bid and becomes the new leader in the $(T+1)^{th}$ period, the new current price will be $p_{T+1} = s(T+1) = s\left(\dfrac{v-c}{s}+1\right) = v-c+s$, and her expected payoff is negative as $\left[v-s(T+1)\right](1-\mu_{T+1})-c = \left[c-s\right](1-\mu_{T+1})-c<0$. Therefore, we must have $\mu_t = 0$ for $t>T$.

For any non-leader in the periods $1<t<T$, placing a bid brings her a positive payoff only if her bid is accepted and no one else bids in the next period. She is indifferent between placing the t^{th} bid and not placing the t^{th} bid, so that $(v-st)(1-\mu_{t+1})-c=0$. Thus the probability of the t^{th} bid occurring in the game is $\mu_t = 1-\dfrac{c}{v-s(t-1)}$, and

$\beta_t = 1-(1-\mu_t)^{\frac{1}{N-1}} = 1-\left(\dfrac{c}{v-s(t-1)}\right)^{\frac{1}{N-1}}$ when $1<t\le T$.

Let $\mu_1 \in [0,1]$, $\mu_t = 1 - \dfrac{c}{v - s(t-1)}$ *for* $1 < t \le T$, *and* $\mu_t = 0$ *for* $t > T$.

Then the strategy profile $\beta_1 = 1 - (1-\mu_0)^{\frac{1}{N}}$ *and* $\beta_t = 1 - (1-\mu_t)^{\frac{1}{N-1}}$ *for* $t > 1$ *constitutes a unique symmetric Markovian Subgame Perfect Equilibrium.* There also exists other symmetric equilibria, in each of which the auction either ends in period 0 with no bidder, or ends in period 1 with one winning bidder.

Using conditional probability μ_{t+1} (the aggregate probability of the $(t+1)^{th}$ bid occurring, given that the t^{th} bid already has occurred), the probability density that the auction ends at exactly t bids is constructed as follows

$$f(t) \equiv (1-\mu_{t+1})\prod_{j=1}^{t}\mu_j = \begin{cases} 1-\mu_1 & \textit{if } t = 0 \\ \dfrac{c}{v-st}\mu_1\prod_{j=2}^{t}\left(1-\dfrac{c}{v-s(j-1)}\right) & \textit{if } 0 < t \le T \end{cases}$$

The expected revenue of the auctioneer $E[Rev] = \sum_{t=1}^{T}(c+s)t\cdot f(t)$, is shown to be equal to v, which is independent of parameters such as bidding fees and bid increments, while variance of the expected revenue increases in c and v, and decreases in s. As it is assumed that the auctioneer has the same valuation v as the bidders, he would get zero expected profit from the auction.

Platt et al. (2013) attempts to explain the significant profitability of penny auctions by relaxing the assumption of risk-neutrality and incorporating preferences towards risk, which seem natural, since placing a bid in a penny auction looks like paying a small fee to gamble that other bidders who will not place the next bid in order to win a big prize. Assume bidders have a Constant Absolute Risk Aversion utility function of $u(W) = \dfrac{1-e^{-\alpha W}}{\alpha}$, where Absolute Risk Aversion equals α. Similarly, the indifference condition between placing a risky bid and not placing a bid is calculated as:

$$(1-\mu_{t+1})\frac{1-e^{-\alpha(W+v-st-c)}}{\alpha}+\mu_{t+1}\frac{1-e^{-\alpha(W-c)}}{\alpha}=\frac{1-e^{-\alpha W}}{\alpha}$$

where $\frac{1-e^{-\alpha(W+v-st-c)}}{\alpha}$ is utility of winning the auction after placing the bid, and $\frac{1-e^{-\alpha(W-c)}}{\alpha}$ is utility of being outbid after placing a bid. The solution of the indifference condition is $\mu_t=\frac{1-e^{\alpha(c+s(t-1)-v)}}{e^{\alpha c}-e^{\alpha(c+s(t-1)-v)}}$ and the corresponding symmetric equilibrium bidding function is $\beta_t=1-(1-\mu_t)^{\frac{1}{N-1}}$. Note that the previous risk-neutrality case corresponds to $\alpha=0$; risk aversion corresponds to $\alpha>0$ and risk seeking corresponds to $\alpha<0$. CARA utility enables us to eliminate the wealth effects in the indifference condition, which ensures wealth differences that arise during bidding do not create heterogeneous bidding incentives or alter equilibrium behaviour.

The probability that the auction ends in period t is

$$f(t)\equiv(1-\mu_{t+1})\prod_{j=1}^{t}\mu_j=\frac{1-e^{\alpha c}}{e^{\alpha c}-e^{\alpha(c+st-v)}}\prod_{j=1}^{t}\left(\frac{1-e^{\alpha(c+s(j-1)-v)}}{e^{\alpha c}-e^{\alpha(c+s(j-1)-v)}}\right)$$

It can be shown that the expected number of bids is decreasing in α,s,c. This implies that expected revenue is decreasing in α, i.e. the auctioneer is able to extract more revenue and profit from more risk-loving bidders whose have risk parameters $\alpha<0$.

The basic model of Augenblick (2012) is similar to Platt et al. (2013), and also restricts attention to symmetric Markovian strategies, but assumes bidders are risk-neutral. The hazard function at time t, conditional on survival until time t or later, is defined as

$$\tilde{h}(t,l_t)\equiv \mathbf{Pr}\left[\textit{every bidder } i \textit{ chooses Not Bid for all } i\neq l_t \middle| \textit{reaching } t \textit{ with leader } l_t\right]$$

i.e. $\tilde{h}(t, l_t) = \frac{\prod_{i=1}^{N}\left(1-\beta_t^i\right)}{1-\beta_t^{l_t}}$ and $\tilde{h}(0)=0$ is arbitrarily chosen. To make smoother empirical predications of hazard rates, on top of the model in discrete time setting, Augenblick (2012) builds a model with a continuous time setting by shrinking the size of the time periods to approach zero. Δt is set to be a small length of time to remodel time points as $t \in \left\{0, \Delta t, 2\Delta t, 3\Delta t ...\right\}$, and the bidding cost is changed to $c\Delta t$. To compare survival and hazard rates across auctions of items with different valuations, time is normalised by item valuation. T is defined as the time that an auction ends, and the normalised time period $\hat{t} = \frac{t}{v}$ and $\hat{T}$ is the normalised time that an auction ends. The survival function is defined as $S(t) = \lim_{\Delta t \to 0} \Pr(T > t)$, and the hazard rate is defined as $h(t) = \lim_{\Delta t \to 0} \frac{S(t) - S(t+\Delta t)}{\Delta t \cdot S(t)}$. Thus, the discrete equilibrium equivalent in the continuous time setting is:

$$h(t) = \frac{c}{v - st} \text{ and } h(\hat{t}) = \frac{c}{1 - s\hat{t}} \text{ for } t < \frac{v}{s},$$

$$\text{when } s > 0,\ S(t) = \left(1 - \frac{st}{v}\right)^{\frac{c}{s}} \text{ and } S(\hat{t}) = \left(1 - s\hat{t}\right)^{\frac{c}{s}} \text{ for } t < \frac{v}{s}.$$

Augenblick (2015) proposes that the naive sunk cost fallacy is the most intuitive explanation of the differences between empirical findings and the theoretical model. He assumes that bidders become less and less willing to leave an auction as they place more bids, even though those costs are "sunk". To capture sunk costs, Augenblick (2015) assumes that each player's perception of the value of the good rises as she spends more money on biding costs. A player i who has placed s_i bids has sunk costs $s_i c$ and perceives the value of the good as $v + \theta s_i c$ with $\theta \geq 0$ defined as the sunk cost parameter. As this parameter rises, the player's sunk costs cause her to bid with a higher likelihood in the auction. If this parameter is zero, the model reverts to the standard risk neutral model above.

Assume that the player is naive about this sunk-cost effect, in the sense that she is unaware that her perception of value might change in the future and also unaware that other players do not necessarily share her perception of value. Without the first type of naivety, players would be aware that they will bid too much in the future and consequently require a compensating premium to play the game initially, thus leading to zero profits for the auctioneer (and violating the empirical observations). Without the second type of naivety, each player would have very complicated higher-order beliefs, being personally unaware of her own future changes in value perceptions, but being aware of other player's changing perceptions and of other players' (correct) beliefs about her own changing perceptions. Furthermore, due to the mechanics of mixed strategy equilibria, each player's bidding probability would largely be determined by the sunk costs of other players rather than her own sunk costs. With this dual naivety assumption, a player simply plays the game as if the value of the good matches her perceived value, which includes a portion of her own sunk costs.

The sunk costs faced by a player at a specific time t depend on the realizations of the player's own mixed decisions, the mixed decisions of the other players, and the realization of the leader's selection process. Define ξ_i^t as the total sunk bids player i has placed up to time t in an auction, so that that bid probabilities β_i^t in the Markovian equilibrium strategies are restricted to only depend on t and sunk costs ξ_i^t.

With sunk costs, for any game with $s > 0$:

When $t > 1$, $\beta_0^i > 0$ and $\beta_1^i > 0$ for all i, in any Markov Perfect Equilibrium, the symmetric Markov Perfect Equilibrium strategies are:

$$\beta_t^i = \begin{cases} 1 & t = 0 \\ 1 - \sqrt[n-1]{\dfrac{c}{v - st + \theta \xi_i^t c}} & for\ 0 < t \leq T \\ 0 & for\ t > T \end{cases} \Bigg\} for\ all\ i$$

It is shown that bidders overbid significantly due to the fallacy of sunk costs, and the auctioneer earns a considerable profit from this overbidding.

Hinnosaar (2016) argues from a different sunk cost perspective from Augenblick (2015) that individuals might not consider c to be at the same monetary scale as v and p, since it is partly sunk. As in practice bidders are able to buy "bid packs" with 50 or 100 bids at a time, therefore, with some probability an individual has marginal cost of next bid less than c. When bidders consider the cost of bid $c_b < c$, and expected revenue is greater or equal to the valuation, that it is possible to earn profit.

Hinnosaar (2016)'s analysis of Swoopo is among the first few working papers on penny auctions. A major difference of Hinnosaar (2016) compared to all other theoretical works is the treatment of simultaneous bidding attempts by two or more bidders if one or more of the non-leaders submit a bid, each of them will be the leader in the next period with equal probability, but each of these players pays the bidding fee to the seller, so the price increments by s times the number of simultaneous bids received. At each period t, the non-leaders simultaneously choose to either *submit a bid* or *pass*. If $K+1>0$ non-leaders submit a bid, each of them will be the leader in the next period with equal probability of $\frac{1}{K+1}$, the price increases to $p_{t+1} = p_t + (K+1)s$ and each of these $K+1$ players pays c dollars to the seller. Since each player makes her decision independently, the probability of player i becoming the new leader after she submits a bid, when each of $N-1$ other non-leaders are bidding with probability β is

$$\sum_{K=0}^{N-1}\binom{N-1}{K}\beta^{K}(1-\beta)^{N-(K+1)}\frac{1}{K+1} ,$$

which is shown to be strictly decreasing in β and N, i.e. one is less likely to become the leader if there are more opponents or if the opponents bid more actively.

The symmetric Stationary Subgame Perfect Nash Equilibria (SSSPNE) are considered, where SSSPNE are that satisfy both ***Symmetry*** (players' identities do not play any role) and ***Stationarity*** (players only condition their behaviour on the current price and number of active bidders, NOT on the whole history of bids or identities of leaders). It is shown that at any SSSPNE, the game is finite, and there exists a point in time where the game has ended with certainty; expected revenue is less than or equal to the valuation of the item; and when the object is sold, revenue is greater and less than the valuation with positive probabilities. However, the observed average profit margin of penny auctions is significantly higher than zero, that the model needs to be extended further to achieve an outcome where expected revenue may be significantly higher than the value of the object.

Hinnosaar (2016) argues that the suggested retail value is higher than the cost to the seller, which is around the value that the customers expect to pay, and that the bidders get some positive utility from participating, v_g, so that $v_b - v_g = v > v_s$, where v_b is the buyer's value, including the participation utility, v_g and v_s is the seller's value. The buyer's value could also be an increasing function of N (the utility value of beating ***N-1*** opponents is increasing in N). Other possible ways to model this entertainment value are: modelling it as a lump-sum value just from participating, as a positive income that is increasing in the number of bids, or assuming that "Saving" money gives some additional happiness that instead of $v-p$, a risk-neutral player would have some increasing and convex utility functions $u(v-p)$. This sort of utility gain is generated through analysis of bidders' risk attitude in Platt et al. (2013).

In all above literature, information asymmetry and asymmetric equilibria in which homogeneous bidders employ different mixed strategies at a particular period, are not considered. Byers et al. (2010) first analyse the impact of information asymmetry broadly, as well as Swoopo's features such as bid packs and the buy-now option specifically, to quantify the effects of imperfect information in these auctions. It was found that even small asymmetries across players (cheaper bids, better estimates of other players' intent, different valuations of items, committed players willing to play "chicken") can increase the auction's duration well beyond that predicted by previous work and thus skews the auctioneer's profit disproportionately, even with fully rational players. If players overlook, or are unaware of any of these factors, the result is outsized profits for pay-per-bid auctioneers. Behavioural factors are also examined through the dataset of live auctions, such as the power of aggressive bidding.

Byers et al. (2010) start with their basic symmetric model and analysis of Swoopo auctions based on Platt et al. (2013) with all players playing identical strategies and two original datasets. One outcomes dataset with 121,419 Swoopo auctions and limited information about an auction such as the product description, the retail price, the final auction price, the bid fee, and the price increment: and one traces dataset of 4,328 live auctions recorded using their own recording infrastructure, with detailed bidding information for each auction including the time and the player associated with each bid.

A rational player's strategy in the basic model depends on his/her assessment of the probability of winning the auction by bidding, based on the auction's parameters such as the current bid, the number of bidders, the bid fee, and the value of the item. Arguing that the information asymmetries arise naturally in players' perceptions of penny auction parameters, the impacts of asymmetry on profitability are analysed together with adoption of the Markov chain for modelling general asymmetries. For instance, although the basic model assumes that the number of players is known to all players in advance, there is no way of knowing exactly how many players are actively participating or monitoring the auction at any time in practice, especially with Swoopo only displaying the list of bidders active in the last 15 minutes. It shows that an underestimation in the number of bidders' increases profitability, as well as certain mixtures of underestimation/overestimation, and even small asymmetries in beliefs in the number of active players can lead to dramatic changes in overall auction revenue, which can grow sharply as the estimates vary from the true number of players.

Other assumptions of the basic model also challenged, are that all players share an identical valuation to an auction item and pay the same bidding fee to place a bid. Different valuations among bidders are obvious, considering Swoopo operates the same auction internationally with a suggested retail price in different currencies that are usually higher than Amazon's list price. The more players that overestimate the item, the larger its impact on profitability. Players who occupy bid credits at a discounted price by winning bidpack auctions, or purchasing sets of prepaid bids beforehand, have lower bidding fees compared to other participants, without their knowledge in general. Players having a lower bidding cost have both a decided information advantage and a tactical advantage. A shill bidder (an extreme case of zero bidding costs) bids on behalf of the auctioneer and does not claim any winning item if he wins, and all revenue from bids placed by other bidders are profit for the seller. While Byers et al. (2010) do not suggest the presence of shill bidding in penny auctions, they do show that it would have a striking impact on seller profitability.

It is claimed that even small asymmetries across players (cheaper bids, better estimates of other players' intent, different valuations of items, committed players willing to play "chicken") can increase the auction's duration significantly and thus skew the auctioneer's profit disproportionately. Intrinsic aspects of the penny auction's mechanism are able to derive profit from even rational, risk-neutral players who correctly model sunk costs, if the players overlook, or are unaware of any of these factors.

Other interesting aspects including behavioural factors that are difficult to models analytically, are examined empirically, e.g. whether aggressive bidding is an effective strategy. Empirical finding shows that the most aggressive bidders contribute the lion's share of profits to Swoopo although aggressive bidders win more often. It is also shown that the profitability of penny auctions is potentially fragile, especially in cases where signalling by committed players willing to play a game of chicken, or collusion between players, can end the auction earlier.

Empirical Literature

Past empirical research with original data collections have shown significant profitability of the penny auction mechanism (Augenblick, 2012; Byers et al., 2010; Goodwin, 2012; Platt et al., 2013), and attempt to explain the excess revenue of auctioneers and bidders' behaviour by different approaches (Augenblick, 2012; Goodman, 2012), discussing sustainability of the mechanism (Wang and Xu, 2012) and bidders' retention (Stix, 2012; Zheng et al., 2011).

Platt et al. (2013) tested their model against the observed bids on over 126,000 Swoopo auctions of 1,958 unique items between September 2008 and May 2010. Their key theoretical prediction is that the final number of bids in a given auction is a random variable with distribution $f(t)$. If a given item type is repeatedly auctioned, whether this sample distribution is consistent with its theoretical counterpart can be tested, by performing statistical tests (Pearsons χ^2 Test and K-S test to compare distributions, t-test to compare means) to quantify how closely the estimated theoretical distribution matches the observed sample distribution. In particular, the role played by the bidders' common valuation and risk preference in replicating the observed behaviour with the theoretical model is examined by estimating the unobserved parameters v and α. The item valuation is not observed, not only because bidders may have private valuations, but also because even retail price such as Amazon's price might not reflect the true valuation, e.g. an auctioned item frequently consists of a package of several goods; some items offered by Swoopo, such as newly-released video games, were at times difficult to find through traditional outlets.

To estimate these parameter values, Maximum likelihood is used; i.e. choosing the parameter(s) to maximise $\sum_m \ln f(t_m; v, \alpha)$, where m represents each observed auction of a given type of item, t_m is the ending number of bids in that auction. Four sets of empirical tests are performed in four settings of model parameters α and v, and summary of the results is shown in Table 2. Each cell under the three goodness-of-fit tests reports the percentage of items for which there is no significant difference between the theoretical prediction and observed data (with p-value greater than .05 or .10), repeated for each of the four specifications. It indicates how often the theory is able to explain observed auctions, depending on how much flexibility is allowed in parameter choices.

Table 2. Statistical tests comparing theoretical estimates in different settings

Specification	N	Pearsons χ^2 Test (compares distributions)		K-S test (compares distributions)		t-Test (compares means)	
		$p \geq .10$	$p \geq .05$	$p \geq .10$	$p \geq .05$	$p \geq .10$	$p \geq .05$
Base: $\alpha = 0$, $v =$ Amazon	172	9.3	13.3	7.0	10.5	8.1	9.3
Value: $\alpha = 0$, $v =$ MLE	172	48.3	54.7	43.6	54.1	73.8	76.7
Risk: $\alpha =$ MLE, $v =$ Amazon	169	56.0	66.9	57.4	69.8	91.1	92.9
Full: $\alpha =$ MLE, $v =$ MLE	169	73.4	76.9	82.2	87.0	96.4	97.0

Notes: The number reported in each cell is the percentage of items for which the particular test statistic has a p-value larger than the threshold indicated in each column. N refers to the number of unique items. K-S refers to the Kolmogorov-Smirnov test.

The Base Specification model (risk-neutrality $\alpha = 0$, $v =$ Amazon price) provides predictions that fit the actual distribution well only for a small fraction of items, and the statistical tests show poor goodness-of-fit. The Value Specification models (risk-neutrality $\alpha = 0$, v is found by Maximum Likelihood) has limited improvement in matching part of the actual distribution, because the risk-neutral model cannot generate the observed hump shape in the distribution of final prices as shown in Figure 1. The statistical tests show a better goodness-of-fit compared to the Base Specification. The Risk Specification model (risk-loving, $\alpha < 0$, found by Maximum Likelihood, $v =$ Amazon price) estimates a risk parameter $\alpha < 0$ for majority of items and $\alpha > 0$ for 23 items. The negative alpha estimators indicate that bidders are mildly risk loving, primarily in the range of -0.003 to -0.03, with a few estimates as low as -0.09. Most were well below the estimated risk preferences of bettors at horse race tracks; only 9 items found α greater than -0.055 estimated by Jullien and Salanie (2000). The improvement in fit over the Base Specification is remarkable; the K-S test now only rejects 30% (rather than 90%) of the items at the 5% significance level as having observed distributions inconsistent with the Risk Specification. Note that the risk-loving model is able to match the hump shape of the distribution, as $f(t)$ will first rise then fall if α is sufficiently negative; thus, early bidders are more likely to bid than risk neutrality would imply. The Full Specification model (α and v are jointly estimated in Maximum

Likelihood) estimated α has a similar range as in the Risk Specification model, typically between -0.001 to -0.03, with only three items requiring $\alpha < -0.055$.[19] In most cases, the estimated valuation remained reasonably close to the Amazon price (on average, 15% greater) and statistical tests show goodness-of-fit improvement.

Figure 1. Theoretical and observed distribution of ending bids in four specifications

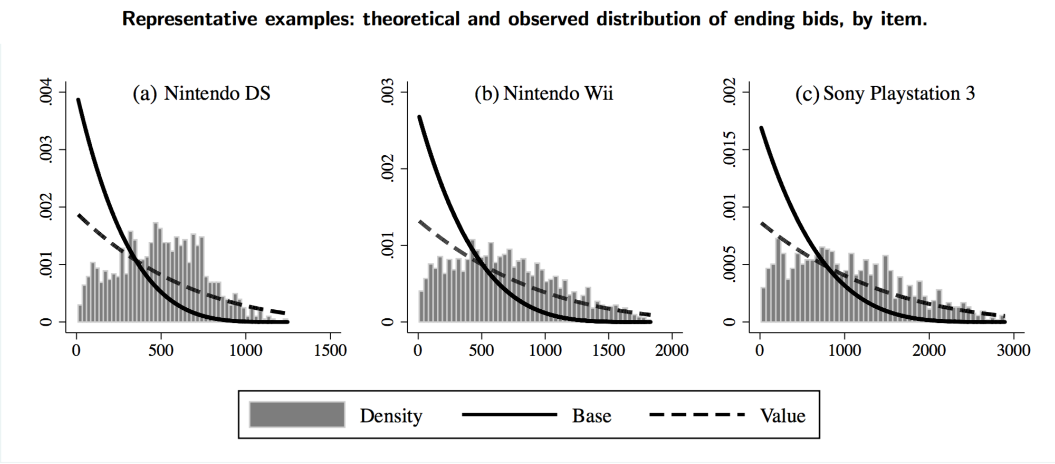

Notes: Density (bars) denotes the observed frequencies. Base (solid line) gives the theoretical frequency using the item's Amazon.com price for its valuation. Value (dashed line) does the same using the MLE-estimated valuation. In both, bidders are assumed to be risk neutral ($\alpha = 0$).

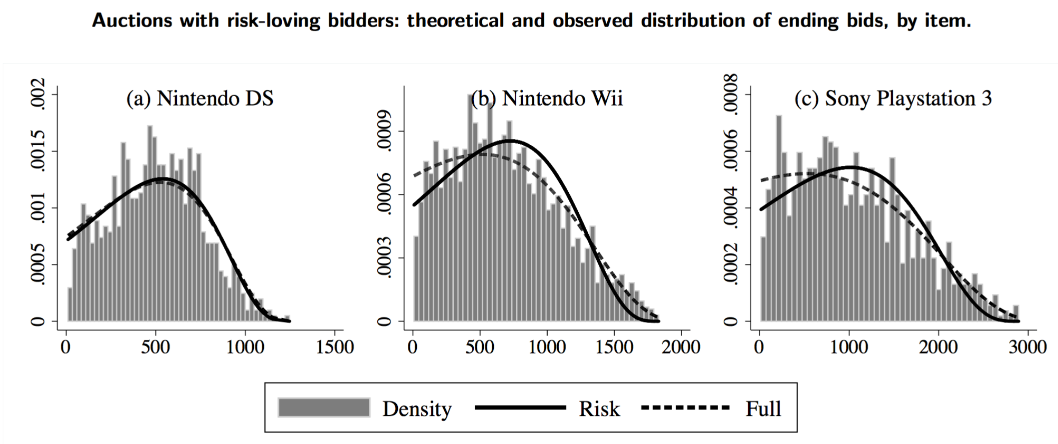

Notes: Density (bars) indicates the observed distribution of ending bids on each item. Risk (solid line) denotes the theoretical frequency using the item's Amazon.com price for its valuation and the MLE-estimated risk parameter. Full (dashed line) does the same using maximum-likelihood estimates for both α and v.

Thus it is concluded that the Base Specification is clearly inadequate, and that risk preferences contribute more towards explaining observed auction outcomes. Adjusting α achieved a much better fit than adjusting v; and even when both were adjusted, the Full Specification only explained a small additional set of mostly cheaper items. The Risk Specification is particularly satisfactory with more reasonable parameters, which explain perceived profits in most auctions.

Augenblick (2012) captured two datasets of Swoopo auctions: an auction-level dataset containing 166,000 auctions from September 2005 to June 2009, and a bid-level dataset containing 18,000 auctions from late February 2009 to June 2009. The first empirical finding shows that a penny auction consistently produces revenue above the market's value, which does not match the zero profitability finding in the theoretical analysis. On average, bidders collectively pay 51% over the adjusted value of the goods in an auction. Retail prices of items sold on Amazon are used as the adjusted value, which are on average 79% of Swoopo's listed prices. Moreover, the profit margin of an auction varies across item types; for instance, consumer goods and bid packages generating average profit margins of 33% and 201% respectively. To understand the deviation from the theoretical model, Augenblick (2012) investigate fit of the data to the predicted hazard function at both the auction level and individual level.

In the auction level data, the empirical hazard function is very close to the equilibrium prediction by equilibrium analysis at the beginning of the auction. However, deviation increases significantly over time. Although the empirical profit margin starts near zero, it rises to over 300% through the course of the auction. If the auction runs long enough, a bidder is willing to pay \$0.75 for a stochastic good with an expected value of only \$0.25. Thus, rather than having a constant profit throughout the auction, the auctioneer makes a large amount of instantaneous profit at the later stages of the auction.

While auction-level empirical hazard rates do not contain information about bidder heterogeneity, the probability of an individual bidder exiting an auction, given the number of bids already placed in the auction up to that time period is observable – this probability is called the *pseudo hazard rate*. We consider a bidder as exiting an auction if he does not place anymore bids in the later periods. The theoretical model suggests that the probability that a player does not bid should rise slightly as the auction progresses, if the number of users in the auction stays constant. However, the observed pseudo hazard rate declines significantly as the number of bids placed in the auction increases, which is not caused by the decline in number of active bidders.

Since the above results are aggregated over all bidders, the concern is whether the heterogeneity of bidders or auctions are the driving causes. Linear regressions are performed to regress the probability of a bidder dropping out of an auction, on the logged number of past bids, various fixed effects and other control variables such as bidder experience levels. The results have shown that the probabilities of auctions ending are less than the theoretical predication, because as the auctions progress bidders are less likely to leave the more bids they have already placed within the auction. Any explanation of these results must include a factor that changes as an auction progresses, since any fixed effect would only drive constant deviations from the theoretical predication. Augenblick (2015) recommends the naive sunk cost fallacy as the most intuitive explanation: bidders become less and less willing to leave an auction as they place more bids, even though those costs are "sunk".

Augenblick (2015) also discusses the two types of learning of bidders: learn to either stop bidding, or learn to bid in a way that generates consistent profit. It was found that learning occurs, but the latter learning process is slow. The majority of bidders learn to stop playing the game quickly (75% of bidders stop bidding before placing 50 bids, and 86% stop bidding before placing 100 bids). The other learning process, according to which more experienced bidders gain higher expected payoffs, was examined by a set of regressions of instantaneous profit on experience and other variables. Regression results indicate a positive, concave relationship between the profit from a bid and the bidder's experience level. For instance, a bidder with no experience has an expected loss of $0.60 cent per bid, while some heavily experienced bidders gain slightly positive expected payoffs per bid (it takes 10,000 prior bids to raise an expected, instantaneous profit per bid to near zero). The most active of bidders (11%) contribute to 50% of Swoopo's profit, and as the learning process is slow the experienced bidders place a large amount of negative-return bids before they learn to break even. There is also evidence showing that experienced bidders learn to apply aggressive bidding strategies to increase their winning probabilities. The aggressive strategy demonstrated by such bidders is to bid immediately whenever possible, showing their determination to win in a war of attrition.

Wang and Xu (2012) argue that a penny auction is not a sustainable selling mechanism, as experienced and strategically sophisticated bidders exploit penny auctions, while inexperienced bidders who happen to be biased receive negative feedback consistently so that they learn to quit quickly. Using a behavioural game theory approach emphasising that players in a new game may be inexperienced and have limited strategic sophistication, the complete bid history at BigDeal, a major penny auction website, is studied. Wang and Xu (2012) argue that bidders' behaviour in penny auctions is better understood through the lens of learning and strategic sophistication than through an equilibrium model that presumes all bidders are fully rational. This suggests that a penny auction is not a sustainable selling mechanism and requires a continuous supply of new bidders.

The empirical evidence shows that the auctioneer profits from a revolving door of new bidders, but loses money to experienced and strategically sophisticated bidders. The vast majority of new bidders who join the website on a given day, only play in 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 reflects the simple logic of individual rationality: no matter how effective a penny auction might be in exploiting bidder biases, it offers an immediate outcome (win or lose) to bidders, so losing bidders can quickly learn to stop participating. A very small percentage of bidders are experienced and strategically sophisticated – winning most of the auctions and earning substantial profits. Thus, penny auction websites cannot survive without continuously attracting new customers.

The fact that most penny auction websites impose win limits likewise suggests that some bidders play better than others. Wang and Xu (2012) attempt to link bidders' winnings or losses with their strategic sophistication by measuring a bidder's strategic sophistication by the frequency of their placing a bid in the middle of the countdown period during an auto-extension of the auction duration. Placing a bid in the middle of the timer is assumed to be inferior to placing the bid at the end of the same period, as the latter has some informational advantage. This measure of strategic sophistication is predictive of experienced bidders' overall winnings or losses, and shows that sophisticated bidders learn to earn more money per auction as they play in more auctions, but unsophisticated ones do not. Two types of learning are discussed: learning to play better and learning to quit. Most new bidders learn to quit quickly, which suggests that a firm's ability to exploit consumer biases is limited by consumer learning.

While most bidders are exploited by penny auctions, some strategically sophisticated bidders take advantage of the penny auction format. For instance, the winner of a penny auction is often not the bidder who places the most bids. 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. 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 have already been placed in the auction. Another basic equilibrium result, that the number of bidders does not affect auction revenue, also does not hold; for instance, when using various proxy variables, it is shown that an auction’s profit increases in the number of potential bidders.

Stix (2012) discusses bidder retention, a long term problem with penny auction mechanisms. Because as soon as the supply of new or inexperienced bidders runs out, the majority of the auctioneer’s income would evaporate. A few potential solutions include setup of win limit rules and special type of auctions that only selected bidders can participate. This reduces the auctioneers’ reliance on inexperienced bidders for profit. For example, beginner auction is a special type of penny auction offered by QuiBids, where only bidders who have not won any auctions may bid. Beginner auctions tend to be numerous and are usually for inexpensive items, as auctioneers attempt to give every new user a win early on in their bidding career, in order to encourage a bidder to become more committed to the website. Introduction of a buy-now price option allows bidders to contribute money they have spent in a lost auction towards the listed price of that item. This provides an extra sense of security to the bidder by limiting their loss to the difference between the marked up listed price and item valuation, (while also limiting an auctioneer’s profit) and bidders are less likely to be discouraged by the immediate winning or losing outcome of the auction. Most penny auction sites recognise the problem of bidder retention.

In analysis of QuiBids’ profitability based on data collected on 37,233 auctions, standard penny auctions without a buy-now option are confirmed to be profitable on average, with an estimated average profit margin of 55.11%. Note that despite the high profit margin, the median profit per auction is still slightly negative, which is likely to be the result of one of QuiBid’s attempts to solve the customer retention problem. The negative median profit encourages new bidders by letting them gain positive payoffs on smaller items in an attempt to get them to bid in auctions for more expensive items, in which QuiBids makes an enormous profit. For instance, the most expensive 2.5% of items auctioned by QuiBids auctions generated almost 43% of its profit. Running many unprofitable auctions of inexpensive items seems to increase customer retention effectively - 76% of bids are submitted by experienced bidders who have placed more than 50 bids - that may give QuiBids large profits in higher-priced auctions later.

Stix (2012) argue that QuiBids survived, while Swoopo failed, due to differences in their initial rules on bidder-auction ratios and voucher bids[4], and claims that some optimal bidder-auction ratios exist.

The original penny auctioneers want the number of bidders per auction to be high enough to generate a significant profit. If this ratio is too high, the proportion of winners will be too small, such that new bidders would soon learn to stop participating, due to repeated losses before they have an opportunity to win an auction. Also, the bidder-auction ratio would not be maintainable due to a shortage of bidders, assuming there is no infinite supply of new bidders. The introduction of a buy-now price leads to both better bidder retention and lower profits. However, utilisation of the buy-now feature limits the profit the auctioneer can generate from each bidder, which may decrease the overall profits. Thus a significant increase in optimal bidder-auction ratios can be expected, and penny auction sites are inclined to increase their ratios to regain lost profits.

The empirical evidence suggests that QuiBids actively maintained their chosen ratio by adjusting the numbers of auctions running at different time periods, and their average of 347 bidders in each auction is significantly larger than the ratio of 42 bidders per auction, prior to the Swoopo's introduction of a buy-now price feature. Swoopo's failure to increase its bidder-auction ratio in penny auctions utilised with buy-now price option would cause a significant drop in profit, which contributed to its bankruptcy.

Stix (2012) shows although the utilisation of the buy-now feature educes 32% of estimated revenue, it has a positive effect of on customer retention such that QuiBids is still able to generate positive profits. QuiBids' rule that voucher bids are ineligible to contribute toward a buy-now price is also considered by Stix. This feature significantly reduces the true value of voucher bids. It is shown that if bidders only care about monetary payoffs, they should avoid bidding in voucher bid auctions. If bidders

[4] Voucher bids are listed as an auction item on QuiBids. Sold in packages, voucher bid auctions allow users to bid in an attempt to win more bids for use on other auctions. Voucher bids are auctioned off using the same penny auction system as other products on QuiBids. This means that a user who wins a voucher bid pack auction on QuiBids gets a predetermined number of voucher bids. The only difference between purchased bids and voucher bids is that voucher bids are not eligible to contribute their cost to the Buy-Now feature.

rationally ceased participating in these bid voucher auctions, QuiBids would be making a loss. The fact of QuiBids being profitable indicates bidders are not behaving rationally.

Zheng et al. (2011) argue that the penny auction format and its extremely low winning price is a double-edged sword, which not only attracts the participation of new bidders, but also makes it challenging to retain existing customers, since only one winning bidder may derive a positive surplus, whereas all other bidderses suffer loss from the bidding fees incurred. Using a field experiment with a 16-week sample period, Zheng et al. (2011) analysed how auction rules can improve overall consumer retention and long-term bidding participation. Three types of restrictions were implemented to reduce the participation capacity of the frequent aggressive bidders, such that other bidders, especially the inexperienced ones, have more opportunities of winning. Firstly, each bidder is allowed to win a maximum of eight auctions within 28 days, directly limiting bidder participation and distributing winning chances to a broader set of customers); Secondly, no bidder is allowed to bid for the same type of item more than once within 28 days, for instance, bidders are unable to get a lot of bid credits at low cost through winning bid pack auctions; and finally each bidder is allowed to participate in only X concurrent auctions where X is 8 minus the number of auctions won in the past 28 days, which further restricts the number of concurrent auctions that a bidder could participate in.

Three regression models were estimated: a simple OLS regression model, a panel logistic regression on bidder retention, and a negative binomial regression on bidder participation with dummy variable of rule changes (indicating if the three rules have been implemented) as an independent variable, and dependent variables including consumer surplus, a modified Gini coefficient to measure disparity, a bidder participation dummy variable to measure consumer retention, and the number of auctions participated in, as well as the number of bids submitted by a unique bidder in one week, to measure auction participation. Control variables included bidder type (frequent bidders and occasional bidders), consumer surplus history, surplus gain and surplus loss that a bidder obtains in a week, cumulative life-time total surplus gain and surplus loss, number of products auctioned in one week, and popularity of products. The results of these analyses show that a skewed or lopsided consumer surplus distribution is highly positively correlated with bidder attrition. The implementation of bidding restriction rules facilitates surplus distribution in a more equal manner, such that the Gini coefficients for consumer surplus drop by 10% after the rule changes. These coefficients ranged between 0.96 and 0.99 before implementation of the rules. Note that 0.99 indicates an extremely unequal distribution, which is consistent with the empirical fact that a small proportion of the bidders earn most of the consumer surplus.

There is evidence that overall customer retention rates are higher after the rule changes. The coefficient of the dummy variable of rule changes has opposite and significant effects on occasional bidders and frequent bidders. It is significantly positive for occasional bidders (marginal effect on bidding probability of 0.11), but not statistically significant for frequent bidders (reducing their participation probability by 4%). Since there is a much larger number of occasional bidders, the overall effect is positive and beneficial to the auctioneer. The total number of auction participants, and the total number of bids placed in auctions, both increase after the rule changes. All estimated model coefficients on the control variables are very significant and in most cases show that implementation of the rules has a positive effect on occasional bidders but a negative effect on frequent bidders, both in terms of the numbers of auctions and bids, while the benefits exceed the loss or reduction in the auction bidding and participation activities by a small group of aggressive, frequent bidders.

Goodman (2012) uses the risk-neutral Nash equilibrium benchmark model of Augenblick (2012), and both papers analyse bid-level data. While Augenblick (2012) describes the market, Goodman (2012) focuses on individual-bidder reputations and behaviour, by identifying and analysing the successful strategies employed by some users. Two datasets of Swoopo auctions are collected with one bid-level dataset of over 52 million submitted bids from 64,000 auctions on Swoopo.com, covering 287,000 unique users, and broader auction-level data for auctions listed by Swoopo from 14 June 2019 to 13 June 2010. There are a few noticeable facts: Swoopo makes positive profit of more than $1 million per month, not only novice users suffer losses, even experienced bidders frequently lose money. Over 90% of 1,387 unique auction items are sold fewer than 100 times, while vouchers for 50 bids and vouchers for 300 bids – the two most common items – are jointly sold over 9,000 times. Only a small proportion of the 287,000 users bid heavily with the top 1% of users (by total bids placed) placing nearly half of all bids, while the bottom half of users only placed 3.4% of all bids, and the median users places 33 bids. Bidders appear to collectively overbid, as the average returns are negative (based on listed price of auction item and bidding fee), and their outcomes are highest for auctions with reliably priced items (that is if, a contemporaneous Amazon price is accessible) and lowest for auctions with unreliably priced items such as bid pack vouchers.

Analyses are executed in three levels: auction-level, bid-level and bidder-auction-level. The auction-level analysis focuses on survivorship, as the duration of an auction determines collective payments from bidders. To allow comparisons among auctions of different price levels and auction parameters, variables are normalised; for example, by measuring the number of bids (or periods) that have elapsed as a proportion of the maximum possible number of periods, and expressing the profit margin for an auction (sum of revenue from bidding fees and winning bid price, minus value of the item) as a percentage of the value of the item. The objective is to analyse whether Swoopo users bid inefficiently. For instance, if auctions last too long on average, the expected total revenue collected from bidders will exceed the value of the auction items. Auctions continue past the beginning stages more often than benchmark equilibrium would suggest, and the overall level of bidding is too high, which implies a negative expected return from bidding. On the other hand, the distribution of the profit margin shows that Swoopo sells goods below valuation 40% of the time, while the mean of the profit margin is 67%. Note that valuation of an auction item is considered to be the Swoopo suggested retail price (in 61% of total auctions) unless contemporaneous Amazon.com pricing can be found.

In bid-level analysis, the expected immediate return on a bid is considered to be the probability of winning the auction in the next period multiplied by the value of winning, minus the cost of bidding. Note that the benchmark equilibrium implies that this return equals zero at all times.

The key independent variables for analysing the expected return on a given bid are broken into four categories: user experience and user-fixed effects that depend only on the bidder; adjusted bid number and bidding runs that measure contemporaneous bid histories; bidding method, bidding speed and time, and day-fixed effects that record how the bid was submitted; and item fixed effects that account for auction characteristics. User experience is useful as a proxy for bidder strategies and characteristics. For a player, user experience measures how many lifetime bids he has submitted. The data provides an average return of -25.5 cents per bid, indicating possible correlation of user experience. This is examined by splitting observations of bids into bins based on the pre-bid experience level of their submitters, showing that inexperienced bidders do considerably worse than veteran bidders, and suggesting a quadratic relationship between mean outcome and log experience. Examination of results suggests that veterans employ superior strategies. New users are largely unaware of which bidding techniques are more successful, but participants who continue bidding figure it out. Learning appears to occur across all types of bidding strategies, and tends to be a more significant factor than survivorship bias in explaining the behaviour discrepancies across experience levels.

Recall that bidders are assumed to not condition on bid histories in the benchmark model implying that hazard rates are history independent. It is more realistic to allow bidders to adjust their bidding decisions based on past histories and future expectations, e.g., if a bidder perceives a high probability of future bidding by some aggressive participant, his expected chance of winning by placing a bid decreases, which would lower his probability of bidding in the current period. Possible signalling devices in a penny auction include frequency of placing bids, bidding runs, bid methods and bidding speed.

Bidding frequently signals an interest in the auction and a willingness to spend. Bidding runs is a running count of how many bids in a given auction the bidder has already submitted, scaled by the maximum possible auction duration measured. Relationship between winning probability and bidding proportion (ratio of bidding runs to total bids) is estimated that when a bidder increases his share of bids within an auction, his odds of winning increase even faster. Bidding speed is a measure of wait time, and records how much time elapsed between a bid and its immediate predecessor, which is difficult to interpret. While short wait times are often signals of strength, alternative bidding strategies such as sniping (waiting until the auction timer approaches zero to submit a bid) make interpretation complicated. A bidding method refers to the three mutually exclusive ways a bid can be submitted in Swoopo: as a single bid (placed manually), a bid placed through BidButler (Swoopo's auto-bidding agent), or by telephone, although an obvious conclusion of the best bidding method cannot be drawn from empirical results: for example, BidButler provides a higher outcome per bid and increases winning probability.

Since the expected immediate return on a bid seems to be correlated with bid histories and bidder characteristics, Goodman (2012) ran five regressions to test their correlation. A regression including only experience variables shows the value of experience is positive and increasing at a log-quadratic rate, which confirms the existence of an experienced premium. After the introduction of bid history and bidding manner variables, the coefficient of log experiences changes from positive to negative, which indicates that they do capture part of the strategies employed by the experienced bidders. This provides evidence that return per bid does vary widely from bid to bid, and its variance can be largely accounted for by the aggressive signalling variables of bidding runs, bidding speeds and bidding methods, while a normalised period of rivals' bidding runs and cross effects have significant impacts as well. All aggressive bidding behaviours are shown to correlate with higher outcomes, while bidding runs appear to be the most effective tool for building an aggressive reputation. Effects of rivals' bidding runs and cross effects are clearly not linear across the categories and more difficult to interpret. One explanation is sluggish adjustment, that bidders are slow in recognising aggression of their opponents, but they eventually notice it and start bidding less themselves. It is of similar nature to the Game of Chicken (if only one player in a penny auction is aggressive, he will have advantage; whereas if multiple players are aggressive, they may all lose in a costly bidding war).

Note that bid-level data do not have independence among observations and an expected immediate return of a bid is not identical to the total impact of a bid in an auction. The expected immediate return of a bid is the probability of winning, times the value of winning, minus the cost of bidding, while there are other impacts of a bid. Analysis of data at the bidder-auction-level, wherein all bids placed by a bidder within the same auction are aggregated, allows for a more comprehensive look at bidder strategies and their impacts, with improved observation independence. Results of bid-level analyses indicate that aggressive strategies are correlated with higher outcomes, and that these bid-level variables need to be presented in an appropriately aggregate manner for bidder-auction-level analysis. For instance, instead of choosing the total outcomes of each bidder (number of bids placed by each bidder vary in most auctions), an outcome per bid (ratio of a bidder's total outcome in an auction and number of total bids that he places in the auction) is chosen as the dependent variable. Bidding methods and bidding speeds are variables specific to a particular bid, which need to be adjusted to fit the current framework. Six regressions are executed similarly as in bid-level analysis, and the results indicate positive effects of experience on outcomes per bid, and regression testing bid-history and bidding-method variables show the strategic variables capture part of the experience premium.

All three levels of analysis show that aggressive bidding is the key: bidding frequently, bidding quickly after the previous bid, and bidding through the automated BidButler service, all significantly increase the likelihood of winning and expected returns. Playing aggressively signals commitment to future participation, as players bid less when they think their chances of winning with a given bid are particularly low, and signals of aggression intimidate rivals into dropping out, thereby increasing one's own win probability.

The impact of two other bidding strategies that are not tested in the above regressions are also considered: first bids and simultaneous bidding. Since all bids are not equally prominent, for instance, the first bid placed by each bidder in an auction is a player's first opportunity in reputation establishment, and empirical evidence shows that how and when the first bid is placed has a considerable impact on overall returns in an auction. Data show that entering an auction early and entering with BidButler are associated with higher win probabilities; for instance, initial bidders win 8.4% of all auctions, earning positive returns of 9.2 cents per bid, and initial bidders who enter auctions with BidButler win 16.2% of all auctions, earning positive returns of 13.7 cents per bid. While the initial bid of an auction generates a negative average immediate outcome of -9.5 cents per bid, the advantages stem from providing informative signals about bidders' likelihood of bidding in the future, which intimidates rivals into dropping out, as on average, a bidder places 25 bids in a penny auction, while initial bidders and initial bidders entering with BidButler place 54 and 108 bids respectively.

Similarly to Augenblick (2009), Goodman (2012) found that aggressive bidding strategies are the key to better outcomes. Although Byers et al. (2010) reached conflicting results—there are not clear gains to assertiveness and the most aggressive bidders are the least successful. Goodman (2012) argues that this discrepancy stems from incongruous definitions of aggression, and claims that aggressive players tend to be more successful overall. The reason is that aggression builds reputations and signals a commitment to future bidding, which intimidates rivals from bidding and raises their own outcomes. Consequently, bidding frequently, which appears to be the most aggressive tactic, bidding quickly, and bidding through BidButler are all associated with higher average returns. Experienced players are also more likely to enter auctions early to avoid competing with other bidders because of simultaneous bidding, thereby increasing their own returns. Crucially, data indicates that these tendencies result largely from learning, not survivorship bias.

The original dataset used in this paper was collected from BidCactus.com, one of the longest standing penny auction websites, by Ruby scripts. Over a time frame of four years, 572,400 auctions from May 2009 to April 2013 were retrieved for the start of BidCactus. There were 253,464 bidders placing 60,795,092 bids in auctions for 1,742 unique items. All publicly observable history for each auction was recorded, including names and Recommended Retail Price (RRP) of auction items, usernames of winners and their winning bids, usernames of all participating bidders and the number of bids they had placed.

All auctions hosted by BidCactus start at an initial price of $0, then bidders may choose to increment the price by $0.01 by placing a bid by either paying a non-refundable $0.75 by credit card, or using their pre-purchased, non-refundable bidding credits, which also extend the duration of the auction by resetting the countdown clock to 30 seconds – if the auction has passed its preset ending time (usually 3 days from the start time). An auction ends when the countdown clock reaches 0, and the last bidder wins and pays the end price, while all participating bidders have paid for every bid placed during the auction's course.

Bidding fees are the primary revenue source of BidCactus, which earned substantial profits over the sample period. As the sole seller of all products on its website, it earned a total revenue of $50,635,419 with a profit margin of 53.38% (cost estimated by the listed RRP in auction), and 90% of the revenue was from bidding costs paid by all participating bidders.[5] A simple example will help to understand the profitability of the mechanism more intuitively. In a BidCactus auction, each bid increments the current price by 1 cent, such that for every $1 increase in the winning price the seller collects additional revenue of $75 from bidding fees. Thus, as long as the winning price reaches 1.32% of the seller's valuation, he would break even, and when the winning price reached 3.95% of the valuation, his profit margin would be 200%. On average, the winning price of a BidCactus auction reaches 1.95% of the items RRP. On the other hand, although some auctions ended with a high profit margin (23% of its auctions

[5] We compute the total profit margin of the auctioneer as the ratio of total profit (total revenue of winning prices and fees paid for bids by bidders in all auctions minus total cost estimated by RRP listed) and total costs, profit margin of any particular auction as a ratio of seller's profit (sum of winning price and bidding fee collected from all participants in the auction, minus RRP) and RRP, average profit margin of a set of auctions (e.g. auctions in a particular category) as the unweighted average of an individual auction profit margin.

resulted in a profit margin of over 100%), not all auctions were profitable and BidCactus made a loss in nearly 60% of its auctions.

Despite its profitability, BidCactus still looks attractive to many bidders, especially the newcomers, who observe that the current price of an auction item is usually remarkably low and the winning price of a closed auction is also usually low. For instance, the median and mean winning price is only 0.96% and 1.95% of the item RRP.

Figure 2. Seller's profit margin distribution in percentages and seller‘s profit distribution in dollars

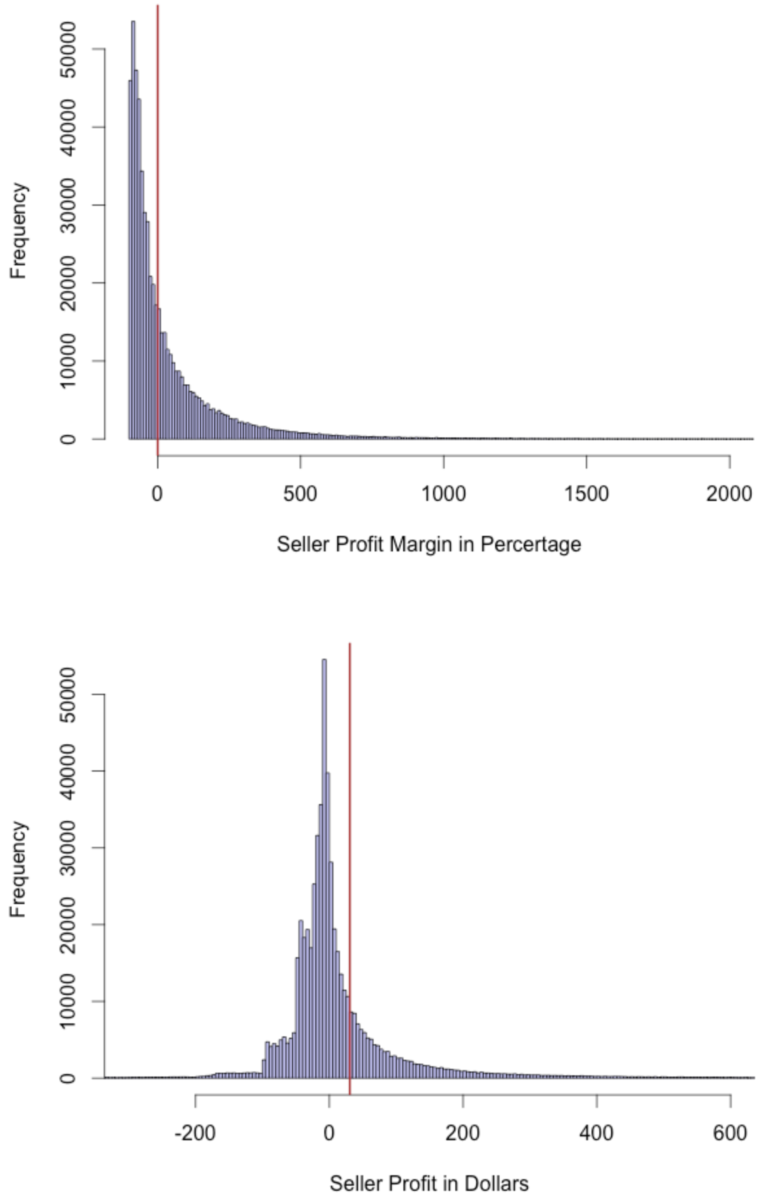

Figure 2 shows the majority of the distribution of a seller's profit margin and profit in dollar terms[6] in each of all BidCactus auctions, with their long tails on both sides cropped, and the red vertical lines indicating the means. The distribution of a seller's profit varies largely from -99.80% to 23,045%, with 90% of the observations[7] falling within the range of -83.15% and 223.09%, and a mean of 48.16%. The distribution of a seller's profit varies largely from -$1,764.71 to $31,966.81, with 90% of the observations[8] falling within the range of -$83.15 and $223.09, and a mean of $30.60.

In total, there were 1,743 unique item prizes that BidCactus sold, and the auctioneer chose to run auctions to sell some of the items more than once, with this number of repeated auctions varying extensively across different types of items. For instance, 6,671 of $50 Visa Gift Cards were sold in 6,671 auctions in four years. The top ten most commonly auctioned prizes including nine bid pack auctions, jointly sold over 200,000 times, while 83% of item types were sold fewer than 100 times. Twenty six types of items have been repeatedly auctioned over 5,000 times with an average profit margin shown in Figure 3. Twenty five of the 26 types generated a positive average profit margin, which are all auctions for bid credit packs and gift cards. The most repeated, listed bid pack auction type is for 100 bid credits, which was listed 30,619 times, generating an average profit margin of 113.68% for the seller. The most repeated, listed gift card auction type is for a $50 Shell gift card, which was listed 9,100 times, generating an average profit margin of 45.52% for the seller.

The types of items are named and sorted into five categories. For instance, a $50 Shell gift card is an item type, which belongs to the gift card category.

1. ***Bidpack***

   Auctions for prizes of packs of bidding credits:
   e.g. 22,887 auctions for 30 bidding credits generate an average profit margin of 77.76%.

---

[6] Seller's profit of an auction is computed as seller's revenue from winning bid and bidding fees from all bids placed in the auction, minus listed RRP.

[7] 90% of observations are taken from auctions with seller's profit margin between its 5% and 95% percentiles.

[8] 90% of observations are taken from auctions with seller's profit (in dollars) between its 5% and 95% percentiles.

2. ***Gift cards***

Gift cards with a clearly defined face value from $10 to $500 (78% with a face value of $50 and below), issued by famous retailers and service providers:
e.g. 1260 auctions of Kmart $25 gift card generate an average profit margin of 26.56%.

3. ***Small goods under $100***

Physical items with RRP under $100
e.g. 197 auctions of a Calphalon Nonstick Stir Fry Pan with RRP $49.95 generates an average profit margin of -50.23%.

4. ***General items***

Physical items with a RRP of $100 and above, excluding consumer electronic items with a RRP of $399 and above:
e.g. 215 auctions of Beats by Dr. Dre Solo Headphones with RRP of $199.95 gives the seller an average profit margin of -9.69%.

5. ***Expensive consumer electronics***

Expensive consumer electronic items from a RRP of $399 go into this category, including Apple iPhones, iPads, iMacs, Game Consoles, DSLR cameras, Notebook computers, HD Televisions and high-end GPS, as they are popular among bidders and often draw aggressive bidding:
e.g. 119 auctions of an Unlocked 16GB Apple iPhone 5 with a RRP of $649 gives the seller an average profit margin of 222.66%.

Figure 3. Items repeatedly auctioned over 5,000 times

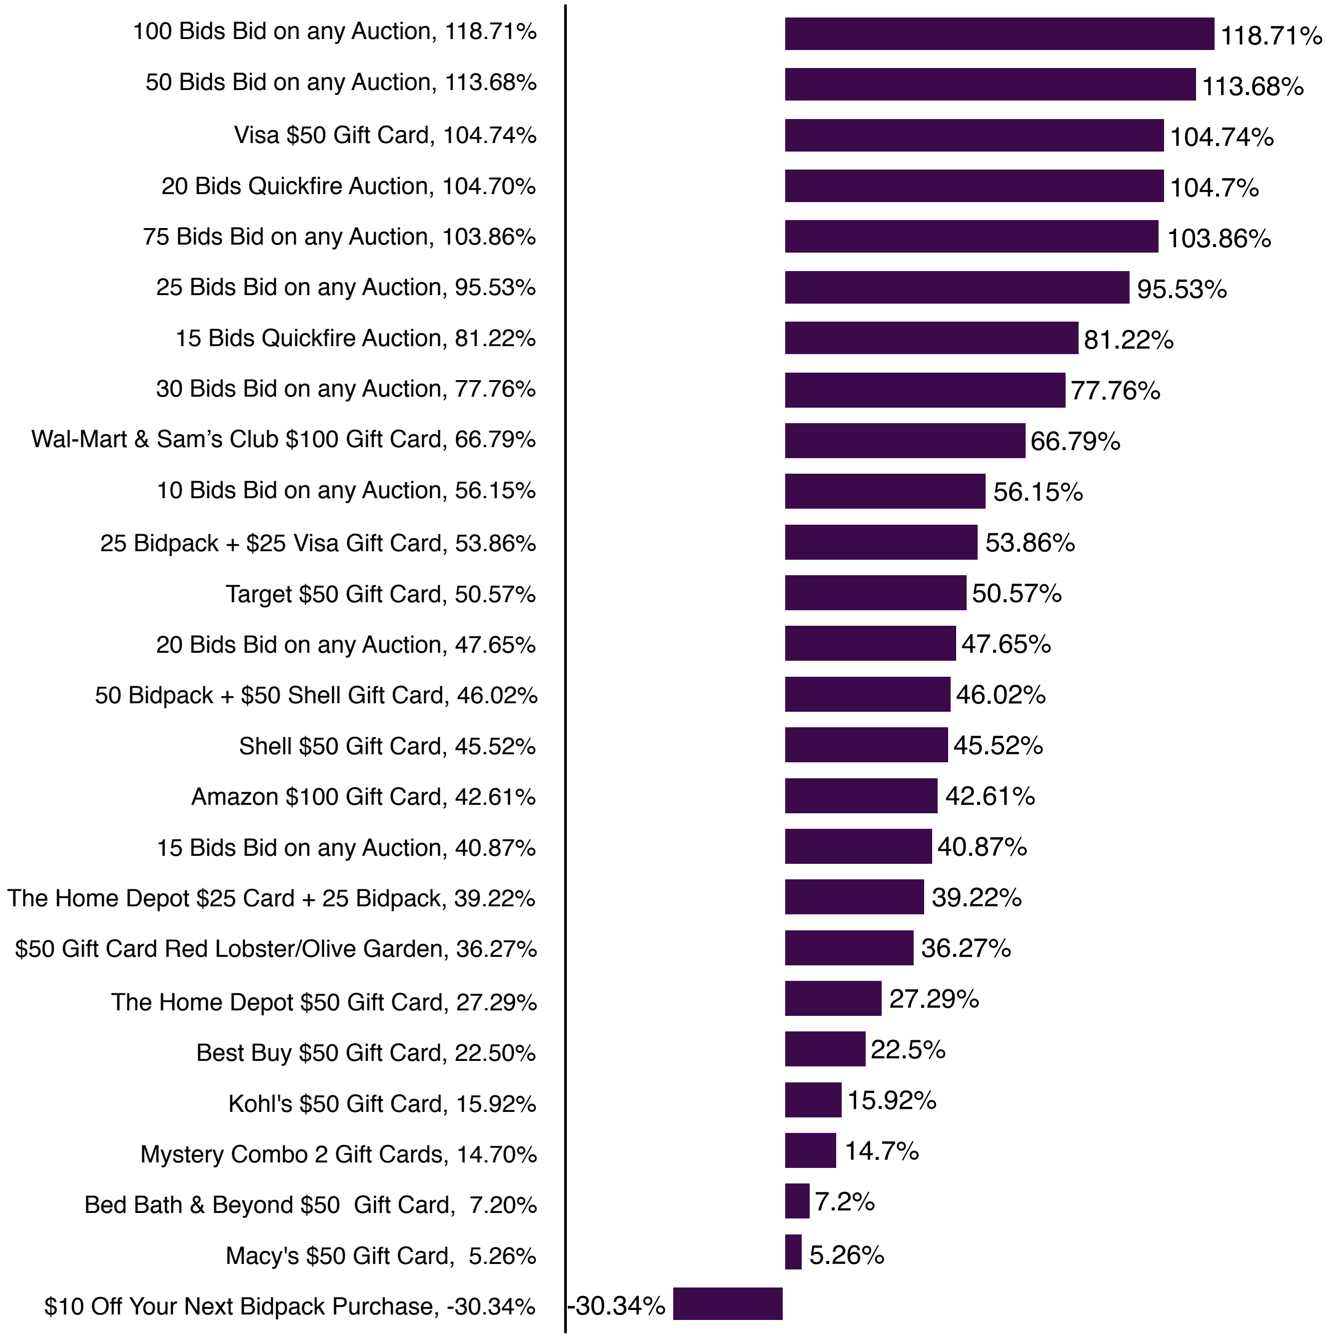

Figure 4 shows the basic facts of each of the five auction item categories. Categories such as general items and small goods under $100 have the most item types of 692 and 654 respectively, while items types in categories bid pack and gift card are listed most frequently with total observations of 287,372 and 243,619 auctions respectively. The seller's average profit margin varies extensively among different item categories. The weighted average profit margin of a category is computed by the ratio of total profit earned in all auction selling items in this category in dollars, and total costs of all items sold based on RRP. The unweighted average profit margin of a category is the average of the profit margin of each individual auction in this category, i.e. average unweighted on valuation of auction items. The expensive consumer electronics category generated the highest weighted average profit margin of 119.19%, and the small goods under $100 category generated the largest loss of -38.21% weighted average profit margin. The top 10 most repeatedly listed item types of each category are included in Appendix 1.

Figure 4. Basic facts of auction item types in five categories

Category	Item Types Counts	Observation Counts	Weighted Average Profit Margin	Unweighted Average Profit Margin	Most Repeatedly Listed Item	Number of Time Repeated	Average Profit Margin
Bidpack	123	287,372	98.64%	72.99%	50 Bids on Any Auctions	30,619	113.68%
Gift Card	172	243,619	36.09%	12.39%	$50 Shell Gift Card	9,100	45.52%
Small Goods Under $100	654	23,660	-38.21%	-37.05%	Keurig Mini Plus Single Cup Brewer	472	-20.22%
General Items	692	18,799	-0.66%	-24.57%	iPod touch 8GB 4th Generation	719	86.11%
Expensive Consumer Electronics	102	2,742	119.19%	111.86%	Apple iPad 2 16GB WiFi + AT & T 3G	312	209.45%

Figure 5. All item types VS those repeated listed 100+ times in five categories

Category	All Item Types Counts	Item Types Counts (Repeated Listed 100+ times)	Weighted Average Profit Margins (All Item Types)	Weighted Average Profit Margin (Repeated 100+ times)
Bidpack	123	74	98.64%	98.66%
Gift Card	172	125	36.09%	36.40%
Small Goods Under $100	654	41	-38.21%	-32.56%
General Items	692	36	-0.66%	26.59%
Expensive Consumer Electronics	102	7	119.19%	164.82%

Figure 5 compares a seller's (weighted) average profit margin in items types of different item categories for all item types and those that were listed for 100 or more items. The seller did choose to list some item types more often than others, e.g. only 7 out of 102 types of expensive consumer electronics were listed 100 times or more, generating a higher average profit margin of 164.82% compared to 119.19% of all item types in this category. Another significant improvement in an average profit margin is in the general items category, where 36 out of 692 types of general items were listed 100 times or more, thus generating a positive average profit margin of 26.59% compared to the overall loss of -0.66% in this category.

Figure 6 compares distributions of a seller's profit margin between -150% to 250%, where profit margins of most auctions fall in all five item categories[9]. The category of small goods under $100 – the one with the lowest average profit margin – has the highest negative peak in profit/loss. Although all five item categories have their peaks at profit/loss between -100% and 0, long tails after 0 yield their overall profit margins upwards.

Figure 6. Distribution of seller's profit margin in item categories

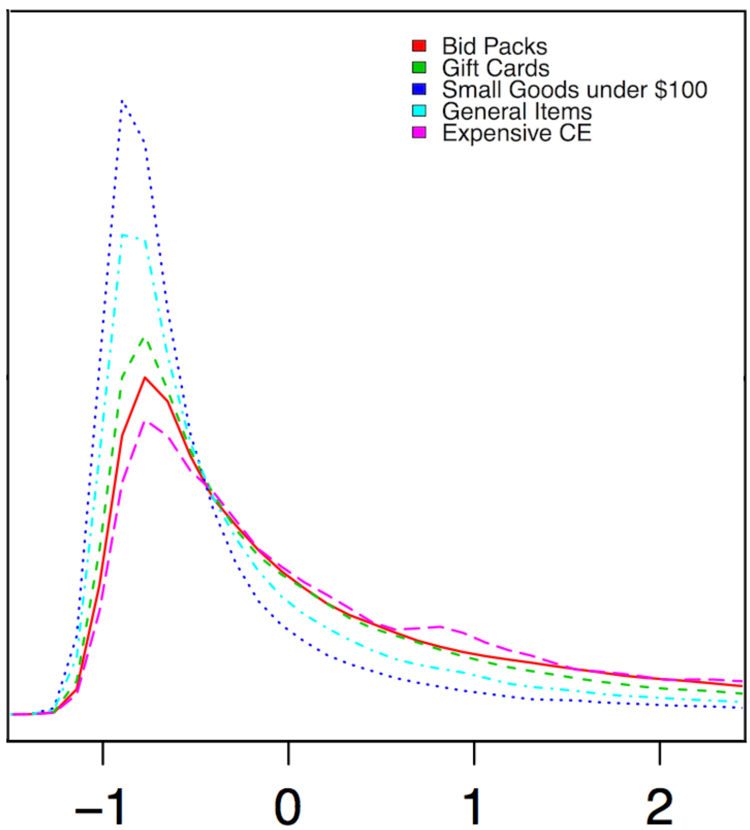

[9] The rest of the graph is trimmed and the full graph is included in Appendix 1.

Figure 7 shows the distribution of the number of auctions participated by bidders. Over half of all bidders never participated in more than six auctions, 12% of bidders quit after only one auction, and 98% of bidders played less than 100 auctions.

Figure 7. Numbers of auctions participated by bidders in total bidder population

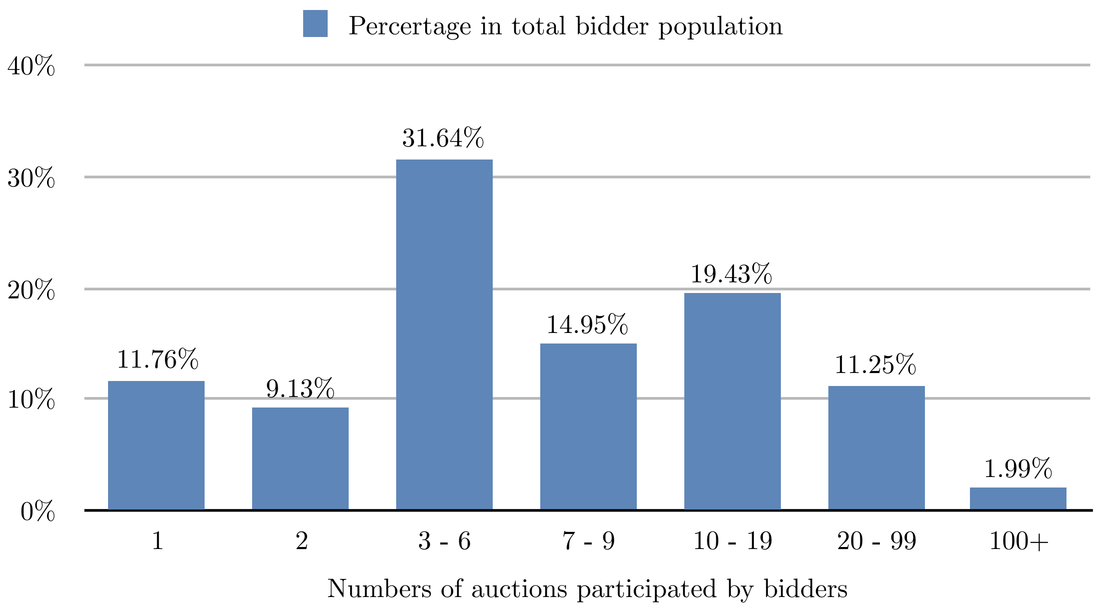

Bidders also appear to learn through experiences; they either learn to quit, or learn to play better. Figure 8 shows on average, the *winning probability* (ratio of number of auctions a bidder has won to the time point, and total number of auctions he/she participated in by placing at least one bid), the *overall net payoff per bid* (ratio of total net payoff and total number of bids placed in all auctions participated, where *total net payoff* is computed by the sum of RRP with winning price deducted from auctions won by a bidder, minus bidding fees for bids placed in all auctions he/she has participated in), and the *bid ratio* (ratio of number of bids placed by a bidder and total number of bids in an auction), which seems to be positively associated with bidder experience levels in terms of numbers of prior auctions participated in.

Figure 8. Bidders' return against experiences (number of prior auctions)

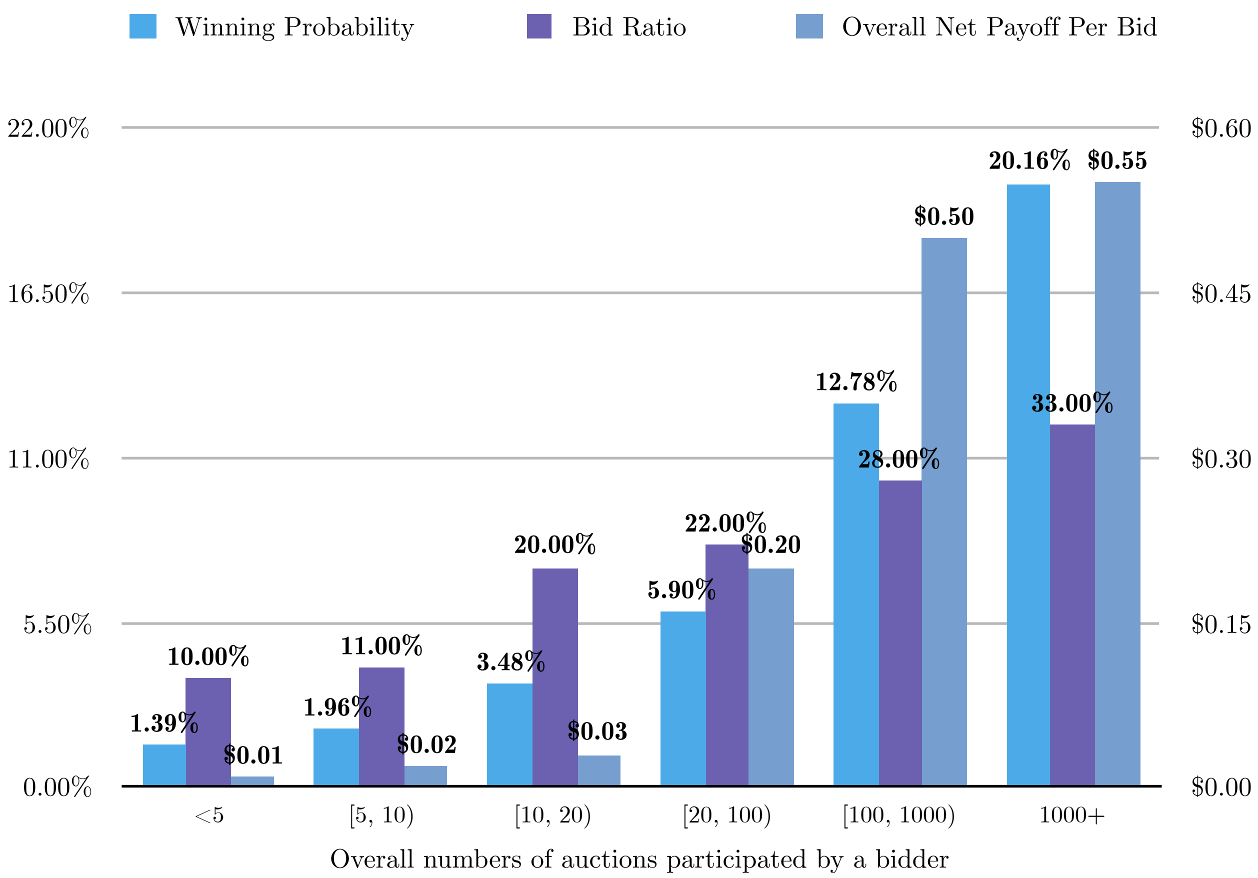

For individual bidders, there are similar findings. An *accumulated net payoff per bid* is the ratio of a bidder's *accumulated net payoff* since he/she starts playing penny auctions (computed as sum of RRP in all the auctions he/she wins, deducted by winning prices in all the auctions he/she wins and bidding fees for bids he/she placed in all auctions participated, including those he/she loses) and the total number of bids he/she submitted in all auctions. The *net payoff per bid* in a winning auction is the ratio of the winner's net payoff (computed as RRP deducted by winning price and bidding fees for bids placed by the winner in the auction) and number of bids placed by the winner. Figure 9 shows how a particular bidder's lifetime net payoff per bid and lifetime ratio of total bids in winning auctions over total bids in losing auctions increases, then declines slightly towards the end in his lifetime experience in terms of number of auctions participated in, although the net payoff per bid in any particular auction he wins appears to be stochastic.

Figure 9. Net payoff per bids in bidder's lifetime experience

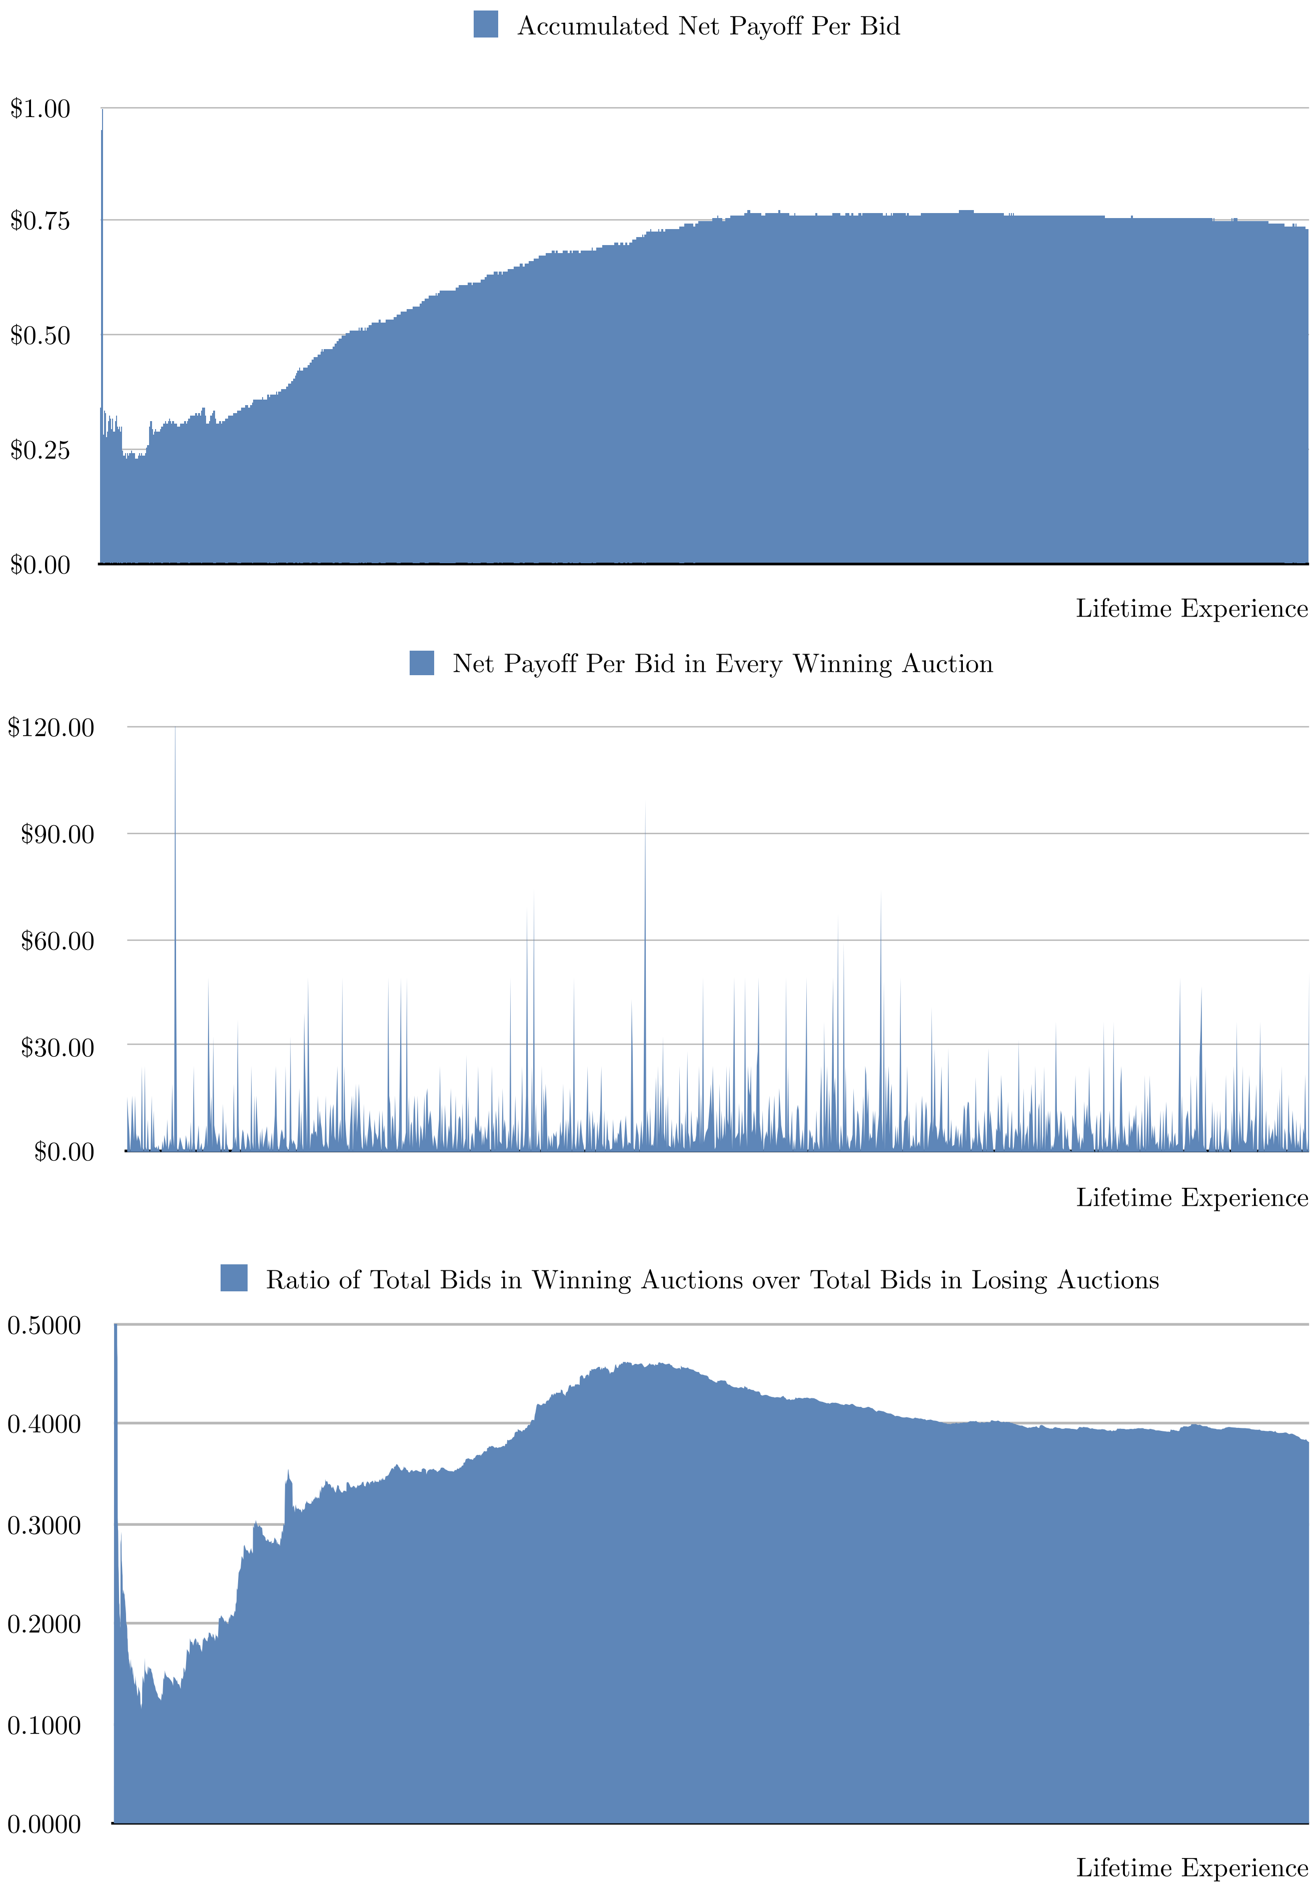

Figure 10 shows distributions of experienced bidder ratios (ratio of number of bidders with 20 or more prior auction experiences and numbers of all bidders in an auction[10]) in all five item categories[11], suggesting there might be a complex interaction between item category and bidder experience. For instance, the expensive consumer electronic category seems to draw different selections of bidders with more inexperienced bidder participation, compared to the other four categories.

Figure 10. Distribution of experienced bidder ratio in item categories

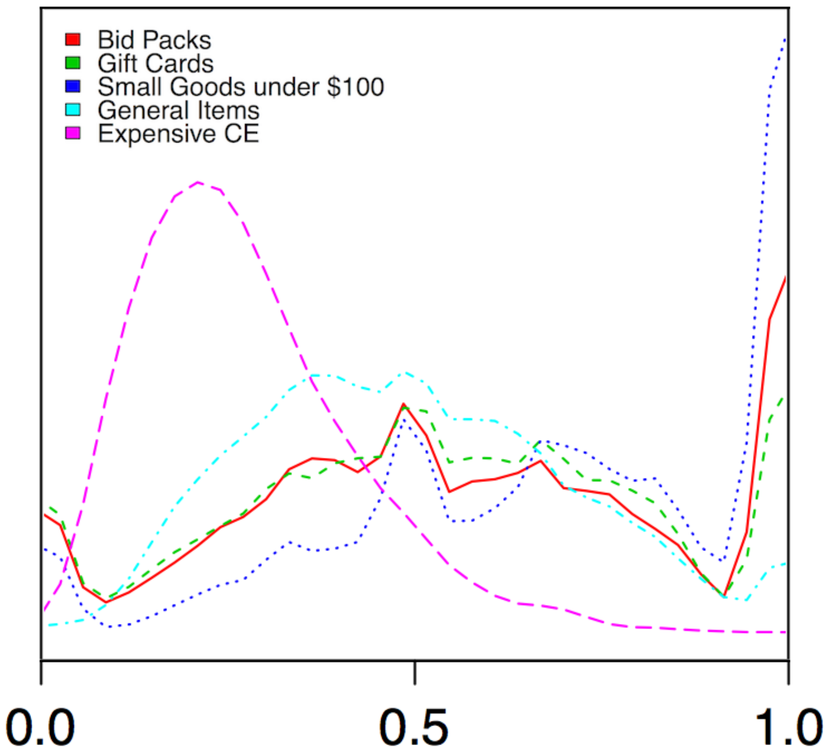

[10] I choose 20 prior auctions as an experience benchmark to separate the bidders. Majority of bidders (87%) play less than 20 auctions in their lifetime, while the rest who participated in 20 or more auctions place 81% of all bids. More details of bidders with different prior auction experiences are included in Appendix 1.

[11] The rest of the graph is trimmed and the full graph is included in Appendix 1.

Regression

Since profitability of a penny auction seems to be correlated with experience of participating bidders and categories of the auction item, we ran regressions to examine their correlation more closely:

$$\boldsymbol{y} = \boldsymbol{B}\beta_1 + \boldsymbol{D}\beta_2 + \varepsilon$$

where $\boldsymbol{y}$ is the seller's profit margin in an observed auction, $\boldsymbol{B}$ contains features of the bidders in the auction (ratio of bidders with experience of 20 prior auctions in all bidders in an auction, and number of bidders in the auction), $\boldsymbol{D}$ includes Dummy variables representing the item categories excluding the Bidpack, the base category, and $\boldsymbol{\varepsilon}$ is the error term.

In Regression (1), we regress only on the number of all bidders in an auction, which is not a factor affecting seller profitability in our theoretical model, though in regression (2), we regress on both the number of bidders and experienced bidder ratio; in regression (3), we regress only on the four dummy variables representing four categories besides bidpacks; in regression (4), we regress on the number of bidders and the experienced bidder ratio together with category dummy variables; in regressions (5) & (6), we include further interacting terms attempting to explain the cross effects between item category and other factors.

The regression results in Table 3 show that both the number of bidders and ratio of experienced bidders are positively correlated with the seller's profit margin. The result of regression (3) coincides with an average profit margin of all categories. Using the bid pack as the base category, we have coefficients of gift cards, small goods under $100, and general items over $100, being negative and only the coefficient of popular consumer electronics being positive, indicating that bid pack auctions generate more profit compared to auctions of all other categories except for the expensive consumer electronics. However, the results of regressions (4) (5) and (6) suggest that there may be more complex interactions between bidder experience and item categories. For instance, although the coefficients on the dummy variable of expensive CE and interacting term of experienced bidder ratio and expensive CE categories are negative, the coefficient on the interacting term of log (number of bidders) and dummy variable of expensive CE is positive, indicates the latter cross effect is large enough to overwhelm the previous two negative effects.

	Number of Bidders only (1)	Experienced Bidder Ratio (2)	Category only (3)	Experience and Category (4)	Cross Effects (5)	Cross Effects (6)
Ln(number of bidders)	1.354 (0.004)	0.992 (0.003)		1.077 (0.003)	1.026 (0.003)	1.112 (0.005)
Ratio of experienced bidders (REB)		0.077 (0.0002)		0.082 (0.0002)	0.110 (0.0003)	0.108 (0.0004)
Gift Card			-0.428 (0.006)	-0.685 (0.005)	-0.421 (0.006)	-0.309 (0.014)
Small Goods under $100			-1.107 (0.013)	-0.599 (0.010)	-0.296 (0.011)	0.636 (0.022)
General Goods			-0.804 (0.017)	-2.171 (0.013)	-1.177 (0.017)	-0.451 (0.046)
Expensive CE			0.556 (0.0437)	-2.880 (0.034)	-1.991 (0.053)	-4.597 (0.163)
REB*Gift Card					-0.033(0.0004)	-0.031(0.0004)
REB*Small Goods under $100					-0.047 (0.001)	-0.040 (0.001)
REB*General Goods					-0.075(0.0007)	-0.071(0.0007)
REB*Expensive CE					-0.050 (0.002)	-0.055 (0.002)
Ln(NOB)*Gift Card						-0.070 (0.007)
Ln(NOB)*Small Goods under $100						-0.612 (0.012)
Ln(NOB)*General Goods						-0.333 (0.019)
Ln(NOB)*Expensive CE						0.747 (0.047)
_cons	-2.061(0.007)	-2.070 (0.007)	0.755 (0.004)	-1.858 (0.007)	-1.983 (0.007)	-2.124 (0.009)
Observations	568,942	568,942	568,942	568,942	576,036	576,036
Adjusted R-squared	0.1915	0.3692	0.01921	0.4138	0.4274	0.4303
p<.001 for all above results Standard errors in parentheses						

Table 3. Seller's profit margin regression results

Significant results may be due to large sample size, so one should be careful here. The positive correlation between the number of bidders, experienced bidder ratios, and seller's profit margin stays significant throughout various regressions, and indicates that more experienced bidders participate in a penny auction, and the more profit the seller makes. This is similar to the game of chicken (if only one player in a penny auction is aggressive, he will often be successful; while if multiple players are aggressive, they may all lose in a costly bidding war).

Since data suggests complex interaction between bidder experience and item categories, we performed separate regressions of a seller's profit margin in auctions of each category. The regression results in Table 4 show that, regardless of category, the number of observed bidders[12] is positively and significantly correlated with a seller's profit margin, while the effects of experienced bidder ratios vary across categories.

	Bid Pack	Gift Cards	Small Goods Under $100	General Goods	Expensive CE
Regress only on Ln(NOB)					
Ln(NOB)	1.622 (0.006)	1.345 (0.005)	0.605 (0.007)	0.991 (0.015)	2.295 (0.074)
_cons	-2.242 (0.012)	-2.224 (0.011)	-1.293 (0.013)	-2.503 (0.038)	-6.806 (0.268)
Number of Observations	275,055	243,610	35,832	18,799	2,742
Adjusted R-squared	0.2131	0.2104	0.1592	0.1963	0.2588
Regress on Ln(NOB) and REB					
Ln(NOB)	1.612 (0.006)	1.375 (0.005)	0.606 (0.007)	0.966 (0.016)	2.457 (0.086)
REB	-0.073 (0.016)	0.310 (0.013)	0.101 (0.021)	-0.191 (0.053)	1.620 (0.432)
_cons	-2.182 (0.017)	-2.451 (0.014)	-1.367 (0.020)	-2.346 (0.058)	-7.823(0.381)
Number of Observations	275,055	243,610	35,832	18,799	2,742
Adjusted R-squared	0.2132	0.2123	0.1597	0.1968	0.2624
Regress on Ln(NOB), REB and Cross effect					
Ln(NOB)	2.023 (0.011)	1.392 (0.009)	0.610 (0.021)	1.676 (0.032)	3.617 (0.138)
REB	1.234 (0.033)	0.367 (0.030)	0.107 (0.040)	3.186 (0.140)	15.961 (1.414)
Ln(NOB) * REB	-0.849 (0.019)	-0.036 (0.017)	-0.004 (0.026)	-1.413 (0.054)	-4.302 (0.405)
_cons	-2.852 (0.023)	-2.478 (0.019)	-1.371 (0.033)	-4.138 (0.089)	-11.918 (0.536)
Number of Observations	275,055	243,610	35,832	18,799	2,742
Adjusted R-squared	0.2188	0.2123	0.1597	0.2248	0.2913
p<.001 for all above results Standard errors in parentheses					

Table 4. Seller's profit margin regressions on observations in each category

[12] Only bidders who submitted at least one bid successfully were observed, which means this number of observed bidders may be equal or less than the real number of all bidders in an auction.

Chapter 3. Basic Model

We start our theoretical analysis of penny auctions from a basic symmetric complete information model, which is a mixture and extension of theoretical models presented in Platt et al. (2013) and Augenblick (2012). Assume the auctioneer is selling a single item to $\boldsymbol{N}$ identical bidders, indexed by $\boldsymbol{i} \in \{1, 2, ..., \mathrm{N}\}$, and all players have common value $\boldsymbol{v}$ for the item. The auction starts at initial price of zero ($\boldsymbol{P}_0 = 0$), and price rises by a bidding increment $\boldsymbol{s} \in \mathbb{R}^+$ in each period, indexed by $\boldsymbol{t} \in \{0, 1, 2, 3, ...\}$. Thus $\boldsymbol{P}_t$, the current price at any period $\boldsymbol{t}$ equals $\boldsymbol{s} \cdot \boldsymbol{t}$, and the value of winning an auction at period $\boldsymbol{t}$ is $\boldsymbol{v} - \boldsymbol{st}$.

At each period $\boldsymbol{t} > 0$, there is only one current leader $\boldsymbol{l}_t \in \{1, 2, 3, ..., \boldsymbol{N}\}$ and $\boldsymbol{N} - 1$ non-leaders; while at the initial period $\boldsymbol{t} = 0$, all $\boldsymbol{N}$ bidders are non-leaders ($\boldsymbol{l}_0 = 0$). Note that the same as in practice, the current leader in any period is not allowed to place a bid in this model. Within each period $\boldsymbol{t}$, all non-leaders simultaneously choose whether to submit a bid or not, and any multiple simultaneous bids would be handled immediately before the game reached the next period $\boldsymbol{t} + 1$. If only one bidder submits a bid, then his bid is accepted and he pays a non-refundable bidding fee $\boldsymbol{c}$; if $\boldsymbol{K} \geq 2$ bidders places their bid in the same period, then only one bid is randomly accepted with probability $1/\boldsymbol{K}$ and the corresponding bidder pays the bidding fee $\boldsymbol{c}$, while all other bidders whose bids are not accepted do not need to pay the bidding fee. If all non-leaders choose not to submit a bid, then the auction ends and the item is sold to the current leader $\boldsymbol{l}_t$ at current price level $\boldsymbol{P}_t = \boldsymbol{s} \cdot \boldsymbol{t}$, while the auctioneer collects a total revenue of $(\boldsymbol{ct} + \boldsymbol{st})$. If the auction ends at $\boldsymbol{t} = 0$, then the seller keeps the item.

As a Complete Information model, we assume that all parameters of the auction are commonly known, and all previous bids and the identity of the current leader in each period are observed by all players who do not discount future consumption. Assume each bidder has the same initial wealth of $\boldsymbol{W}$ and identical strictly increasing Von Neumann–Morgenstern utility function of $\boldsymbol{u}(\bullet)$, a participants placing no bid maintains his utility $\boldsymbol{u}(\boldsymbol{W})$, a non-winning bidder who places $\boldsymbol{k}$ bids gets $\boldsymbol{u}(\boldsymbol{W} - \boldsymbol{kc})$, and a winning bidder who wins at period $\tilde{\boldsymbol{t}}$ and places $\tilde{\boldsymbol{k}}$ bids in total gets $\boldsymbol{u}(\boldsymbol{W} + \boldsymbol{v} - \boldsymbol{s}\tilde{\boldsymbol{t}} - \tilde{\boldsymbol{k}}\boldsymbol{c})$.

We also assume $c < v - s$, otherwise no one could bid in period 0; and $\mathbf{mod}(v-c,s)=0$ [13] for simplicity.

Since bidding is a binary activity, let bidder i's strategy set consist of $\beta^i_{H_t} \in [0,1]$, where $\beta^i_{H_t}$ is the probability that non-leader i attempts to place a bid at period t with history of $H_t = \{(0,l_0),(1,l_1),...,(t-1,l_{t-1})\}$ and $H_1 = \{(0,l_0 = auctioneer)\}$. Let $\mu_{H_{t+1}}$ be the probability of $(t+1)^{th}$ bid being placed by any non-leader at period t as leader at any period is not allowed to place a bid, given the t^{th} bid has occurred, i.e. the probability that the auction survives at period $t+1$, given period t is reached. Thus we have $\mu_{H_{t+1}} \equiv 1-(1-\beta_{H_t})^{N-1}$ $\forall t > 0$, and $\mu_{H_1} \equiv 1-(1-\beta_0)^N$.

We define the (discrete) hazard function $h(t,H_t)$ as the probability that all non-leaders choose not to bid at a period, given it is reached, and mapping each period to the probability that the game ends at that state, given the state is reached. Thus, we have

$$h(t,H_t) = \mathbf{Pr}\left[\beta^i_{H_t} = 0 \ \forall i \neq l_t \mid auction\ reaches\ period\ t\right] = \frac{\prod_{i=1}^{N}\left(1-\beta^i_{H_t}\right)}{1-\beta^{l_t}_{H_t}} = \left(1-\beta_{H_t}\right)^{N-1} \ \forall t > 0 \ ,$$

and the hazard rate of the initial period represents the probability of the auction ending with no bidder, $h(0) = \prod_{i=1}^{N}\left(1-\beta^i_0\right) = \left(1-\beta_0\right)^N$.

That is, we have $h(t,H_t) = 1-\mu_{H_{t+1}}$ $\forall t > 0$, and

$$\begin{aligned} h(t,H_t) &= \mathbf{Pr}\left[\beta^i_{H_t} = 0 \ \forall i \neq l_t \mid auction\ reaches\ period\ t\right] \\ &= \frac{\prod_{i=1}^{N}\left(1-\beta^i_{H_t}\right)}{1-\beta^{l_t}_{H_t}} = \left(1-\beta_{H_t}\right)^{N-1} \quad \forall t > 0 \end{aligned}$$

Assume each player's next move is predicted by the last move of the other players, NOT by the earlier history of moves, we have $h(t) = \left(1-\beta_t\right)^{N-1}$ $\forall t > 0$

---

[13] That is, the reminder of devision of v-c by s is zero, which means v-c is an integer multiple of s.

Define the hazard rate of an auction ending without a bidder $h(0) = \mu_{H_1}$, and

$T \equiv \frac{v-c}{s} - 1$ is an integer by previous assumption of $\mathbf{mod}(v-c, s) = 0$ [14], we claim that:

Theorem 1

In any Subgame Perfect Equilibrium where the auction survives beyond period 1 with some positive probability, we have

$$h(t) = \begin{cases} 0 & t = 0 \\ \dfrac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)} & \textit{for } 0 < t \le T \\ 1 & \textit{for } t > T \end{cases}$$

Since we each bidder has identical strictly increasing vNM utility function of $u(\bullet)$ and $T \equiv \frac{v-c}{s} - 1$, we have $h(t) = 1 \ \forall t > T$. No bidder is willing to place any bid at period $T+1$ and beyond, as if she places a bid in period $T+1$ and becomes the new leader at period $T+2$, the best possible outcome for her is that the auction ends at $T+2$ and she wins the item, which still gives her a payoff less than her current wealth, as $W_T - c + \left[v - s(T+2)\right] = W_T - s < W_T$. Any non-leader at period T is different between placing a bid and not - both giving her a payoff of zero. If she places a bid and her bid is accepted, she becomes leader at period $T+1$ and wins the auction, as no one else is willing to overbid her, thus her payoff is $W_T - c + \left[v - s(T+1)\right] = W_T$. She also gets W_T if another non-leader enters period $T+1$ and wins the auction.

Now start our backward induction from period T. Consider any non-leader at period $T-1$, she must be indifferent between placing a bid and not placing a bid in any Subgame Perfect Equilibrium. If she places a bid and her bid is accepted, she becomes the current leader at period T, a period in which all non-leaders getting zero payoffs,

[14] Augenblick 2012 shows that the case with T not being an integer can be remedied in an ε-perfect equilibrium, which approximates the equilibrium we propose below.

regardless whether they bid or not. Then she gets $W_{T-1}+v-sT-c$ if the auction ends at period T, or $W_{T-1}-c$ if she enters period $T+1$ as a non-leader by induction. She gets a W_{T-1} payoff if she enters period T as a non-leader, if her bid is not accepted or she chooses not to place a bid. Thus, we have the indifference condition

$$h(T)u\left[W_{T-1}+v-sT-c\right]+\left(1-h(T)\right)u\left[W_{T-1}-c\right]=u\left[W_{T-1}\right]$$

for any non-leader at period $T-1$ in any Subgame Perfect Equilibrium, which derives

$$h(t)=\frac{u(W_{T-1})-u(W_{T-1}-c)}{u(W_{T-1}+v-st-c)-u(W_{T-1}-c)} \quad \text{when } t=T$$

By backward induction, we can find the equilibrium indifference conditions for any non-leader at periods $t=T-2,\ T-3,\ \ldots,\ 0$,

$$h(t+1)u\left[W_t+v-s(t+1)-c\right]+\left(1-h(t+1)\right)u\left[W_t-c\right]=u\left[W_t\right]$$

deriving a hazard rate of periods $1\le t<T$,

$$h(t)=\frac{u(W_{t-1})-u(W_{t-1}-c)}{u(W_{t-1}+v-st-c)-u(W_{t-1}-c)}$$

Thus when $0<t\le T$,

$$h(t)=\frac{u(W_{t-1})-u(W_{t-1}-c)}{u(W_{t-1}+v-st-c)-u(W_{t-1}-c)}$$

Since we are looking at Subgame Perfect Equilibrium where the auction survives the first two periods with positive probability, we choose $h(0)=0$ arbitrarily, i.e. some bidding always occurs at period 0 in equilibrium, so that the auction reaches period 1 with certainty. For instance, in case that $0<h(0)<1$, where there is some positive probability that the auction ends with no bidder and does not reach period 1, the auctioneer can repeatedly run the auction, until some players bid, leading the hazard rate of the initial period to effectively be zero.

Other symmetric equilibria also exist, such as all players always bidding in even periods and always passing in odd periods, or vice versa, which all produce the same result: the auction either ends in period 0 with no bidder, or ends in period 1 with one winning bidder, which are concluded as: in any other symmetric subgame perfect equilibrium, the auction will conclude either in period 0 with no bidders, or in period 1 with one bidder.

Suppose that there exists a period $\boldsymbol{t}$ where bidders randomise with

$$\beta_t > 1 - \left[\frac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)} \right]^{\frac{1}{N-1}}$$

and follow the strategy profile $\beta_t = 1 - \left[h(t) \right]^{\frac{1}{N-1}}$ for all periods beyond $\boldsymbol{t}$, which is permissible as bidders are indifferent between bidding at $\boldsymbol{t}$ and not bidding. In this case, non-leaders in period $t-1$ strictly prefer not bidding, as an expected payoff from bidding is negative, strictly less than a zero payoff from not bidding:

$$(v - sq)(1 - \mu_{q+1}) - b = (v - sq)(1 - \beta_q)^{n-1} - b < 0$$

Thus only $\beta_{t-1} = 0$ can occur, which in turn requires $\beta_{t-2} = 1$, i.e. everyone strictly prefers an attempt to bid in period $t-2$ as no one will bid in $t-1$. We continue the backward induction to show the alternative choices of bidders over earlier periods, such that if $\boldsymbol{t}$ is even, the equilibrium outcome concludes with a single bid (since $\beta_1 = 0$), and if $\boldsymbol{t}$ is odd, the outcome concludes with no bid occurring (since $\beta_0 = 0$). Alternatively, if randomising

$$\beta_t < 1 - \left[\frac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)} \right]^{\frac{1}{N-1}}$$

a reverse outcome results for even or odd $\boldsymbol{t}$. Therefore, if indifference is broken at any period, given that bidders follow the strategy profile $\beta_t = 1 - \left[h(t) \right]^{\frac{1}{N-1}}$ for all periods beyond it, all earlier periods would follow the alternating strategy resulting in the auction concluding with zero or one bidder.

Theorem 1 is proved by backward induction that, at any $\boldsymbol{t}$, non-leading bidders are indifferent from submitting a bid or not, and expect zero gain in expected utility from the continuation of the game, which means that they expected a zero gain in expected utility from the beginning of the auction. Under risk neutrality, this means the seller must expect zero profit, since the total surplus to be shared is zero as the seller and buyers place the same value on the object. Whereas, if bidders are risk-averse or risk-seeking with concave or convex utility functions, the hazard rate will be higher or lower than the hazard rate with risk-neutral bidders, and the expected profit of a seller will be negative or positive respectively. Theorem 1 also shows that history does not affect hazard rates in equilibrium (though individual bidder strategies may still be history-dependent)

We can also prove the zero profitability property under risk neutrality, by constructing the probability density function that the auction ends at period t (i.e. in total, t bids are accepted) $f(t)$. Recall that μ_{t+1} is the probability of $(t+1)^{th}$ bid being placed by any non-leader at period t, given the t^{th} bid has occurred, i.e. the probability that the auction survives at period $t+1$, given period t is reached, which equals to $1-h(t)$. Thus, we have

$$f(t) \equiv (1-\mu_{t+1})\prod_{m=1}^{t}\mu_m = h(t)\prod_{m=1}^{t}\left[1-h(m-1)\right]$$

if bidders have linear utility fuction,

$$f(t)=\begin{cases} h(0) & \textit{if } \ t=0 \\ \dfrac{c}{v-st}\left[1-h(0)\right]\prod_{m=2}^{t}\left(1-\dfrac{c}{v-s(m-1)}\right) & \textit{if } 0<t\le T \end{cases}$$

i.e. probability of the auction ending at period t equals to the product of probability of $(t+1)^{th}$ bid not occurring and probability of all past t bids being placed, i.e. the product of the hazard rate of period t and probability that the auction survives through all periods prior to period t. We can compute the expected revenue of the auctioneer as $E[Rev]=\sum_{t=1}^{T+1}(c+s)t\cdot f(t)$, which sums to v, and is independent of auction parameters such as a bidding fee and bidding increment. The sum is complicated to prove, while we can obtain the same result via the following indirect method. Assuming that the seller places no intrinsic value on the item, the total expected surplus of the auction is equal to v times the probability that the auction item is sold, i.e. probability of auction does not end before period one, i.e. one minus hazard rate of period one. This expected surplus is split between the seller and bidders, given the expected surplus of the bidders is zero, the seller's expected revenue is $v\left(1-h_1\right)$. We can reasonably set $h_1 = 0$, as the item has not been sold. The seller is able to immediately initiate a new auction at practically no cost, repeating it until the item is sold. Thus the seller's expected revenue is v, and gets zero expected profit in the complete-information subgame perfect equilibrium under risk neutrality.

However, the variance of revenue does depend on these parameters. Adopting the approach of Platt et al 2012, when c is small relative to v, by treating t as a continuous variable, $f(t)$ can be approximated by a *Generalized Beta Distribution of the First Kind* (setting GB1 parameters as $a = 1$, $\gamma = 1$, $\delta = \frac{c}{s}$, $\rho = \frac{v-c}{s}$), and

$$f'(t) \approx f(t+1) - f(t) = \frac{c-s}{v-s-st} f(t)$$

derives a unique solution $f(t) = c(v-s)^{-\frac{c}{s}}(v-s-st)^{\frac{c}{s}-1}$ when s is very small relative to v, where the constant of integration is determined such that $\int_0^{\frac{v-s}{s}} f(t)dt = 1$. Using properties of GB1, we can compute variance of the expected revenue $Var[Revenue] = \frac{c}{c+2s}(v-s)^2$, which increases in c and v, and decreases in s.

Figures 11 and 12 show the probability distribution of the final bid and the seller's expected revenue within a period, $(c+s)t \cdot f(t)$ in a numeric example of v=\$10, c=\$1, s=\$0.2. Note that the area under the per-period expected revenue curve equals to the total expected revenue that the seller can gain throughout the auction.

Figure 11. Probability distribution of final bid

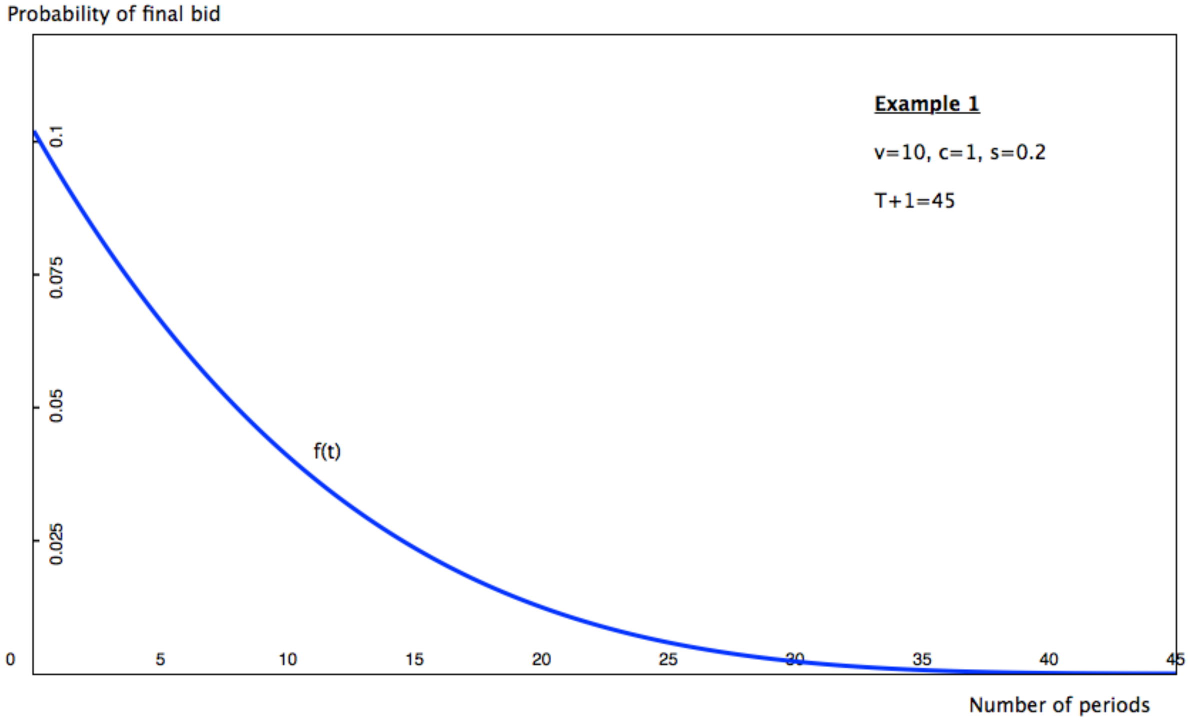

Figure 12. Seller's expected revenue per period

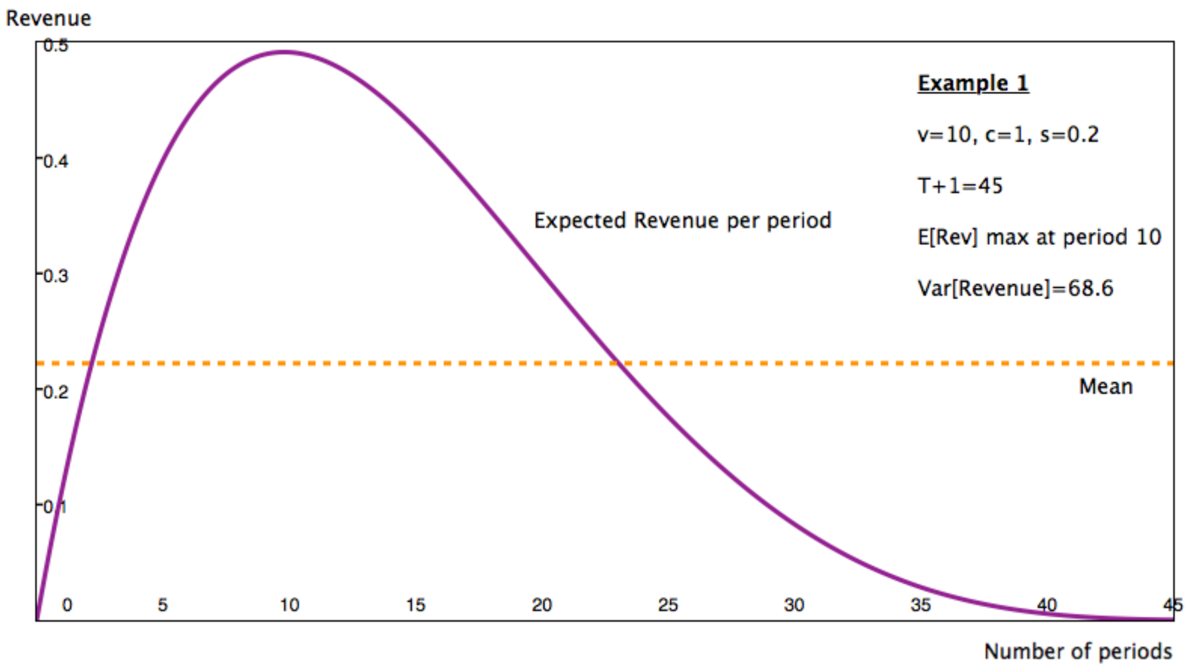

Theorem 2

In any Markov Perfect Equilibrium where the auction survives beyond period 1 with some positive probability, and $\beta_0^i > 0$, $\beta_1^i > 0 \ \forall i$, *we have symmetric bidding strategy that* $\beta_t^i = \beta_t^j \ \forall i \neq j$, *where*

$$\beta_t^i = \begin{cases} 1 - \sqrt[N-1]{\dfrac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)}} & \textbf{\textit{for}}\ 1 \le t \le T \\ \\ 0 & \textbf{\textit{for}}\ t > T \end{cases} \qquad \forall i$$

In Theorem 1, we discuss the property of a hazard rate of our interested set of Subgame Perfect Equilibrium, and we discuss property of bidders' strategy profiles in a more interested subset of the this Subgame Perfect Equilibrium set in Theorem 2, the symmetric Markov Perfect Equilibria, in which all bidders play symmetric Markov strategies. Markov Perfect Equilibrium is refinement of Subgame Perfect Equilibrium, and consists of a set of mixed strategies for each of the players where each player's mixed strategy can be conditioned on the state of the game, which only encode payoff relevant information and exclude strategies that depend on the identity of the players. For instance, bidders only condition their strategies on time points (period) of the auction, not the entire history, and change in identity of any current leader at any period will not change the payoff of the strategies, so we can simplify our description of the strategy profiles by excluding the identity of the current leader of any particular period, e.g. write as β_t^i instead of β_{t,l_t}^i. We also focus only on symmetric Markov Perfect Equilibria here, in which the players have strategy and action sets which are mirror images of one another, and payoffs for playing a particular strategy depend only on the other strategies employed, not on who is playing them, i.e. $\beta_t^i = \beta_t^j \ \forall i \neq j$.

We argue that in any Markov Perfect Equilibrium where the auction survives beyond period 1 with positive probability and all bidders bid with positive probability in the initial periods (period 0 and 1), bidders play the symmetric strategies stated in Theorem 2. The additional constraint that every player bids with some positive probability at $t = 0$ and $t = 1$ helps us to exclude the unwanted asymmetric equilibria, e.g. the equilibria in which one player effectively leaves the auction after period 1, and the equilibria in which some player is always the leader in a specific period. First consider the idea of symmetric strategies in equilibrium. Suppose there exists $\beta_t^i \neq \beta_t^j$ at some period of $t > 1$, then $h(t, l_t = i) \neq h(t, l_t = j)$, i.e. the two players face different hazard rates if each of the two is the leader at period t, which leads to at least one of the hazard rate unequal to $\dfrac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)}$, violating Theorem 1; and it would cause players to bid in a way that keeps the history off the equilibrium path, e.g. players would behave differently in $t - 1$ or $t - 2$ from what they would do in the equilibrium.

Recall that $h(t) = 1 \ \forall t > T$ by Theorem 1, i.e. the auction ends with certainty and no player is willing to place a bid at any period beyond T, thus we have $\beta_t^i = 0 \ \forall \ i \ and \ \forall t > T$. Similarly as $h(0) = 0$, i.e. the probability that the auction ends at period 0 is zero, by symmetry, thus we have $\beta_t^i = 1 \ \forall \ i$ when $t = 0$.

When $0 < t \leq T$, $\beta_0^i > 0$ and $\beta_1^i > 0$ for all i by induction from period 1, we can prove

$$\beta_t^i = 1 - \sqrt[N-1]{\frac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)}} \quad \forall i$$

Recall $h(t) = \dfrac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st - c) - u(W_{t-1} - c)}$ when $0 < t \leq T$ by Theorem 1. Since $\beta_0^i > 0$ and $\beta_1^i > 0 \ \forall i$, every player enters period 1 and period 2 as leader with some positive probability in equilibrium, thus we have

$$h(t = 1, l_1 = i) = h(t = 1, l_1 = j) = \frac{u(W_0) - u(W_0 - c)}{u(W_0 + v - s - c) - u(W_0 - c)}$$

and

$$h(t=2,l_2=i)=h(t=2,l_2=j)=\frac{u(W_1)-u(W_1-c)}{u(W_1+v-2s-c)-u(W_1-c)}$$

Suppose there exists some players i,j that $\beta_1^i \neq \beta_1^j$, by definition of the hazard rate, we have $h(t=1,l_1=i)\neq h(t=1,l_1=j)$ as $\frac{\prod_{m=1}^{N}\left(1-\beta_1^m\right)}{1-\beta_1^i}\neq\frac{\prod_{m=1}^{N}\left(1-\beta_1^m\right)}{1-\beta_1^j}$, contradicting to our previous finding that $h(t=1,l_1=i)=h(t=1,l_1=j)$. Thus, there does not exist any two player i,j that $\beta_1^i \neq \beta_1^j$ in equilibrium, and we must have $\beta_1^i=\beta_1^j\ \forall i\neq j$, such that $h(t=1)=\left(1-\beta_{t=1}^i\right)^{N-1}$, which generate a unique solution of

$$\beta_1^i=1-\sqrt[N-1]{\frac{u(W_0)-u(W_0-c)}{u(W_0+v-s-c)-u(W_0-c)}}\quad\forall i$$

Similarly, we can show that

$$\beta_2^i=1-\sqrt[N-1]{\frac{u(W_1)-u(W_1-c)}{u(W_1+v-2s-c)-u(W_1-c)}}\quad>0\ \forall i$$

Since $\beta_1^i>0\ \forall i$ and $\beta_2^i>0\ \forall i$, every bidder enters period 3 as the current leader with some positive probability in equilibrium, and similarly we can prove that

$$\beta_3^i=\beta_3^j=1-\sqrt[N-1]{\frac{u(W_2)-u(W_2-c)}{u(W_2+v-3s-c)-u(W_2-c)}}\quad\forall i\neq j$$

By extending to the following periods, we can show that

$$\beta_t^i=\beta_t^j=1-\sqrt[N-1]{\frac{u(W_{t-1})-u(W_{t-1}-c)}{u(W_{t-1}+v-st-c)-u(W_{t-1}-c)}}\quad\forall i\neq j\ \text{when}\ t=4,\ t=5,\ldots,t=F$$

We can also prove that there is no strictly profitable deviation from this strategy profile if the auction survives beyond period 1. Consider non-leader i 's expected continuation payoff at t from bidding and not bidding in the subgame starting at $1\le t\le T$. By induction and Theorem 1, we can see that she gets zero benefit from not bidding, and also zero benefit from placing a bid, as

$$u\left(W_t - c\right) + h(t+1,i)u\left(W_t + v - s(t+1)\right) + \left[1 - h(t+1,i)\right] \bullet u\left(W_t\right) = u\left(W_t\right)$$

Since the continuation payoffs from bidding and not bidding are both zero, there is no strictly profitable deviation from the stated mixed strategies. Therefore, we have proved Theorem 2.

Theorem 3.

In any symmetric Markov Perfect Equilibrium, we have either all bidders playing the symmetric strategy profile described in Theorem 2, or alternating strategies such that the auction ends in period 0 with no bid, or ends in period 1 with one bidder.

We are aware that there are other Markov equilibria, in which an auction always ends at period 0 or period 1. In these equilibria, players believe that some players will bid with very high probability in period 1 or 2, which leads them to strictly prefer not to bid.

Suppose that there exists a period $t^* \le T$ where bidders randomise with some

$$\beta_{t^*} > 1 - \sqrt[N-1]{\frac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st^* - c) - u(W_{t-1} - c)}} \text{ and } \beta_t = 0 \ \ \forall t > T$$

which is permissible as bidders are indifferent between bidding at t^* and not bidding. In this case, non-leaders in period $t^* - 1$ strictly prefer not bidding, as the expected payoff from bidding is strictly less than the payoff from not bidding, as

$$u(W_{t^*-1} + v - st^* - c)h(t^*) + \left[1 - h(t^*)\right]u(W_{t^*-1} - c) < u(W_{t^*-1})$$

Thus only $\beta_{t^*-1} = 0$ can occur, which in turn requires $\beta_{t^*-2} = 1$, i.e. everyone strictly prefers an attempt to bid in period $t^* - 2$ as no one will bid in $t^* - 1$. We can continue the backward induction to show the alternative choices of the bidders over earlier periods, such that if t^* is even, the equilibrium outcome concludes with a single bid (since $\beta_1 = 0$), and if t^* is odd, the outcome concludes with no bid occurring (since $\beta_0 = 0$). Alternatively, if we randomise

$$\beta_{t^*} < 1 - \sqrt[N-1]{\frac{u(W_{t-1}) - u(W_{t-1} - c)}{u(W_{t-1} + v - st^* - c) - u(W_{t-1} - c)}}$$

we get a reverse outcome for even or odd t^*. Therefore, if the indifference condition is broken in as any period, given that bidders follow strategy profile in Theorem 2 for all periods beyond it, all bidders would follow the alternating strategy in all earlier periods resulting in outcomes that the auction either ends in period 0 with no bid, or ends in period 1 with one bidder, and there is no profitable deviation intuitively.

Chapter 4. Risk Model - Single Bidder Type

We have shown that the auctioneer gets zero expected profit in the complete-information subgame perfect equilibrium under risk neutrality, and we can show that the auctioneer gets positive/negative expected profit when bidders have convex/concave utility functions.

For simplicity, we assume that all bidders have an identical Constant Absolute Risk (CARA) utility function of $u(W)=\dfrac{1-e^{-\alpha W}}{\alpha}$, where α is the Absolute Risk Aversion, and every bidder has the same initial wealth W, such that initial wealth terms are cancelled out in the calculation.

By Theorem 1, we find the hazard rate $h(t)=\dfrac{e^{\alpha c}-1}{e^{\alpha c}-e^{\alpha(c+st-v)}}$ $0<t\leq T$ at period $0<t\leq T$ in any Subgame Perfect Equilibrium under an additional constraint that every bidder bids with positive probability at initial periods, so that we can exclude those uninteresting asymmetric equilibria.

Similarly we can also calculate the probability density as:

$$f(t)= h(t)\prod_{m=1}^{t}\left[1-h(m-1)\right]=\frac{e^{\alpha c}-1}{e^{\alpha c}-e^{\alpha(c+st-v)}}\prod_{m=1}^{t}\left(\frac{1-e^{\alpha(c+s(m-1)-v)}}{e^{\alpha b}-e^{\alpha(c+s(m-1)-v)}}\right)$$

We can see that the support of $f(t)$ is the same as before, $f(t)>0$ for all t when $0<t\leq T$, and the increase in v has essentially the same effect as in the risk-neutral case (increasing the support and flattening the distribution). When $\alpha>0$ (i.e. risk aversion), $f(t)$ has a similar convex shape to the risk-neutral density function, only with greater curvature as α rises. However, the distribution behaves quite differently for $\alpha<0$ (i.e. risk-loving). When α is very close to zero, an inflection point $\hat{t}$ is introduced near zero such that $f''(t)<0$ is below $\hat{t}$; thus $f(t)$ is no longer strictly convex. As α decreases, this inflection point takes on higher values, and eventually creates a hump-shaped distribution, where $f(t)$ is maximised at

$$t = \max\left\{0, \frac{v-c}{s} + \frac{1}{\alpha s}\ln\left(\frac{1-e^{\alpha b}}{1-e^{\alpha s}}\right)\right\}$$

, which increases as α becomes more negative or v becomes larger.

While we cannot analytically solve for expected revenue in this case, we can derive the effect of the parameters on the expected revenue by evaluating the comparative statics of each parameter on the hazard rate $h(t)$. By direct calculation, there are: $\frac{\partial h(t)}{\partial \alpha} > 0$, $\frac{\partial h(t)}{\partial s} > 0$, $\frac{\partial h(t)}{\partial c} > 0$, i.e. the hazard rate increases in α, s, c.

Since $h(t) = 1 - \mu_{t+1}\ \forall t > 0$, we can see that the average final number of bid decreases in α, s, c, which implies that expected revenue decreases in α. The finding is quite intuitive, a decrease in α means bidders are more risk-loving and gaining more utility from participating in the risky game, i.e. the auctioneer is able to extract more revenue from bidders' risk preference.

The empirical findings confirmed the negative risk parameter α primarily in the range of -0.001 to -0.04 (with a few estimates as low as -0.09), indicating that bidders are mildly risk loving. [15]

The seller's expected revenue in a penny auction is the sum of his expected revenue within each period, which depends on the probability of the auction ending at that period. Figures 13 and 14 show how the probability of the final bid and seller's expected revenue within each period, $(c+s)t \cdot f(t)$, varies when bidders have different risk attitudes in numeric examples. Note that since the seller's expected revenue within each period sums up to the seller's expected revenue of the auction, the area under the per-period expected revenue curve equals to the total expected revenue the seller can gain throughout the auction. The dashed line of Mean (item valuation v divided by $\mathbf{T}$) is included as a guideline, as seller breaks even when the area under the per-period expected revenue curve equals to be area under the dashed line of mean.

[15] Horserace participant's risk attitude as benchmark

In Example 1, we set v=\$10, c=\$1, s=\$0.2, and bidders are risk-neutral ($\alpha = 0$). In Example 1A and 1B, we will use the same auction parameters, while setting bidders to be risk-loving ($\alpha = -0.05$) and risk-averse ($\alpha = 0.05$).

Figure 13. Probability distribution of final bids in numerical examples

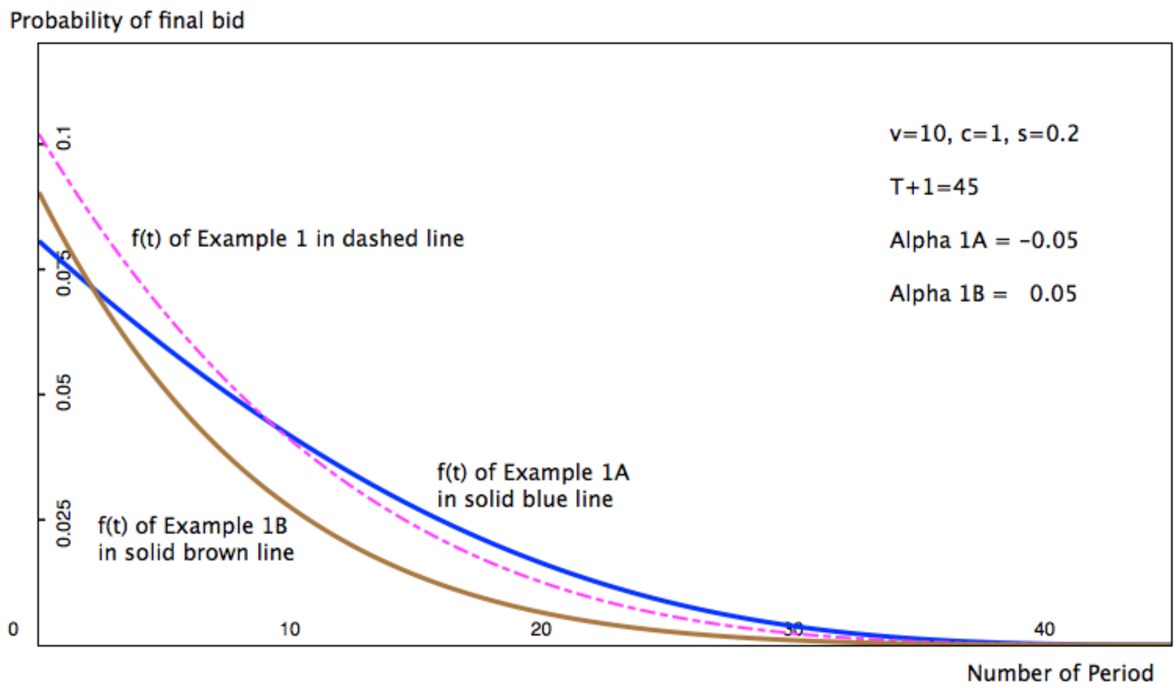

Figure 14. Seller's expected revenue within each period in numerical examples

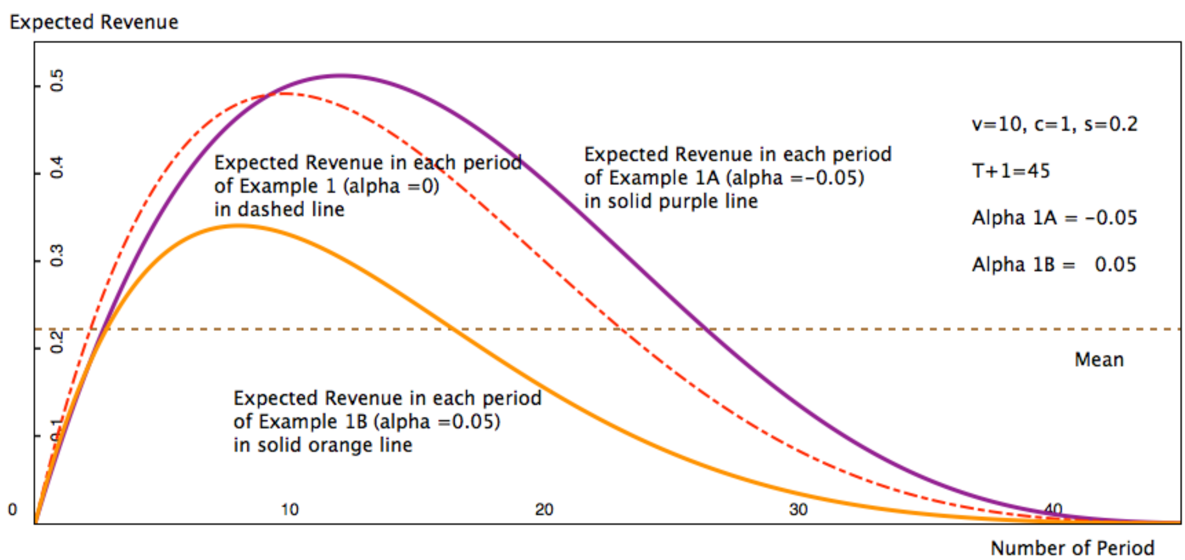

From the graphs we can see that the positive alpha indicating risk-aversion in Example 1B, drags down the probability density distribution of the final bid and expected revenue within each period, which shrinks the area under the curve and reduces the seller's expected revenue to below the item valuation. While the negative alpha indicating risk-loving in Example 1A increases the seller's expected revenue, which now exceeds the item's valuation.

We conclude that the average final number of bids decrease and the expected revenue of the auctioneer decreases in α , so that the seller earns more from more risk-loving bidders, who gain more utility from participating in the risky game, and earns less from more risk-averse bidders. It is intuitive to understand that as higher risk tolerance of bidders increases their bidding aggression, it extends the length of the auction and improves the profit of the seller.

We use our equilibrium bidding model to estimate risk attitudes of bidders using our original dataset of auctions collected on BidCactus. The empirical strategy of Platt et al. (2013) will be followed (with one change to be noted below), and in addition, we will estimate separate coefficients of absolute risk aversion for experienced and inexperienced bidders, to test whether there is a statistically significant difference.

Our key theoretical prediction is that the final number of bids in a given auction is a random variable with distribution $f(t,\alpha)$ [16]. If a given item is repeatedly auctioned many times, we will have enough observations to estimate risk parameters, and to see whether our risk model does a better job at explaining the observed outcomes compared to the basic model; whether there is a significant difference between risk attitudes depending on experience. By performing statistical tests (Pearsons χ^2 test and K-S test to compare distributions, t-test to compare means), we can quantify how closely the estimated theoretical distribution matches the observed sample distribution. Maximum likelihood estimates of α for each regularly-auctioned item are generated.

[16] Although the theory predicts that this distribution should be independent of **N** (the number of bidders), our earlier regressions suggest otherwise, and this is a matter for future investigation.

Choose the parameter α to maximise $\sum_j \ln f(t_j;\alpha)$, where j represents each observed auction of that item, t_j is the ending number of bids in that auction, and f is the discrete theoretical distribution in the CARA model

$$f(t,\alpha) = \frac{e^{\alpha c} - 1}{e^{\alpha c} - e^{\alpha(c+s*t-v)}} \prod_{m=1}^{t} \left(\frac{1 - e^{\alpha(c+s(m-1)-v)}}{e^{\alpha c} - e^{\alpha(c+s(m-1)-v)}} \right)$$

The following examples illustrate the actual and fitted distributions for items in two different categories. We first compute a ML estimate of alpha for each item, then we compare the theoretical distributions of final bids under risk neutrality (alpha set to be zero) and those with risk parameters adopting the ML estimate of alpha, against the observed sample distributions. Table 5 shows that implementation of the MLE in the model explains the observation better when compared to setting alpha to zero (assuming risk neutrality). The tests show the improvement in fit is statistically significant, and the graphs give an intuitive impression of the difference in fit.[17]

Calphalon Nonstick Stir Fry Pan (RRP$49.95, Item#697, Category: Small Good under $100)	
Number of Observations	197
Estimated Alpha by MLE	0.03178
Average Profit Margin	- 50.23%
t-test (p>=0.05)	Basic: 8.1
	MLE: 91.1
Chi-Sq test (p>=0.05)	Basic: 9.3
	MLE: 73.4
K-S test (p>=0.05)	Basic: 7.0
	MLE: 82.2

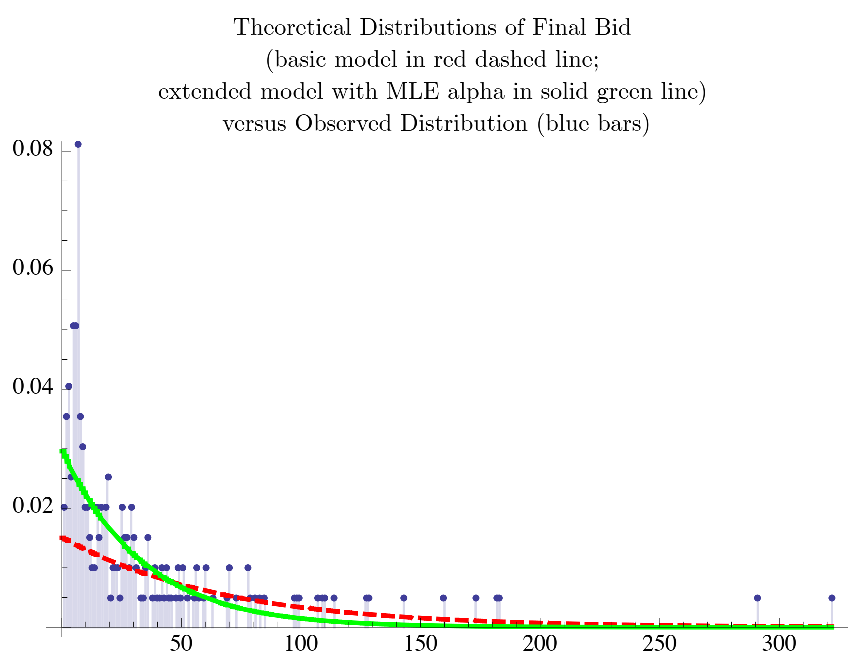

[17] In the Kmart $25 Gift Card graph, datas beyond period 300 are trimmed to show a clearer idea of the earlier stage of the game, while our CARA risk model does a much fit for the long tail.

<table>
<tr><th colspan="3">Kmart $25 Gift Card (RRP$25.00, Item#1122, Category: Gift Card)</th></tr>
<tr><td>Number of Observations</td><td>1,260</td><td rowspan="9">Theoretical Distributions of Final Bid
(basic model in red dashed line;
extended model with MLE alpha in solid green line)
versus Observed Distribution (blue bars)
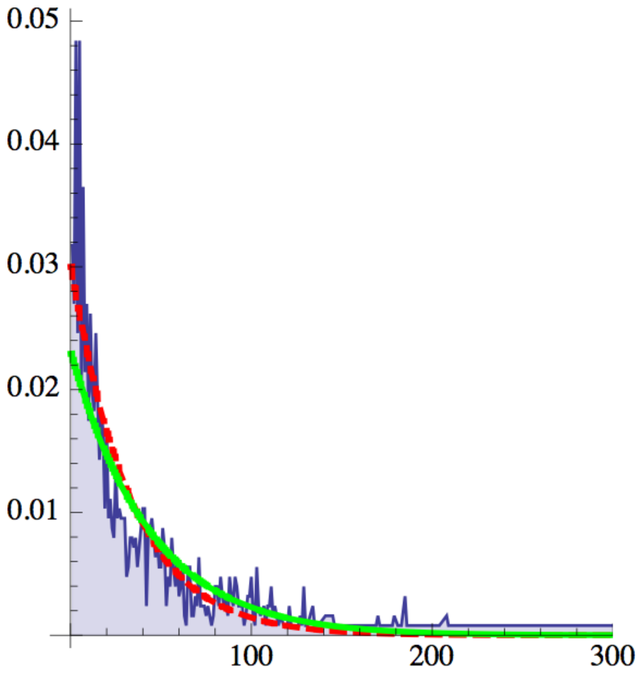
</td></tr>
<tr><td>Estimated Alpha by MLE</td><td>-0.0212377</td></tr>
<tr><td>Average Profit Margin</td><td>26.56%</td></tr>
<tr><td rowspan="2">t-test (p>=0.05)</td><td>Basic: 9.3</td></tr>
<tr><td>MLE: 92.9</td></tr>
<tr><td rowspan="2">Chi-Sq test (p>=0.05)</td><td>Basic: 13.3</td></tr>
<tr><td>MLE: 76.9</td></tr>
<tr><td rowspan="2">K-S test (p>=0.05)</td><td>Basic: 10.5</td></tr>
<tr><td>MLE: 87.0</td></tr>
</table>

Table 5. Comparison of theoretical and observed distributions of final bids

Our test results match our prior theoretical conclusion that bidders' risk-loving/risk-aversion either increase/decrease the seller's expected revenue, and the signs of the alpha Maximum Likelihood estimators coincide with the signs of the observed average seller's profit margin, which is also intuitive to understand in practice. For instance, a $25 Kmart gift card has a well-defined market value close to the RRP, which is useful to most bidders; while a stir frypan is likely to value less than its RRP and not many bidders would be interested in purchasing this good, i.e., auctions of those two items are likely to attract different groups of participants. We assumed all bidders had identical risk attitudes in the model setting for simpler analysis, while it is unlikely to be true in reality, and we may consider our alpha estimators as an average risk-attitude indicator of the participating bidders. Let us say, a gift card auction may attract more experienced bidders with more preferences toward risk, which leads to more aggressive bidding behaviours, a longer auction course and higher seller revenue. Our empirical findings also suggest that an experienced level of participating bidders in a penny auction is likely to be associated with the seller's revenue, and we will attempt to extend the model further to capture the effect of bidder experience in the next section.

Statistical tests for each item that have been repeatedly auctioned over 100 times are performed to compare theoretical distribution with observed data for basic model of $\alpha = 0$ and the extended model with MLE α, and a sample of the result is attached in the Appendix 2. We conclude that the basic model is inadequate when compared to the extended model with a risk attitude, which contributes towards explaining observed auction outcomes.

Our empirical findings show that the majority (86.91%) of bidders never play more than 20 auctions, while the minority (13.09%), who participated in 20 or more auctions, played more aggressively in the auctions, placing 81.36% of total bids[18]. The bidders who eventually gain 20 or more, prior to auction experience, place 13.30 bids per auction on average in their lifetime, while those who never participated in 20 or more auctions only placed 8.15 bids per auction on average in their lifetime. Intuitively, a variety of risk attitudes exist in bidders, which affects their bidding behaviour. For instance, more risk-averse bidders are more likely to bid less and are less likely to enrich their bidding experience, while more risk-seeking bidders, who get extra utility gain from bidding due to their love of risk-taking, become more experienced. The observed correlation between bidders' behaviour and their prior bidding experiences motivated us to analyse bidders in separate groups by their experiences.

Total Auctions Participated	Proportion in Bidder Population	Proportion of Bids in Total Bids Placed	Average Number of Bids Placed per Auction
Less than 20	86.91%	18.64%	8.15
20 and more	13.09%	81.36%	13.30

Table 6. More experienced and inexperienced bidder comparisons

On the other hand, a seller's profit margin also appears to correlate with bidders' experiences. The average profit margin of auctions with bidders that have less than 20 prior auction experiences is - 58.21% (loss), while the average profit margin of auctions with at least one bidder with 20 or more prior auction experiences is 51.51%. In Figure 15, we plotted the average seller's profit margin (ASPM) against bidders' experiences in terms of a rounded average prior auction (APA) participated in two auction item categories (Bid Packs and Gift Cards), two item types (50 Bid Packs and $50 Home Depot Gift Cards), and spotted some overall polynomial trends, where an ASPM grows in an APA at the start, then gradually drops as the APA becomes very large.[19] A

[18] All bids placed by these bidders are counted, including the bids placed in their first 20 auctions.

[19] Average prior auctions (APA) are defined as take an average of all bidders' prior auction counts in an auction, e.g. if there are 2 bidders who have played 2 and 4 auctions before, the APA is 3. We rounded APA raw results to whole numbers. The concerning small number of observations may result from randomness, as we focused on auctions of item types that were repeatedly auctioned 30 or more times. There are 41 auction items in 3 categories (19 in Bid Packs category, 21 in Gift Cards category) with enough observations.

similar trend was observed in last three diagrams of Figure 16 when we plotted an ASPM against rounded APAs of all auctions.

Figure 15. Average seller's profit margin against average bidder experience - 2 item categories and 2 item types

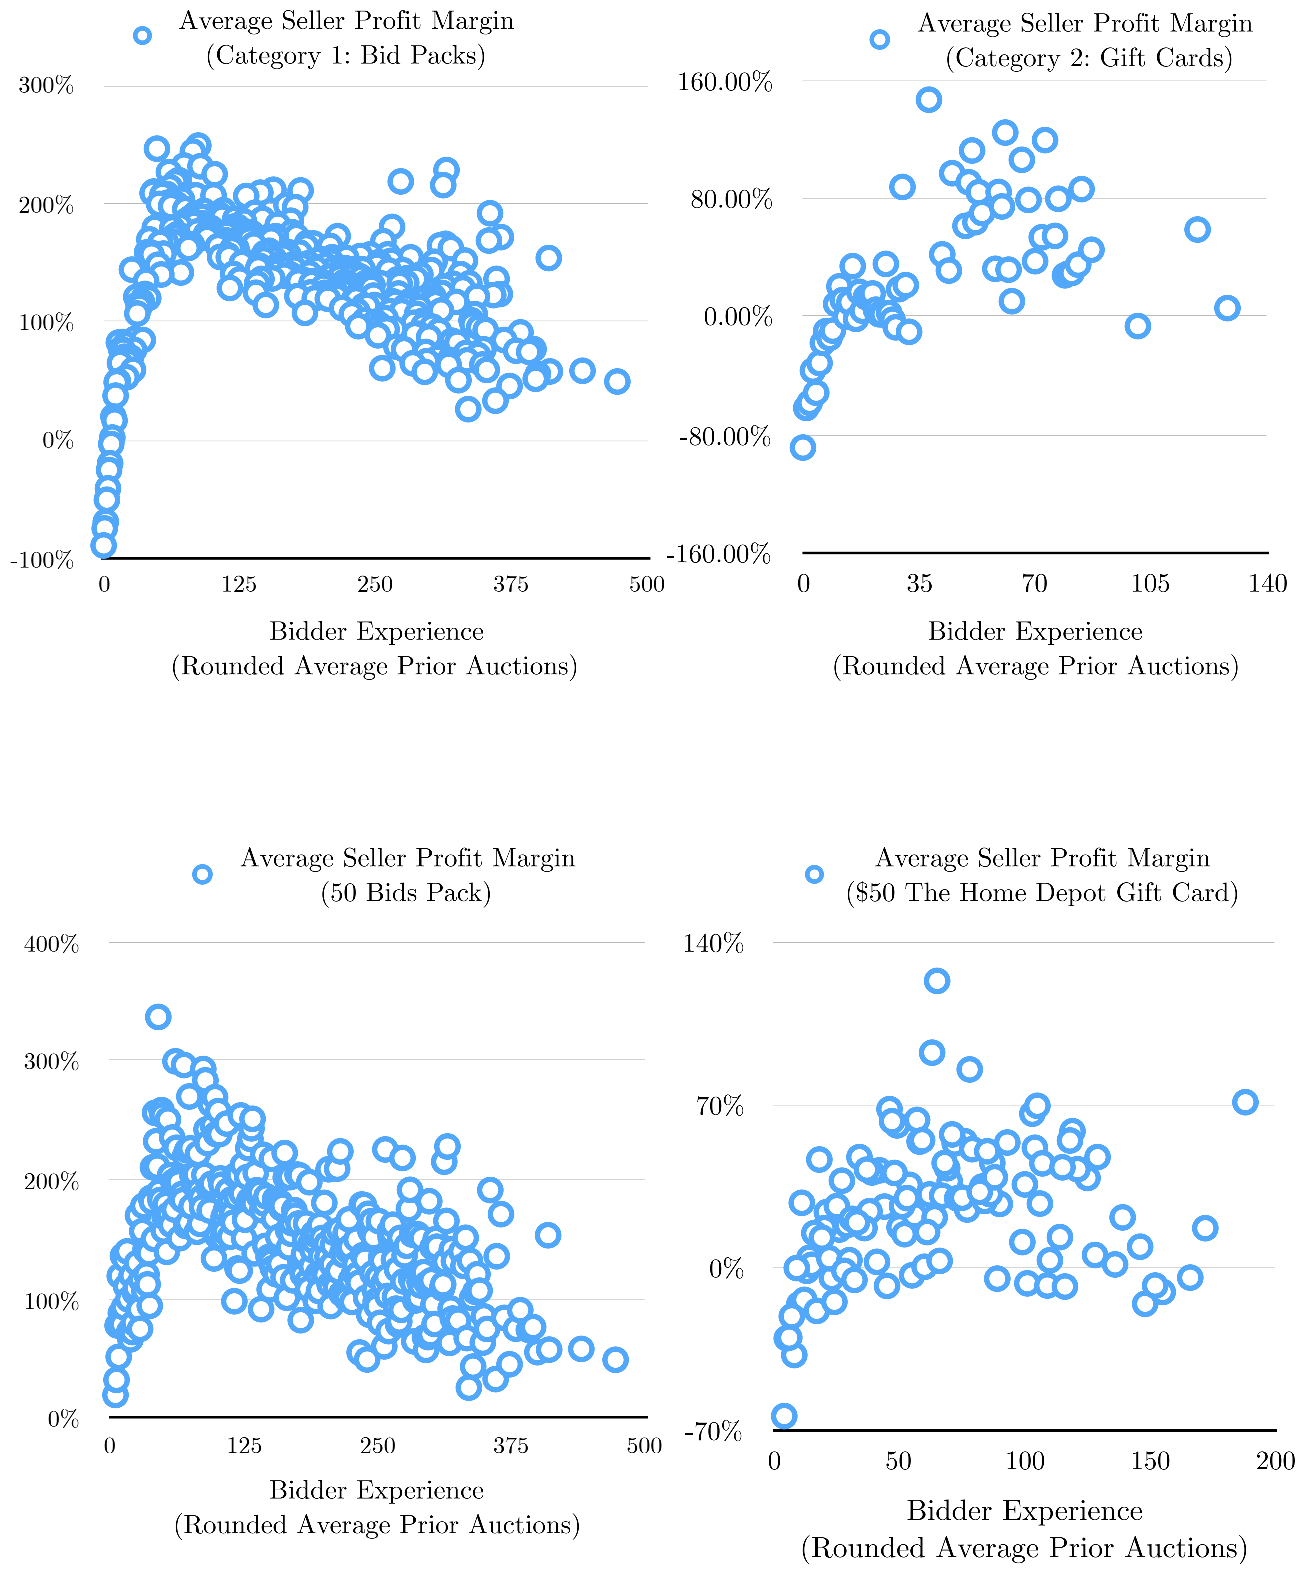

Figure 16. Average seller's profit margin against average bidder experience
- All auctions and zoomed in

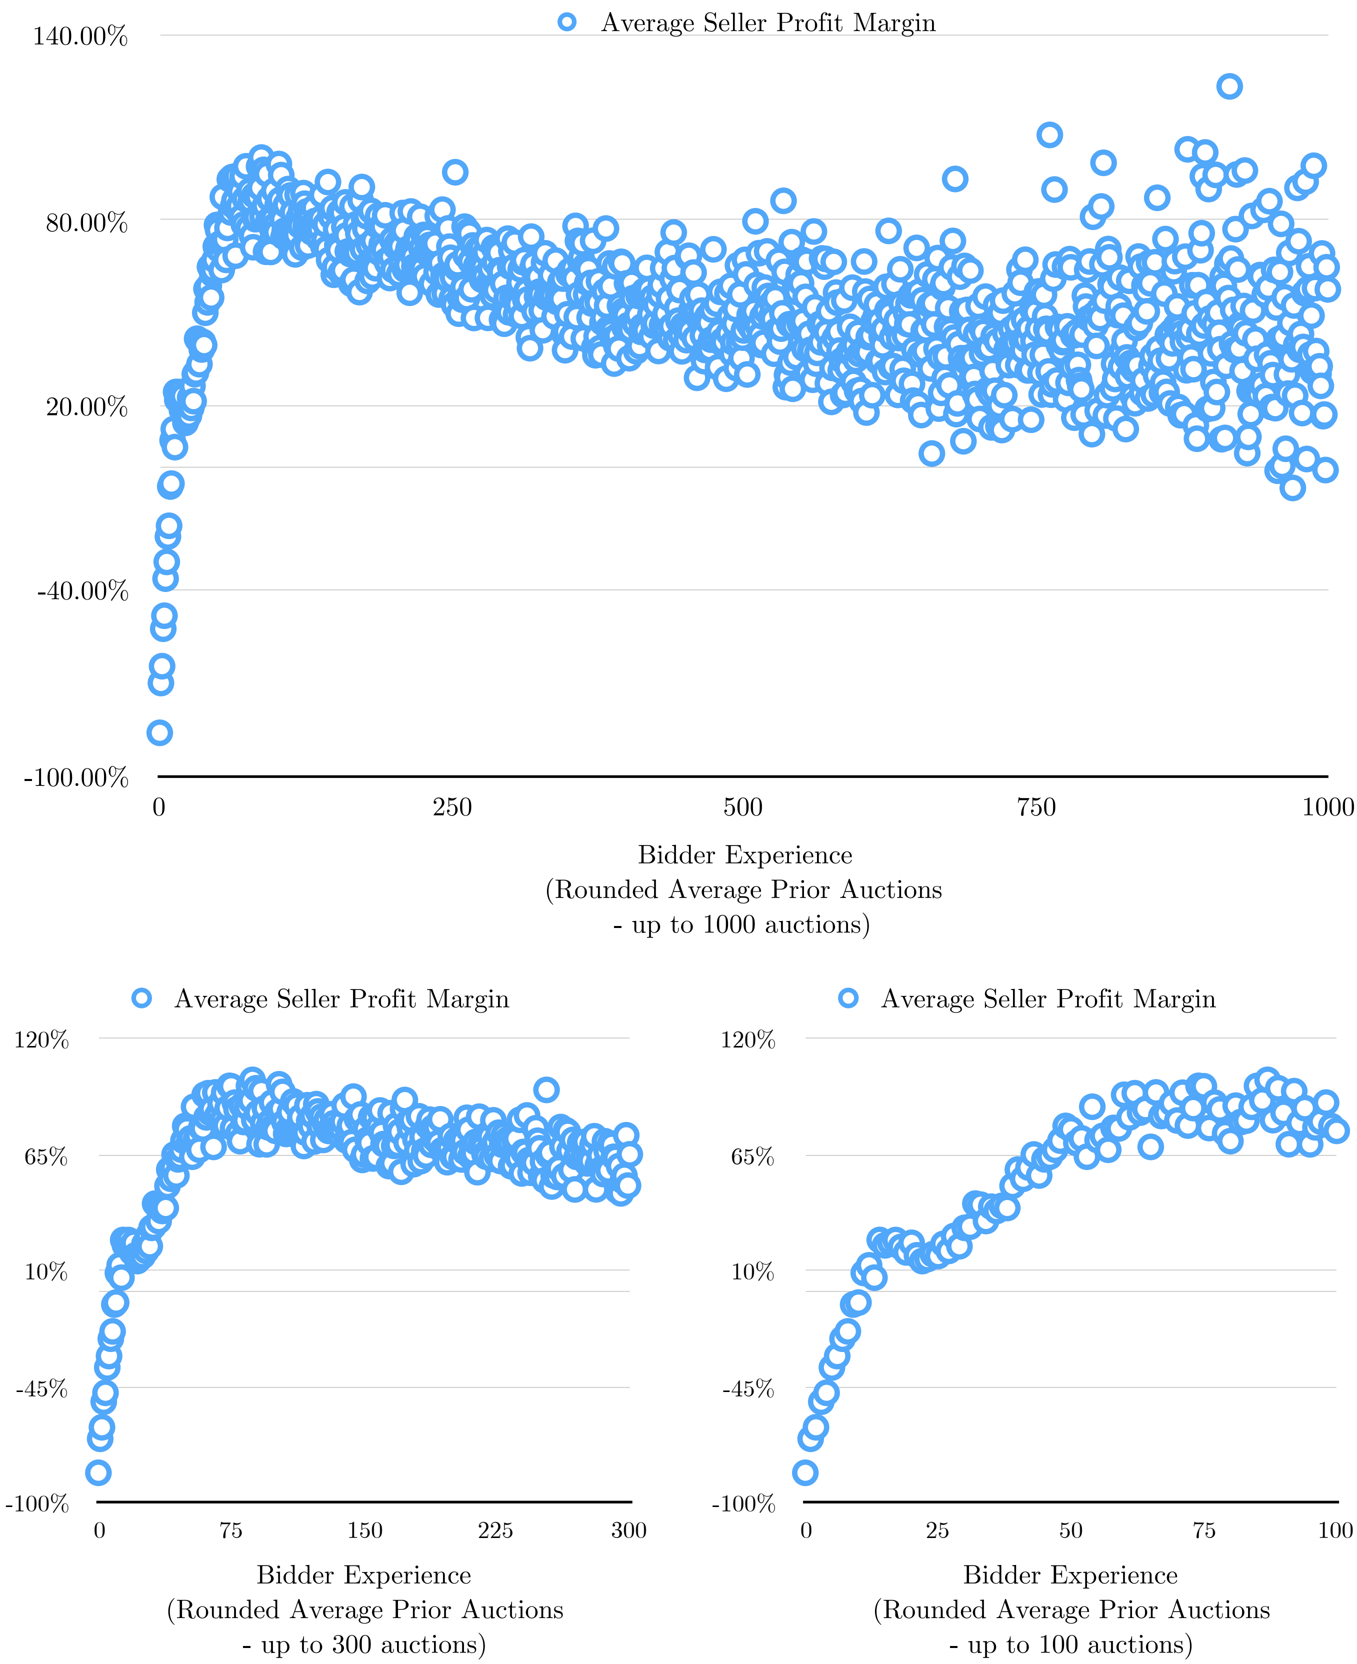

We defined our bidder experience groups by prior auction experience, and our experienced bidders as those who have played 20 or more auctions before participating in a particular auction, and treated the rest as inexperienced bidders. We chose 20 prior auctions to distinguish bidder types, as the majority of 86.91% of bidders never play more than 20 auctions, and as shown in Table 7, other prior auction experience ranges have much less bidders, e.g. only 3.48% of all bidders played between 20 and 25 auctions in their lifetime.

Total Auctions Participated	Numbers of Bidders	Ratio in Bidder Population	Average Number of Auctions Won	Average Number of Bids Placed per Auction	Winning Percentage
<5	97,514	38.20%	0.03	11.72	1.39%
[5,7)	36,589	14.33%	0.10	8.19	1.76%
[7,10)	38,174	14.95%	0.17	7.26	2.17%
[10,20)	49,600	19.43%	0.46	7.30	3.48%
[20,25)	8,884	3.48%	1.07	7.91	4.92%
[25,30)	5,326	2.09%	1.42	7.93	5.29%
[30,50)	9,428	3.69%	2.31	8.57	6.18%
[50,100)	5,083	1.99%	5.37	10.43	7.94%
[100,200)	2,098	0.82%	15.30	13.38	11.15%
[200,500)	1,443	0.57%	43.26	15.61	14.04%
[500,750)	378	0.15%	98.23	16.29	16.28%
[750,1000)	587	0.23%	119.70	16.29	17.19%
[1000,2000)	331	0.13%	283.67	18.72	20.90%
[2000,4000)	147	0.06%	527.69	15.57	18.92%
[4000,5000)	31	0.01%	905.03	13.05	20.50%
[5000,10000)	53	0.02%	1,369.66	13.16	20.37%
10000+	15	0.01%	2,445.93	7.66	19.23%
Total	255,304	100.00%	2.26	11.90	11.28%

Table 7. Bidders with different prior auction experiences

For items that are very frequently auctioned, we were able to find a large number of auctions with participation by newcomers only and those with participation by bidders with rich experience only. Thus, a straightforward hypothesis to test is whether the newcomers and experienced ones have the same risk attitudes, as we assumed in our previous model settings. We selected items that had been repeatedly auctioned over 100 times and had at least 30 auctions with experienced-only bidders, and 30 auctions with inexperienced-only bidders. Recall that we compute the maximum likelihood estimators (MLE) of alphas by choosing the parameter α to maximise $\sum_j \ln f(t_j;\alpha)$, where j

represents each observed auction of that item, t_j is the ending number of bids in that auction, and f is the discrete theoretical distribution in the CARA model. We did not only compute the maximum likelihood estimators of alphas for each experience group, but we also tested whether the alphas of the two groups were significantly different from each other by using the Likelihood Ratio test. Our null hypothesis is that bidders in different experience groups have the same risk attitude.

For all items that we had sufficient numbers of observed auctions to test, we rejected the null hypothesis, i.e. the empirical findings indicated that bidders in different experience groups had significantly different risk attitudes, thus experiences of bidders affected the seller's expected revenue, which coincided with our previous regression results. Table 8 shows an example of a single auction item we have tested – $25 Kmart Gift Card. Not only do the alpha estimates of each experience group look significantly different, but it is also supported by the result of the Likelihood Ratio test.[20]

$25 Kmart Gift Card	**Result**
Alpha MLE of Auctions with Inexperienced bidders only (subset of 118 auctions with only Inexperienced bidders)	0.0795
Alpha MLE of Auctions with Experienced bidders only (subset of 324 auctions with only Experienced bidders)	0.0277
Alpha MLE of Auctions with single-type bidders (442 auctions combining the above two subsets)	0.0382
Likelihood Ratio Test Result - to compare goodness of fit of either treating two type bidders with two alphas (alternative model), or treating everyone has one alpha (null model)	Test statistic 19.82 exceeds critical value 3.841 Hypothesis rejected

Table 8. Example of risk estimates for both risk experience groups

[20] A likelihood ratio test (LR test) is a statistical test used for comparing the goodness of fit of two statistical models, where the null model here (every bidder has same risk parameter), is a special case of the alternative model (experienced bidders and inexperienced bidders have different risk parameter). The test is based on the likelihood ratio, which expresses how many times more likely the data are under one model than the other. This likelihood ratio, or equivalently its logarithm, is compared to a critical value to decide whether or not to reject the null model. For instance, in $25 Kmart Gift Card case, since our test statistic 19.82 exceeds the critical value, it is concluded that the alternative model explains the data better.

In conclusion, the empirical findings suggest a correlation between experience and bidding aggressiveness, which are being picked up as differences in risk attitudes in the estimation. Bidders in different experience groups in our dataset are shown to have significantly different risk attitudes, with the more experienced group being more risk-seeking and bidding more aggressively, which increases the seller's expected revenue. Our result raises a puzzle by suggesting an apparent correlation between risk attitude and experience. One intuitive explanation would be that bidders with risk-seeking preferences are more likely to become addicted to this risky mechanism and gain experience through time (bidders who have participated in 1,000 and more auctions place an average of 16.63 bids per auction, while those who have never participated in 20 or more auctions place an average of 10 bids per auction), while risk-averse players are likely to be scared away from penny auction websites after a few attempts. If there are only inexperienced bidders with risk-aversion in a penny auction, sellers would earn negative payoffs, which coincides with the observation of an average profit margin in inexperienced-bidder-only auctions, being -74.79%, while auctions with experienced participants bring a seller an average profit margin of 49.76%.

Figure 17. Average seller's profit margin vs number of bidders in the auction

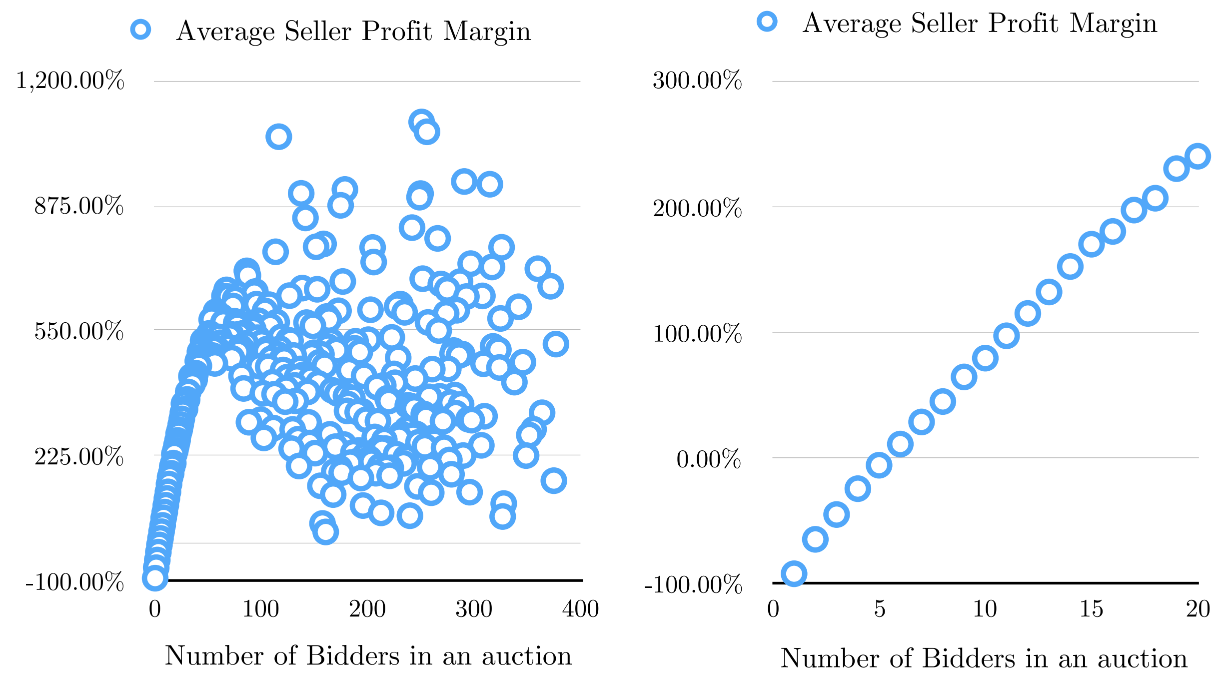

We recall that when every bidder has an identical (Constant Relative Risk Aversion) utility function (regardless of their risk preferences), a seller's expected revenue is independent of the number of bidders $\boldsymbol{N}$ in our current model as $\boldsymbol{N}$ is cancelled out in the calculation of final bid probability density function, which does not reflect the significant correlation between a seller's profit margin and the number of bidders observed in our dataset[21]. Figure 17 shows the average seller's profit margin against the number of bidders in all auctions, and a separate diagram focusing on auctions with up to 20 bidders, as there are on average 8.87 bidders competing in an auction, and 94.78% of total auctions have 20 or less bidders participating.

We have analysed subsets of auctions with single types of bidders separately, while the rest of the auctions, those with two types of bidders coexisting, were not included. In total, we have three types of auctions by bidder experience: auctions with only experienced bidders; auctions with only inexperienced bidders; and those with mixed types of bidders (at least one experienced bidder and one inexperienced bidder). Proportions of the three types of auctions with different numbers of participants are not fixed, but vary significantly. Figure 18 demonstrates the relative proportion distribution of the three types of auctions across subsets of auctions with numbers of bidders from 2 to 10, where proportion of auctions with single type of bidders decrease in $\boldsymbol{N}$. [22] In 65% of auctions with two observed bidders, both participants are experienced bidders; while in all auctions with four or more observed bidders, over half of the observations include mixed types of bidders; and 94% of auctions with ten observed bidders include mixed types of bidders.

On the other hand, the distribution of the number of bidders per auction is not uniform, but vary significantly. As shown in Figure 19, out of all the auctions, 97% have less than 25 participants and the peak is when there are 61,813 auctions with five bidders participating. Most auctions with only one type of bidder have less than 20 bidders participating.

---

[21] While our observed number of bidders in an auction is the only indicator we have for $\boldsymbol{N}$, it could be equal or less than the real $\boldsymbol{N}$, as we only acknowledge existence of a bidder in an auction when observing him/her place a bid successfully, and those who never place a successful bid are not observed. A more aggressive and risk-seeking bidder will, therefore, appear to "participate" more often in our dataset.

[22] We only look at subsets of auctions with numbers of bidders up to 10, as 77% of all auctions have 10 or less bidders participating and we do not have a large number of observations for some subsets for $N>10$.

Figure 18. Distribution of three types of auctions with varying numbers of bidders

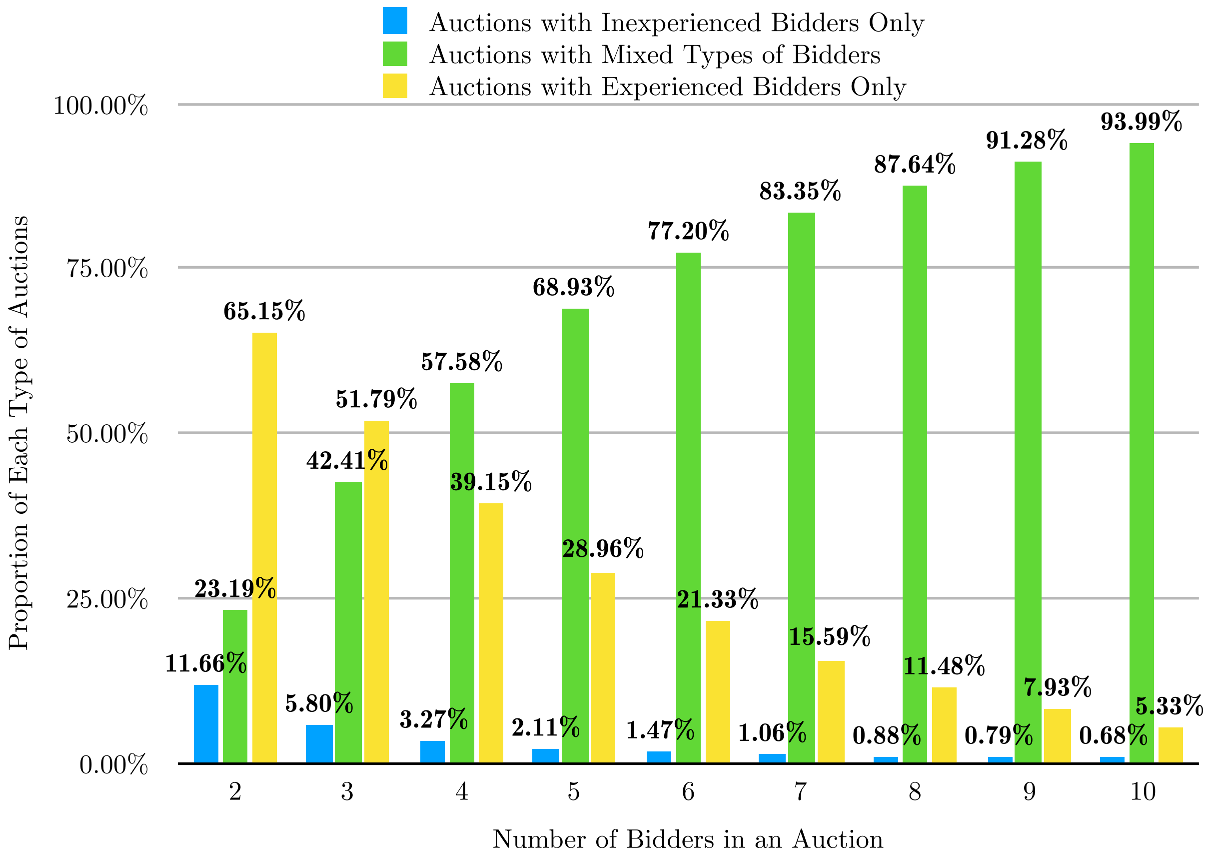

Figure 19. Distribution of number of bidders in three types of auctions

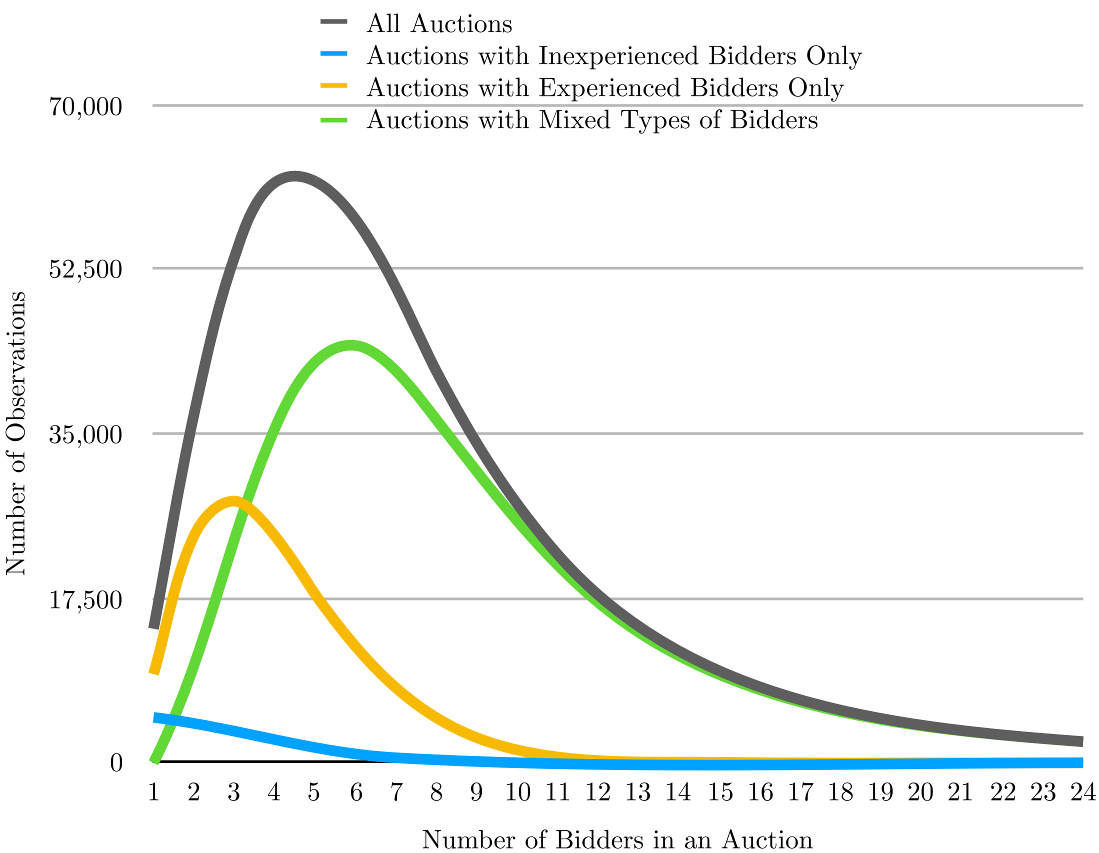

Table 9. Average of seller's profit margin (SPM) for different types of auctions with different numbers of bidders

N	SPM in Auctions with Experienced Bidders Only	SPM in Auctions with Mixed Types of Bidders	Ni	Ne	SPM in Auctions with Inexperienced Bidders Only
2	**-68.38%**	**-75.96%**	**1**	**1**	**-79.91%**
3	**-48.78%**	**-56.89%**	**1.28**	**1.72**	**-64.88%**
		-54.55%	1	2	
		-62.91%	2	1	
4	**-31.57%**	**-40.48%**	**1.55**	**2.45**	**-45.72%**
		-35.94%	1	3	
		-44.61%	2	2	
		-51.53%	3	1	
5	**-11.58%**	**-24.24%**	**1.87**	**3.13**	**-34.85%**
		-16.30%	1	4	
		-25.07%	2	3	
		-37.14%	3	2	
		-39.08%	4	1	
6	**6.87%**	**-9.39%**	**2.22**	**3.78**	**-13.39%**
		1.33%	1	5	
		-6.53%	2	4	
		-17.05%	3	3	
		-29.46%	4	2	
		-26.77%	5	1	
7	**27.22%**	**7.55%**	**2.61**	**4.39**	**-6.21%**
		20.11%	1	6	
		14.19%	2	5	
		2.31%	3	4	
		-4.61%	4	3	
		-13.76%	5	2	
		-8.39%	6	1	

Ni as Number of Inexperienced Bidders, Ne as Number of Experienced Bidders

Table 8 continues to next page

N	SPM in Auctions with Experienced Bidders Only	SPM in Auctions with Mixed Types of Bidders	Ni	Ne	SPM in Auctions with Inexperienced Bidders Only
8	**36.49%**	**26.93%**	**3.04**	**4.96**	**12.05%**
		44.82%	1	7	
		37.44%	2	6	
		28.63%	3	5	
		15.53%	4	4	
		8.66%	5	3	
		-3.22%	6	2	
		1.05%	7	1	
9	**47.49%**	**43.12%**	**3.54**	**5.46**	**31.44%**
		64.47%	1	8	
		59.24%	2	7	
		51.72%	3	6	
		39.39%	4	5	
		24.35%	5	4	
		15.79%	6	3	
		5.43%	7	2	
		18.77%	8	1	
10	**57.66%**	**56.25%**	**4.03**	**5.97**	**46.23%**
		84.41%	1	9	
		77.71%	2	8	
		67.90%	3	7	
		56.56%	4	6	
		45.10%	5	5	
		34.67%	6	4	
		24.36%	7	3	
		15.86%	8	2	
		27.96%	9	1	
11	**85.43%**	**76.72%**	**4.54**	**6.46**	**49.67%**
		105.40%	1	10	
		125.18%	2	9	
		90.33%	3	8	
		81.42%	4	7	
		67.21%	5	6	
		55.38%	6	5	
		50.01%	7	4	
		35.22%	8	3	
		40.57%	9	2	
		44.23%	10	1	

Table 9 shows the changes in a seller's average profit margin in all gift card auctions, across the three types of auctions (auctions including only inexperienced bidders, auctions including only experienced bidders, auctions including both inexperienced and experienced bidders), with the same number of observed bidders for all $\boldsymbol{N}$ between 2 and 10. For instance, in auctions with two bidders, the seller makes an average loss of -68.38% if both bidders are experienced, an average loss of -79.91% if both bidders are inexperienced, and an average loss of -75.96% if one experienced bidder and one inexperienced bidder coexist; in auctions with 11 bidders, the seller achieves an average profit margin of 49.67% if all bidders are experienced, an average profit margin of 85.43% if both bidders are inexperienced, and an average profit margin of 76.72% if both inexperienced and experienced bidders coexist, and the average number of observed inexperienced bidders and experienced bidders are 4.54 and 6.46.

As shown in Table 8, we observe that a seller's average profit margin increased as the number of bidders increased for each of the three types of auctions, e.g. a seller makes an average profit margin of -48.78%, -31.57%, -11.58%, 6.87%, 27.22% and 36.49% when there are only experienced participants and the number of bidders are 3, 4, 5, 6, 7 and 8 respectively. For auctions with the same number of bidders, a seller's average profit margin is highest in auctions with experienced bidders only, followed by auctions with mixed types of bidders, and is lowest in auctions with inexperienced bidders only. For auctions with mixed types of bidders, we observed a trend where the seller's average profit margin increases in the number of experienced bidders, for a fixed number of bidders, e.g. in auctions with five mixed types of bidders, a seller's average profit margin increases from -39.80% to -16.30% as the numbers of experienced bidders increase from 1 to 4. While for auctions with a larger number of bidders, the trend in profit margin looks potentially non-monotonic, as average profit margins fall in the number of inexperienced bidders, except for the last row when there is only one experienced bidder, where it rises. One initial explanation could be that when there are two or more experienced bidders, they compete more aggressively with rivals of the same type.

The above empirical findings motivated us to extend our model to capture the differences of various numbers of bidders with different prior auction experiences, allowing for the existence of only one type of participant and co-existence of both types of participants in one auction.

Chapter 6. Bidders of Two Types

Now we extend our model to fit auctions with both experienced and inexperienced bidders participating. Let Ne be the number of experienced bidders in an auction, and Ni be the number of inexperienced Bidders, then we have $N = Ni + Ne$. We assume all bidders have a Constant Absolute Risk Aversion (CARA) utility function of $u(W) = \frac{1-e^{-\alpha W}}{\alpha}$, where α is the Absolute Risk Aversion. Let $k \in \{i,e\}$ present a participant's own bidder type, bidders of the same type have identical alphas α^k, and $\beta_t(H_t) = \beta_t^{k,l_t}$.

Hazard rates given the current leader $l_t \in \{i,e\}$ are of either two types:

$$\begin{cases} h(t,l_t = e) = \left[1-\beta_t^{e,e}\right]^{N_e-1}\left[1-\beta_t^{i,e}\right]^{N_i} \\ h(t,l_t = i) = \left[1-\beta_t^{e,i}\right]^{N_e}\left[1-\beta_t^{i,i}\right]^{N_i-1} \end{cases}$$

By backward induction, if any Subgame Perfect Equilibrium exists where the auction survives beyond period 1 with some positive probability, it should satisfy the condition of indifference where any non-leader at period t $for\ 0 < t \le T$ must be indifferent between placing a bid and not placing a bid:

$$\begin{cases} h(t,l_t = e)u\left[W_{t-1}+v-st-c\right]+\left(1-h(t,l_t=e)\right)u\left[W_{t-1}-c\right] = u\left[W_{t-1}\right] \\ h(t,l_t = i)u\left[W_{t-1}+v-st-c\right]+\left(1-h(t,l_t=i)\right)u\left[W_{t-1}-c\right] = u\left[W_{t-1}\right] \end{cases}$$

Solving for hazard rates, we get

$$\begin{cases} h(t,l_t = e) = \left[1-\beta_t^{e,e}\right]^{N_e-1}\left[1-\beta_t^{i,e}\right]^{N_i} = K_t^e \\ h(t,l_t = i) = \left[1-\beta_t^{e,i}\right]^{N_e}\left[1-\beta_t^{i,i}\right]^{N_i-1} = K_t^i \end{cases}$$

where $K_t^e = \frac{e^{\alpha^e c}-1}{e^{\alpha^e c}-e^{\alpha^e(c+st-v)}}$ and $K_t^i = \frac{e^{\alpha^i c}-1}{e^{\alpha^i c}-e^{\alpha^i(c+st-v)}}$ at period $0 < t \le T$

K_t^k depends on the risk attitude parameter α^k of bidder type $k \in \{i,e\}$ for any given t, and $K_t^k \in [0,1]$ as $v \geq c + st$. The equilibrium condition is a set of two equations with four β_t^{k,l_t} that need to be solved, such that multiple equilibria exist.

Likewise, since we are looking at Subgame Perfect Equilibrium where the auction survives beyond period 1 with positive probability, we choose $h(0) = 0$ arbitrarily, i.e. some bidding always occurs at period 0 in equilibrium, so that the auction reaches period 1 with positive probability, as the auctioneer can run the auction repeatedly until some players bid in period 0, leading the hazard rate of the initial period to effectively be zero. Other symmetric equilibria also exist that we are not interested in, such as the auction either ending in period 0 with no bidder, or ending in period 1 with one winning bidder.

To calculate the seller's expected revenue

$$E[Rev] = \sum_{t=1}^{T+1} (c+s)t \cdot f(t)$$

we build the probability density function that the auction ends at period t

$$f(t) = \tilde{h}(t) \prod_{m=1}^{t} \left[1 - \tilde{h}(m-1)\right]$$

where the hazard rate at period t

$$\tilde{h}(t) = \mathbf{Pr}\left[l_t = e\right] h\left(t, l_t = e\right) + \mathbf{Pr}\left[l_t = i\right] h\left(t, l_t = i\right) = \mathbf{Pr}\left[l_t = e\right] K_t^e + \mathbf{Pr}\left[l_t = i\right] K_t^i$$

and probability that the current leader at period t is a particular bidder type $\mathbf{Pr}\left[l_t = e\right]$ and $\mathbf{Pr}\left[l_t = i\right]$ depend on $\left\{\beta_0^k, \beta_1^{k,l_1}, \ldots, \beta_t^{k,l_t}\right\}$.

Since we have multiple solutions of β_t^{k,l_t} in equilibrium condition, it is mathematically too difficult to compute *f(t)* directly. In order to proceed to the application of the new model, we go through two approaches: a special case when $N=2$ where there are existing unique solutions of β_t^{k,l_t} in the equilibrium we are interested in, and a general case when $\boldsymbol{N}>\boldsymbol{2}$, with added restriction on β_t^{k,l_t} to pin down some interesting equilibria.

N=2 Two-bidder Model

In 6% of the total auctions we observe, there are only 2 bidders participating. When ***N=2*** and the two bidders are of different types[23], we have a unique subgame perfect equilibrium where the auction survives beyond period 1 with some positive probability that satisfies

$$\begin{cases} h(t,l_{t=e}) = 1-\beta_t^i = K_t^e \\ h(t,l_{t=i}) = 1-\beta_t^e = K_t^i \end{cases}$$

Similarly, to compute expected revenue, we construct probability density function that the auction ends at period t, $f(t)=\tilde{h}(t)\prod_{m=1}^{t}\left[1-\tilde{h}(m-1)\right]$, where the hazard rate $\tilde{h}(t)=\mathbf{Pr}\left[l_t=e\right]h\left(t,l_t=e\right)+\mathbf{Pr}\left[l_t=i\right]h\left(t,l_t=i\right)$. Let $\theta=\mathbf{Pr}\left[l_1=i\right]$, the probability that the current leader at period t is a particular bidder type

$$\mathbf{Pr}\left[l_t=e\right]=\begin{cases} \theta\prod_{n=1}^{\frac{t}{2}}\beta_{2n-1}^e\prod_{n=1}^{\frac{t}{2}-1}\beta_{2n}^i & \textit{when } t \textit{ is even} \\ \left(1-\theta\right)\prod_{n=1}^{\frac{t-1}{2}}\beta_{2n-1}^i\prod_{n=1}^{\frac{t-1}{2}}\beta_{2n}^e & \textit{when } t \textit{ is odd} \end{cases}$$

$$\mathbf{Pr}\left[l_t=i\right]=\begin{cases} \left(1-\theta\right)\prod_{n=1}^{\frac{t}{2}}\beta_{2n-1}^i\prod_{n=1}^{\frac{t}{2}-1}\beta_{2n}^e & \textit{when } t \textit{ is even} \\ \theta\prod_{n=1}^{\frac{t-1}{2}}\beta_{2n-1}^e\prod_{n=1}^{\frac{t-1}{2}}\beta_{2n}^i & \textit{when } t \textit{ is odd} \end{cases}$$

---

23 When both bidders are of the same type, the case is simplified to a single-type bidder problem

Thus, we have

$$f(t)=\begin{cases}\theta \prod_{n=1}^{\frac{t}{2}} \beta_{2n-1}^{e} \prod_{n=1}^{\frac{t}{2}-1} \beta_{2n}^{i} \left(1-\beta_t^i\right) + \left(1-\theta\right) \prod_{n=1}^{\frac{t}{2}} \beta_{2n-1}^{i} \prod_{n=1}^{\frac{t}{2}-1} \beta_{2n}^{e} \left(1-\beta_t^e\right) & \textbf{\textit{when t is even}} \\ \theta \prod_{n=1}^{\frac{t-1}{2}} \beta_{2n-1}^{e} \prod_{n=1}^{\frac{t-1}{2}} \beta_{2n}^{i} \left(1-\beta_t^e\right) + \left(1-\theta\right) \prod_{n=1}^{\frac{t-1}{2}} \beta_{2n-1}^{i} \prod_{n=1}^{\frac{t-1}{2}} \beta_{2n}^{e} \left(1-\beta_t^i\right) & \textbf{\textit{when t is odd}}\end{cases}$$

And seller's expected revenue $E[Rev]$ decreases in bidders' risk attitudes α^k, i.e. the more risk seeking the bidders are, the higher expected revenue will be. Seller makes profit when both bidders are risk-seeking with negative α^i and α^e, makes loss when both bidders are risk-averse with positive α^i and α^e, and may make profit or loss when one bidder is risk-seeking and the other bidder is risk-averse.

Recall that our key theoretical prediction is that the final number of bids in a given auction is a random variable with distribution $f(t,\alpha)$. If a given item type is repeatedly auctioned many times, we will have enough numbers of observation to estimate risk parameters. Maximum likelihood estimates of $\alpha=\left(\alpha^i,\alpha^e\right)$ are estimated by choosing the parameter α to maximise $\sum_j \ln f(t_j;\alpha)$, where j represents each observed auction of that item, t_j is the ending number of bids in that auction.

Note that when ***N=2***, the two bidders may be of the same or different types. In the first case, the problem is simplified to a single-type bidder case, where the probability density function of final bids ending at t for auctions with only one type of bidder $k\in\left\{i,e\right\}$, is rewritten as below for ***N=2*** model for simplicity:

$$f^k(t;\alpha^k)=\prod_{m=1}^{t-1}\beta_m^k\left(1-\beta_t^k\right)\ \ \forall t\geq 2$$

And when the bidders are of different types, the probability density function of final bids ending at t and is rewritten as below, instead of the one above:

$$f^{mixed}(t;\alpha^i,\alpha^e)=\begin{cases}\prod_{n=1}^{\frac{t}{2}}\beta^e_{2n-1}\prod_{n=1}^{\frac{t}{2}-1}\beta^i_{2n}\ \left(1-\beta^i_t\right) & \textit{when t is even, } l_t=e\\ \prod_{n=1}^{\frac{t}{2}}\beta^i_{2n-1}\prod_{n=1}^{\frac{t}{2}-1}\beta^e_{2n}\ \left(1-\beta^e_t\right) & \textit{when t is even, } l_t=i\\ \prod_{n=1}^{\frac{t-1}{2}}\beta^e_{2n-1}\prod_{n=1}^{\frac{t-1}{2}}\beta^i_{2n}\ \left(1-\beta^e_t\right) & \textit{when t is odd, } l_t=i\\ \prod_{n=1}^{\frac{t-1}{2}}\beta^i_{2n-1}\prod_{n=1}^{\frac{t-1}{2}}\beta^e_{2n}\ \left(1-\beta^i_t\right) & \textit{when t is odd, } l_t=e\end{cases}$$

Thus when choosing $\alpha=\left(\alpha^i,\alpha^e\right)$ to maximise $\sum_j \ln f(t_j;\alpha)$, where $j\in\left\{j^i,j^e,j^{mixed}\right\}$ represents each observed auction of that item that j^i represent each observation having two inexperienced bidders, j^e represent each observation having two experienced bidders, j^{mixed} represent each observation having one inexperienced bidder and one experienced bidder, t_j is the ending number of bids in that auction, and

$$f(t_j;\alpha)=\begin{cases}f^i(t;\alpha^i) & \textit{when both bidders are type i}\\ f^e(t;\alpha^e) & \textit{when both bidders are type e}\\ f^{mixed}(t;\alpha^i,\alpha^e) & \textit{when bidders are different types}\end{cases}$$

, we maximise $\sum_j \ln f(t_j;\alpha)=\sum_{j^i}\ln f^i(t_{j^i};\alpha^i)+\sum_{j^e}\ln f(t_{j^e};\alpha^e)+\sum_{j^m}\ln f^{mixed}(t_{j^m};\alpha^i,\alpha^e)$

We look at 134 auction item types where each has 30 or more auctions with only 2 bidders participating, and compute alpha MLE using two models - ***N=2*** Single-Bidder-Type model and ***N=2*** Two-Bidder-Type model. Table 10 shows an example of one auction item type we have tested - $25 CVS Gift Card, where there are 257 repeatedly run auctions with only 2 bidders. We can see the two alpha estimates of each type of bidders in the Two-Bidder-Type model $\alpha = \left(\alpha^i, \alpha^e\right) = \left(0.2095, 0.1160\right)$, which are significantly different from the alpha estimate $\alpha = 0.1390$. in the Single-Bidder-Type model as in previous chapter. When treating the single-type model as the null model and the two-type model as the alternative model, our likelihood test statistic is greater than the critical value thus we reject the null model being a better fit of the data, suggesting that the Two-Bidder-Type model explains the data better than the Single-Bidder-Type model, and the two types of bidders have different risk attitudes.[24]

MLE in different models	**Result**
Using N=2 Two-Bidder-Type Model - Alpha MLE of Experienced Bidders	0.1160
Using N=2 Two-Bidder-Type Model - Alpha MLE of Inexperienced Bidders	0.2095
Using N=2 Single-Bidder-Type Model - Alpha MLE of Inexperienced Bidders	0.1390
Likelihood Ratio Test Result - to compare goodness of fit of either treating two type bidders with two alphas (alternative model), or treating everyone has one alpha (null model)	Test statistic 22.64 exceeds critical value 3.841 Hypothesis rejected

Table 10. Example of alpha MLE in N=2 two-type model vs single-type model

In our study of all 134 item types, we have likelihood test statistics exceeding the critical value, indicating that our two-type model matches our data better, and we always have the risk parameter of the experienced bidder group smaller than the risk

[24] A likelihood ratio test (LR test) is a statistical test used for comparing the goodness of fit of two statistical models, where the null model, the single-type model in our problem, is a special case of the alternative model the two-type model. The test is based on the likelihood ratio, which expresses how many times more likely the data are under one model than the other. This likelihood ratio, or equivalently its logarithm, is compared to a critical value to decide whether or not to reject the null model.

parameter of the inexperienced bidder group, suggesting that bidders in different experience groups in our dataset are shown to have significantly different risk attitudes, with the more experienced group being more risk-seeking and bidding more aggressively. Our application also suggests that an auction with two experienced bidders participating generates the highest expected revenue; an auction with two inexperienced bidders participating generates the lowest expected revenue; and the expected revenue of an auction with one experienced bidder and one inexperienced bidder falls in between, which matches our observations.

N > 2 Two-bidder-type Model with Restriction

Since we do not have a unique solution of bidders' bidding strategy profile β_t^{k,l_t} in our two-bidder-type general model for $N>2$, we attempt to add a restriction to pin down a subset of solutions that we are more interested in. We recall the basic backward induction logic of our two-type equilibrium: we have that the hazard rate of a given period must depend on the leader's type and $K_t^k = \frac{e^{\alpha^k c} - 1}{e^{\alpha^k c} - e^{\alpha^k (c+st-v)}}$ increases in α^k, and we have that the more risk-averse the current leader is, the higher the probability that the auctions ends, since the hazard rate must be such that the current leader is indifferent between bidding in the previous period (to become the next leader) and not bidding.

Let us consider an equilibrium in which $\beta_t^{e,l_t} = \beta_t^{i,l_t} = \beta_t^{l_t}$ for all $l_t \in \{e,i\}$, that bidders' bidding strategies depend only on the type of the current leader, but not his/her own type.[25] Thus, our equilibrium conditions are simplified to:

$$\begin{cases} \left[1-\beta_t^e\right]^{N_e-1}\left[1-\beta_t^e\right]^{N_i} = K_t^e \\ \left[1-\beta_t^i\right]^{N_e}\left[1-\beta_t^i\right]^{N_i-1} = K_t^i \end{cases}$$

which provide a unique solution

$$\begin{cases} \beta_t^e = 1-\left(K_t^e\right)^{1/(N-1)} \\ \beta_t^i = 1-\left(K_t^i\right)^{1/(N-1)} \end{cases}$$

The probability density function that the auction ends at period t

$$f(t) = \tilde{h}(t)\prod_{m=1}^{t}\left[1-\tilde{h}(m-1)\right]$$

, where the hazard rate at period t

[25] I have also tried another restriction that bidders choose their bidding strategies depending on their own types, not on the type of the current leader, restricting four betas to two. The approach was abandoned as the equilibrium solutions of hazard rates are not always between zero and one.

$$\tilde{h}(t) = \mathbf{Pr}\left[l_t = e\right] K_t^e + \mathbf{Pr}\left[l_t = i\right] K_t^i$$

where $\mathbf{Pr}\left[l_t = e\right]$ and $\mathbf{Pr}\left[l_t = i\right]$ depend on $\left\{\beta_0, \beta_1^{l_1}, \ldots, \beta_t^{l_t}\right\}$. Since it is too complicated to write and compute the unconditional probabilities of the current leader being a certain type at period t in bidding strategies, we used recurring relationships between conditional probabilities of the current leader being a certain type at period t given such a period is reached.

Let the conditional probability of the current leader being experienced at period t given t is reached, is denoted by

$$\Pi_t^e = \mathbf{Pr}\left[l_t = e \mid t \textit{ reached}\right]$$

and conditional probability of the current leader being inexperienced at period t given t

$$\Pi_t^i = \mathbf{Pr}\left[l_t = i \mid t \textit{ reached}\right] = 1 - \Pi_t^e$$

that when $0 < t \le T$,

$$\Pi_t^e = \frac{\Pi_{t-1}^e\left(1 - K_{t-1}^e\right)\left(\dfrac{N^e - 1}{N - 1}\right) + \Pi_{t-1}^i\left(1 - K_{t-1}^i\right)\left(\dfrac{N^e}{N - 1}\right)}{\Pi_{t-1}^e\left(1 - K_{t-1}^e\right) + \Pi_{t-1}^i\left(1 - K_{t-1}^i\right)}$$

$$\tilde{h}(t) = \Pi_t^e \cdot K_t^e + \left(1 - \Pi_t^e\right) \cdot K_t^i$$

Similarly, maximum likelihood estimates of $\alpha = \left(\alpha^i, \alpha^e\right)$ are estimated by choosing the parameter α to maximise $\sum_j \ln f(t_j; \alpha)$. When $N > 2$, an auction may have all bidders of one type or mixture of two types, then we maximise the following: [26]

$$\sum_j \ln f(t_j;\alpha) = \sum_{j^i} \ln f^i(t_{j^i};\alpha^i) + \sum_{j^e} \ln f^e(t_{j^e};\alpha^e) + \sum_{j^m} \ln f^{mixed}(t_{j^m};\alpha^i,\alpha^e)$$

where

[26] Although an auction with one type of bidders is a special case of an auction with two types of bidders, here we used the probability density function of one-type risk model as in earlier chapters, due to computation problem when applying the two-type-bidder model onto one-type-bidder obsecrations.

$$f^i(t;\alpha^i) = K_t^i \prod_{m=1}^{t} \left(\frac{1 - e^{\alpha^i (c+s(m-1)-v)}}{e^{\alpha^i c} - e^{\alpha^i (c+s(m-1)-v)}} \right)$$

$$f^e(t;\alpha^e) = K_t^e \prod_{m=1}^{t} \left(\frac{1 - e^{\alpha^e (c+s(m-1)-v)}}{e^{\alpha^e c} - e^{\alpha^e (c+s(m-1)-v)}} \right)$$

$$f^{mixed}(t) = \tilde{h}(t) \prod_{m=1}^{t} \left[1 - \tilde{h}(m-1) \right]$$

$$\tilde{h}(t) = \Pi_t^e \cdot K_t^e + \left(1 - \Pi_t^e\right) \cdot K_t^i$$

Seller's expected revenue $E[Rev]$ decreases in bidders' risk attitudes α^k, i.e. the more risk seeking the bidders are, the higher expected revenue will be. Seller makes profit when bidders are risk-seeking with negative α^i and α^e, makes loss when bidders are risk-averse with positive α^i and α^e, and may make profit or loss when some bidders are risk-seeking and the others are risk-averse.

We looked at 69 auction item types where each had 30 or more auctions with only inexperienced participants; 30 or more auctions with only experienced participants; and 30 or more auctions with both inexperienced and experienced participants. We computed alpha MLE using two models – Single-Bidder-Type model as in Chapter 4 and $N>2$ restricted the Two-Bidder-Type model as above. Table 11 shows examples of five auction item types we have tested. We can see α^i and α^e, the two alpha estimates of each type of bidder in the restricted Two-Bidder-Type model, which are significantly different from the alpha estimate in the Single-Bidder-Type model. For instance, in the Subway $10 Gift Card auctions, we can see the two alpha estimates of each type of bidder in the Two-Bidder-Type model $\alpha = \left(\alpha^i, \alpha^e\right) = \left(0.0869, \text{-}0.1631\right)$, which are significantly different from the alpha estimate in the Single-Bidder-Type model $\alpha = -0.1306$.

Treating the single-type model as the null model and restricted two-type model as the alternative model, our likelihood test statistics are all greater than the critical value where we reject the null model being a better fit of the data, suggesting that the restricted Two-Bidder-Type model explains the data better than the Single-Bidder-Type model, and the two types of bidders have different risk attitudes.

In our study of all 69 item types, we have likelihood test statistics exceeding the critical value, indicating that our restricted two-type model matches our data better, and we always have risk parameters of the experienced bidder group smaller than the risk parameter of the inexperienced bidder group, suggesting that bidders in different experience groups in our dataset are shown to have significantly different risk attitudes, with the more experienced group being more risk-seeking and bidding more aggressively.

Expected revenue $\boldsymbol{E}[\boldsymbol{Rev}] = \sum_{t=1}^{T+1} (\boldsymbol{c}+\boldsymbol{s})\boldsymbol{t}\cdot\boldsymbol{f}(\boldsymbol{t})$, and $\boldsymbol{f}(\boldsymbol{t}) = \tilde{\boldsymbol{h}}(\boldsymbol{t}) \prod_{m=1}^{t} \left[1-\tilde{\boldsymbol{h}}(\boldsymbol{m}-1)\right]$, where the $\tilde{\boldsymbol{h}}(\boldsymbol{t}) = \mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{e}\right]\boldsymbol{K}_t^e + \mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{i}\right]\boldsymbol{K}_t^i$. When there are only experienced participants in an auction, $\mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{e}\right] = 1$ that $\boldsymbol{h}^e(\boldsymbol{t}) = \boldsymbol{K}_t^e$; when there are only inexperienced participants in an auction, $\mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{i}\right] = 1$ that $\boldsymbol{h}^i(\boldsymbol{t}) = \boldsymbol{K}_t^i$; when there are both experienced participants and inexperienced participants in an auction, $\mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{e}\right] \in (0,1)$ and $\mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{i}\right] \in (0,1)$, that $\boldsymbol{h}^{mixed}(\boldsymbol{t}) = \mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{e}\right]\boldsymbol{K}_t^e + \mathbf{Pr}\left[\boldsymbol{l}_t = \boldsymbol{i}\right]\boldsymbol{K}_t^i \in \left(\boldsymbol{K}_t^e, \boldsymbol{K}_t^i\right)$, as experienced bidders are found to be more risk seeking. Thus $\boldsymbol{h}^e(\boldsymbol{t}) < \boldsymbol{h}^{mixed}(\boldsymbol{t}) < \boldsymbol{h}^i(\boldsymbol{t})$, indicating $\boldsymbol{E}^e\left[\boldsymbol{Rev}\right] < \boldsymbol{E}^{mixed}\left[\boldsymbol{Rev}\right] < \boldsymbol{E}^i\left[\boldsymbol{Rev}\right]$ when $\alpha^e < \alpha^i$.

Thus our model suggests that an auction with experienced participants only generates the highest expected revenue; an auction with inexperienced participants only, generating the lowest expected revenue; and the expected revenue of an auction with both experienced and inexperienced participants falling in between. This does not match our observations if we do not remove the effects of the number of bidders (as in examples shown in Table 11, e.g. average profit margin in straightly mixed auctions of Mobil $50 Gift Card auctions is 78.62% greater than -30.62%, average profit margin in auctions with experienced bidders only), and match our observations well if we study subsets of observations with the same number of bidders (Table 12 shows an example of separating observations of Best Buy $50 Gift Card into subsets by number of bidders. For any subsets with the same number of bidders, we observed that the average profit margin of auctions with both experienced and inexperienced participants falls between an average profit margin of auctions with inexperienced participants only and an average profit margin of auctions with experienced participants only. For instance, an average profit margin of auctions with five straightly mixed experienced and inexperienced bidders is -41.87%, which is between -46.24% and -20.32%, and average profit margin of auctions with five inexperienced bidders and an average profit margin of auctions with five experienced bidders). We recall that the number of bidders in an auction does not affect the seller's expected revenue in our one-type-bidder model, which does not match our empirical findings of an average profit margin of one-type-bidder observations increasing in the number of bidders. This matter will be discussed further in a later chapter.

Since $\boldsymbol{f}^{mixed}(\boldsymbol{t}) = \tilde{\boldsymbol{h}}(\boldsymbol{t})\prod_{\boldsymbol{m}=1}^{\boldsymbol{t}}\left[1-\tilde{\boldsymbol{h}}(\boldsymbol{m}-1)\right]$ with $\tilde{\boldsymbol{h}}(\boldsymbol{t}) = \Pi_t^e \cdot \boldsymbol{K}_t^e + \left(1-\Pi_t^e\right)\cdot \boldsymbol{K}_t^i$, and Π_t^e increases in $\boldsymbol{N}^e$ when $\boldsymbol{N}$ is fixed, $\boldsymbol{E}^{mixed}\left[\boldsymbol{Rev}\right]$ increases in $\boldsymbol{N}^e$ with fixed $\boldsymbol{N}$. That is, for a fixed number of bidders $\boldsymbol{N}$, the expected revenue of an auction with both experienced participants and inexperienced participants increases in $\boldsymbol{N}^e$, number of experienced participants and decreases in $\boldsymbol{N}^i$, number of inexperienced participants, which match the majority of our observations, e.g. finding about Shell $50 Gift Card as shown in earlier chapter, except for some auctions with only one experienced participant and two or more inexperienced participants.

Note that Π_t^e also increases in $\boldsymbol{N}$ when ratio $\boldsymbol{N}^e/\boldsymbol{N}$ is fixed, and that $\boldsymbol{E}^{mixed}\left[\boldsymbol{Rev}\right]$ increases in $\boldsymbol{N}$ with fixed $\boldsymbol{N}^e/\boldsymbol{N}$. That is, for a fixed ratio of experienced bidders in all bidders $\boldsymbol{N}^e/\boldsymbol{N}$ in an auction with both experienced participants and inexperienced participants, expected revenue increases in $\boldsymbol{N}$, which match our observations. As shown in Table 12, average ratios of $\boldsymbol{N}^e/\boldsymbol{N}$ are all between 55% and 60% in subsets of Best Buy $50 Gift Card auctions with $\boldsymbol{N}$=3, 4, …, 8, and the average profit margins of the subsets increase as $\boldsymbol{N}$ increases from 3 to 8.

Table 11. Example of alpha MLE in N ≥ 2 mixed-type model vs single-type model

Auction Item	Golfsmith $25	Omaha Steak $25	Toys R Us $20	Mobil $50	Subway $10
Total Auction Count	622	684	1,818	855	1,658
Average Seller Profit Margin	-41.59%	-20.08%	6.01%	55.50%	81.43%
Average Number of Bidders	4.51	4.30	4.75	7.82	5.33
Treat All Bidder Identical					
Alpha Estimate	0.0453	0.0165	-0.0099	-0.0179	-0.1306
Auctions with Inexperienced Participants Only					
Average Seller Profit Margin	-89.70%	-85.05%	-72.27%	-44.65%	-56.34%
Average Number of Bidders	1.73	2.00	2.52	5.83	2.64
Alpha_i Estimate	0.3448	0.2463	0.1644	0.0261	0.1731
Auctions with Experienced Participants Only					
Average Seller Profit Margin	-51.27%	-33.08%	-19.48%	-30.62%	29.07%
Average Number of Bidders	3.49	3.92	4.12	4.92	4.61
Alpha_e Estimate	0.0638	0.0324	0.0189	0.0153	-0.0666
Auctions with Straightly Mixed Types of Participants					
Average Seller PM	-9.88%	6.82%	43.56%	78.62%	153.97%
Average Number of Bidders	7.27	5.18	5.76	8.57	6.43
Alpha_i Estimate	0.1773	0.1207	0.0575	-0.0176	0.0688
Alpha_e Estimate	-0.0797	-0.0968	-0.1026	-0.0259	-0.3702
All Auctions - General Two-Type Model					
Alpha_i Estimate	0.2244	0.1597	0.0967	0.0074	0.0869
Alpha_e Estimate	0.0104	0.0205	-0.0246	-0.0069	-0.1631
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0

Table 12. Example of alpha MLE in subsets with different number of bidders

Best Buy $50	**N=3**	**N=4**	**N=5**	**N=6**	**N=7**	**N=8**
Total Auction Count	681	847	1,013	958	933	835
Average Seller PM	-64.89%	-50.48%	-39.18%	-23.96%	-7.28%	15.99%
Average N^e / N	72.69%	66.79%	64.26%	61.60%	58.69%	55.55%
Treat All Bidder Identical						
Alpha Estimate	0.0524	0.0320	0.0215	0.0111	0.0027	-0.0066
Auctions with Inexperienced Participants Only						
Average Seller PM	**-70.28%**	**-57.87%**	**-46.24%**	**-42.72%**	**-40.13%**	**0.32%**
Alpha_i Estimate	0.0637	0.0411	0.0277	0.0244	0.0223	-0.0007
Auctions with Experienced Participants Only						
Average Seller PM	**-61.79%**	**-40.12%**	**-20.32%**	**-16.31%**	**7.67%**	**92.97%**
Alpha_e Estimate	0.0470	0.0223	0.0091	0.0069	-0.0036	-0.0256
Auctions with Straightly Mixed Two Types of Participants						
Average Seller PM	**-66.84%**	**-53.98%**	**-41.87%**	**-24.33%**	**-7.43%**	**13.96%**
Average N^e / N	57.45%	58.04%	60.02%	59.14%	57.62%	54.90%
Alpha_i Estimate	0.0134	0.0963	0.0700	0.0391	0.0286	0.0088
Alpha_e Estimate	0.1203	-0.0125	-0.0102	-0.0092	-0.0189	-0.0186
All Auctions - General Two-Type Model						
Alpha_i Estimate	0.0678	0.0462	0.0327	0.0251	0.0215	
Alpha_e Estimate	0.0489	0.0231	0.0097	0.0050	-0.0055	-0.0238
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0

Chapter 7. Bidders with Three types

Another interesting fact found in the data is that a tiny proportion of bidders (576 counted) placed 11,920,373 bids at the time they had 1,000 or more prior auction experiences, and that 0.2% of total bidders placed 20% of the total bids. Their risk attitudes (or strategy profiles) are more difficult to be analysed in a straightforward fashion, as it seems they play differently in different types of auctions. In super-experienced bidders only auctions, they often tend to play conservatively – acting more risk-averse compared to Ordinary-Experienced Bidders (20-1,000 p.a.) when they play against rivals of the same type, especially when the number of bidders playing in an auction is small. Auctions with mixed types of bidders play more aggressively.

Now we extend our model to fit auctions with three types of participants:

- *inexperienced bidders* who have participated less than 20 prior auctions,
- *ordinary-experienced bidders* who have participated more than 20 and less than 1000 prior auctions; and
- *super-experienced bidders* who have participated 1000 or more prior auctions.

Let $\boldsymbol{Ni}$ be the number of inexperienced bidders, $\boldsymbol{Ne}$ be the number of ordinary-experienced bidders in an auction, $\boldsymbol{Ns}$ be the number of super-experienced bidders in an auction, so that we have $\boldsymbol{N = Ni + Ne + Ns}$. We assume all bidders have a Constant Absolute Risk Aversion (CARA) utility function of $\boldsymbol{u(W)} = \frac{1-\boldsymbol{e}^{-\alpha \boldsymbol{W}}}{\alpha}$, where α is the Absolute Risk Aversion, and bidders of the same type have identical alphas. Let $\boldsymbol{k} \in \{\boldsymbol{i,e,s}\}$ present a participant's own bidder type, that we have $\beta(\boldsymbol{H}_t) = \beta_t^{\boldsymbol{k},\boldsymbol{l}_t}$. Similarly, let us add a restriction and consider an equilibrium in which $\beta_t^{i,l_t} = \beta_t^{e,l_t} = \beta_t^{s,l_t} = \beta_t^{l_t}$ for all $\boldsymbol{l}_t \in \{\boldsymbol{i,e,s}\}$, that bidders' bidding strategies depend on the type of the current leader, but not his/her own type. Thus, our equilibrium conditions are now:

$$\begin{cases} h(t, l_t = i) = \left[1-\beta_t^i\right]^{N_e}\left[1-\beta_t^i\right]^{N_i-1}\left[1-\beta_t^i\right]^{N_s} = K_t^i \\ h(t, l_t = e) = \left[1-\beta_t^e\right]^{N_e-1}\left[1-\beta_t^e\right]^{N_i}\left[1-\beta_t^e\right]^{N_s} = K_t^e \\ h(t, l_t = s) = \left[1-\beta_t^s\right]^{N_e}\left[1-\beta_t^s\right]^{N_i}\left[1-\beta_t^s\right]^{N_s-1} = K_t^s \end{cases}$$

which provide a unique solution

$$\begin{cases} \beta_t^i = 1-\left(K_t^i\right)^{1/(N-1)} \\ \beta_t^e = 1-\left(K_t^e\right)^{1/(N-1)} \\ \beta_t^s = 1-\left(K_t^s\right)^{1/(N-1)} \end{cases}$$

where $K_t^{l_t} = \dfrac{e^{\alpha^{l_t} c} - 1}{e^{\alpha^{l_t} c} - e^{\alpha^{l_t}(c+st-v)}}$ for all $l_t \in \{i, e, s\}$.

We compute expected revenue

$$E[Rev] = \sum_{t=1}^{T+1} (c+s)t \cdot f(t)$$

by building the probability density function that an auction ends at period t

$$f(t) = \tilde{h}(t) \prod_{m=1}^{t}\left[1-\tilde{h}(m-1)\right]$$

where the hazard rate at period t

$$\tilde{h}(t) = \mathbf{Pr}\left[l_t = i\right] K_t^i + \mathbf{Pr}\left[l_t = e\right] K_t^e + \mathbf{Pr}\left[l_t = s\right] K_t^s$$

where $\mathbf{Pr}\left[l_t = i\right]$, $\mathbf{Pr}\left[l_t = e\right]$ and $\mathbf{Pr}\left[l_t = s\right]$ depend on $\left\{\beta_0, \beta_1^{l_1}, \ldots, \beta_t^{l_t}\right\}$.

Similarly, it is too complicated to write and compute the unconditional probabilities of the current leader being a certain type at period t in bidding strategies, so we use recurring relationships between conditional probabilities of the current leader being a certain type at period t given such period is reached instead.

Let the conditional probabilities of the current leader being each type at period $\boldsymbol{t}$ given $\boldsymbol{t}$ is reached

$$\Pi_t^i = \mathbf{Pr}\left[l_t = i \mid t \ reached \right]$$
$$\Pi_t^e = \mathbf{Pr}\left[l_t = e \mid t \ reached \right]$$
$$\Pi_t^s = \mathbf{Pr}\left[l_t = s \mid t \ reached \right] = 1 - \Pi_t^i - \Pi_t^e$$

so that we have

$$\Pi_t^i = \frac{\Pi_{t-1}^i \left(1 - K_{t-1}^i\right)\left(\dfrac{N^i - 1}{N - 1}\right) + \Pi_{t-1}^e \left(1 - K_{t-1}^e\right)\left(\dfrac{N^i}{N - 1}\right) + \Pi_{t-1}^s \left(1 - K_{t-1}^s\right)\left(\dfrac{N^i}{N - 1}\right)}{\Pi_{t-1}^i \left(1 - K_{t-1}^i\right) + \Pi_{t-1}^e \left(1 - K_{t-1}^e\right) + \Pi_{t-1}^s \left(1 - K_{t-1}^s\right)}$$

$$\Pi_t^e = \frac{\Pi_{t-1}^i \left(1 - K_{t-1}^i\right)\left(\dfrac{N^e}{N - 1}\right) + \Pi_{t-1}^e \left(1 - K_{t-1}^e\right)\left(\dfrac{N^e - 1}{N - 1}\right) + \Pi_{t-1}^s \left(1 - K_{t-1}^s\right)\left(\dfrac{N^e}{N - 1}\right)}{\Pi_{t-1}^i \left(1 - K_{t-1}^i\right) + \Pi_{t-1}^e \left(1 - K_{t-1}^e\right) + \Pi_{t-1}^s \left(1 - K_{t-1}^s\right)}$$

$$\tilde{h}(t) = \Pi_t^i \cdot K_t^i + \Pi_t^e \cdot K_t^e + \left(1 - \Pi_t^i - \Pi_t^e\right) \cdot K_t^s$$

We can show that expected revenue increases in all three alphas. A seller's expected revenue $E[Rev]$ decreases in bidders' risk attitudes α^k, i.e. the more risk-seeking the bidders are, the higher the expected revenue will be. A seller makes a profit when bidders are risk-seeking with negative α^i, α^e and α^s, and makes a loss when bidders are risk-averse with positive α^i, α^e and α^s, and may make a profit or a loss when some bidders are risk-seeking and the others are risk-averse. Assuming $\alpha^i < \alpha^e, and\ \alpha^i < \alpha^s$, the expected revenue also increases in N^e, *and* N^s for fixed N.

Similarly, maximum likelihood estimates of $\alpha = \left(\alpha^i, \alpha^e, \alpha^s\right)$ for an auction with any combination of three types of bidders, are estimated by choosing the parameter α to maximise $\sum_j \ln f(t_j; \alpha)$. When $N > 2$, an auction may have all bidders of one type, or a mixture of two or three types, i.e. seven possible scenarios that we maximise[27]

$$\sum_j \ln f(t_j;\alpha) = \sum_{j^i} \ln f^i(t_{j^i};\alpha^i) + \sum_{j^e} \ln f^e(t_{j^e};\alpha^e) + \sum_{j^s} \ln f^s(t_{j^s};\alpha^s)$$
$$+ \sum_{j^m} \ln f^{mixed\ i\&e}(t_{j^m};\alpha^i,\alpha^e) + \sum_{j^m} \ln f^{mixed\ i\&s}(t_{j^m};\alpha^i,\alpha^s) + \sum_{j^m} \ln f^{mixed\ e\&s}(t_{j^m};\alpha^e,\alpha^s)$$
$$+ \sum_{j^m} \ln f^{mixed3types}(t_{j^m};\alpha^i,\alpha^e,\alpha^s)$$

I am unable to solve the above maximisation due to computational difficulty, as we are unable to apply a general three-type-bidder model onto all observations, while we can compute MLE of the three alphas within each of the seven subsets, each representing one of the seven scenarios, e.g. auctions with three types of straight-mixed participants.

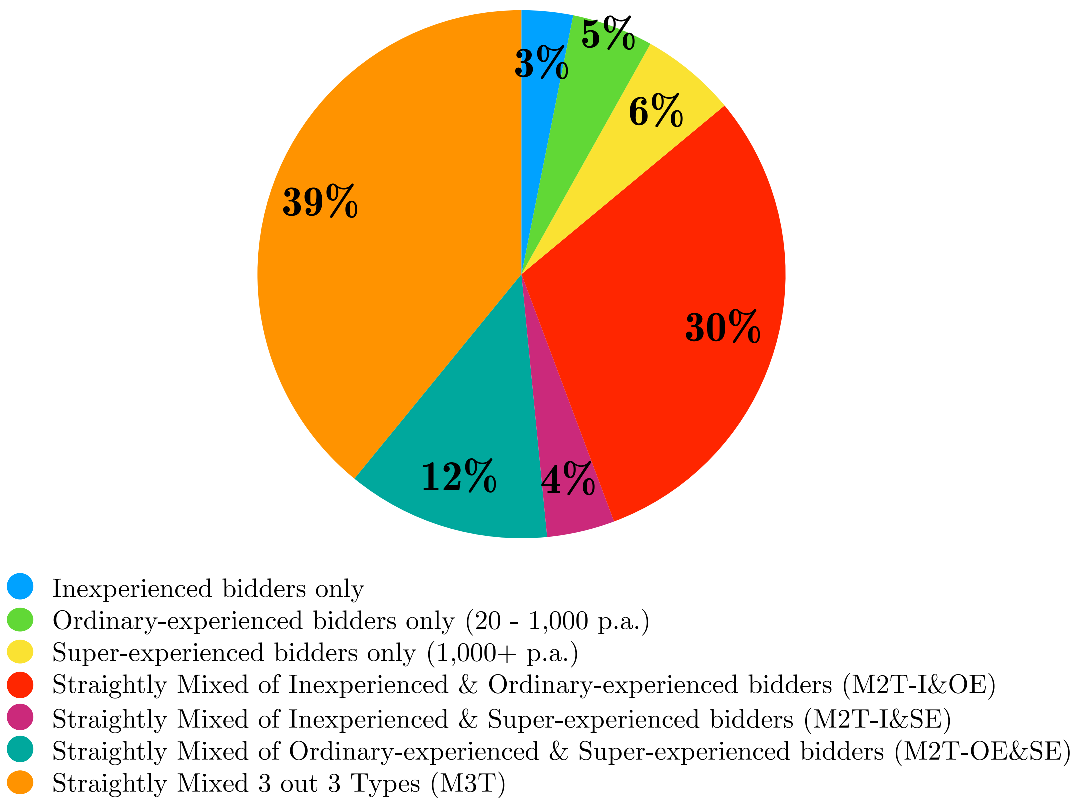

Figure 20. Distribution of observations of seven scenarios

[27] Although an auction with one or two types of bidders are special cases in an auction with three types of bidders, we have computation difficulty when applying the three-type-bidder model onto obsecrations with one-type bidders only, or a mixture of two types.

Figure 20 shows the observation distribution of the seven scenarios so the majority of the observations are of two scenarios: 39% of the total observed auctions have three types of straight-mixed participants (M3T), and 30% of the total observed auctions have straight-mixed, inexperienced and ordinary-experienced participants only (M2T-I&OE).

Table 13 shows a finding of three item types, each with seven subsets. For instance, by analysing subsets of Omaha Steak $25 Gift Card auctions, we observe:

- Auctions with only inexperienced bidders (1T-I) generate the largest average loss -85.05% for the seller, followed by -64.41% of auctions with only ordinary-experienced bidders (1T-OE), and -57.69% of auctions with only super-experienced bidders (1T-SE), with alpha estimators $\alpha^s < \alpha^e < \alpha^i$ as 0.0783 < 0.1000 < 0.2463;

- Observations with only two types of participant (M2T) auctions with super-experienced and inexperienced bidders (M2T-I&SE) generate the highest average loss of -51.91%, followed by -39.47% in auctions with ordinary-experienced and inexperienced bidders (M2T-I&OE), and -13.54% in auctions with super-experienced and ordinary-experienced bidders (M2T-OE&SE), with an alpha estimator α^i higher than α^e and α^s, and $\alpha^e < \alpha^s$[28];

- Auctions with all three types of participants (M3T) generate the highest average profit margin, and we have alpha estimators $\alpha^e < \alpha^s < \alpha^i$ as -0.1673 < -0.1027 < 0.0882, showing that ordinary-experienced bidders are the most risk-seeking, or least risk-averse when all three types of bidders are playing against each other.

- All types of bidders are bidding more aggressively when there are rivals of other types present.

[28] One intuitive explanation is that super-experienced bidders are likely to be able to tell bidder types of their rivals from observing their usernames and play accordingly, such that they play more aggressively, as in the Game of Chicken when facing rivals of their own types, and slightly less aggressively playing against ordinary-experienced bidders only, and much less aggressively playing against inexperienced bidders.

– Our Likelihood Ratio tests suggest that all three types of bidders have different risk parameters.

Table 13. Example of alpha MLE in N>2 model with three types of bidders

Item Name	**Golfsmith $25**	**Omaha Steak $25**	**Subway $10**
Total Auction Count	622	684	1,658
Average Number of Bidders	7.27	5.18	6.43
Average Seller Profit Margin	-41.59%	-20.08%	81.43%
Auctions with Inexperienced Participants Only			
Average Seller Profit Margin	**-89.70%**	**-85.05%**	**-56.34%**
Average Number of Bidders	1.73	2.00	2.64
Alpha_i Estimate	0.3448	0.2463	0.1731
Auctions with Ordinary Experienced Participants Only			
Average Seller Profit Margin	**-71.09%**	**-64.41%**	**-41.99%**
Average Number of Bidders	2.08	2.21	2.10
Alpha_oe Estimate	0.1275	0.1000	0.1024
Auctions with Super Experienced Participants Only			
Average Seller Profit Margin	**-59.02%**	**-57.69%**	**-19.71%**
Average Number of Bidders	3.23	4.10	4.02
Alpha_se Estimate	0.0830	0.0793	0.0297
Auctions with Inexperienced and Ordinary Experienced bidders only			
Average Seller Profit Margin	**-48.28%**	**-39.47%**	**74.22%**
Average Number of Bidders	3.50	3.77	4.14
Mean (Ni)	1.50	1.95	2.31
Mean (Ne)	2.00	1.81	1.83
Alpha_i Estimate	0.1648	0.0958	-0.0846
Alpha_oe Estimate	-0.0036	-0.0054	-0.1496
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0

Item Name	Golfsmith $25	Omaha Steak $25	Subway $10
Auctions with Inexperienced and Super Experienced bidders only			
Average Seller Profit Margin	**-54.04%**	**-51.91%**	**52.73%**
Average Number of Bidders	4.00	3.73	5.05
Mean (Ni)	1.18	1.27	1.42
Mean (Ns)	2.82	2.45	3.63
Alpha_i Estimate	0.2807	0.2246	0.1022
Alpha_se Estimate	0.0077	-0.0043	-0.1792
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0
Auctions with Ordinary Experienced and Super Experienced bidders only			
Average Seller Profit Margin	**-39.27%**	**-13.54%**	**62.42%**
Average Number of Bidders	4.26	4.50	5.12
Mean (Ne)	1.90	1.85	1.66
Mean (Ns)	2.36	2.65	3.45
Alpha_oe Estimate	0.0426	-0.0137	-0.1307
Alpha_se Estimate	0.0501	0.0343	-0.0572
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0
Auctions with Three Type of bidders			
Average Seller Profit Margin	**29.14%**	**57.85%**	**190.72%**
Average Number of Bidders	10.91	6.68	7.10
Mean (Ni)	1.54	1.42	1.67
Mean (Ne)	6.43	2.51	1.98
Mean (Ns)	2.94	2.75	3.45
Alpha_i Estimate	0.0778	0.0882	-0.1875
Alpha_oe Estimate	-0.0841	-0.1673	-0.2318
Alpha_se Estimate	-0.0507	-0.1027	-0.2157
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0

Note that the highest average number of bidders may contribute to the highest average profit margin of M3T auctions also. Assuming ordinary-experienced bidders are more aggressive than super-experienced bidders when playing in auctions existing rivals of different types of bidders, as suggested by our findings in alpha MLE, then for auctions with same number of bidders, the expected revenue of M3T auctions should be

- lower than expected revenue of M2T-I&OE auctions when

$$\boldsymbol{N}^{e,m2t-i\&oe} / \boldsymbol{N}^{m2t-i\&oe} = \left(\boldsymbol{N}^{e,m3t} + \boldsymbol{N}^{s,m3t}\right) / \boldsymbol{N}^{m3t}$$

- and higher than expected revenue of M2T-I&SE auctions when

$$\boldsymbol{N}^{s,m2t-i\&se} / \boldsymbol{N}^{m2t-i\&se} = \left(\boldsymbol{N}^{e,m3t} + \boldsymbol{N}^{s,m3t}\right) / \boldsymbol{N}^{m3t}$$

Table 14 shows an example of comparing three subsets of $50 Visa Gift Cards in three scenarios (M2T-I&OE, M2T-I&SE, M3T), with a fixed number of bidders of six in each of the subsets. In each auction of the M2T-I&OE subset, the number of ordinary-experienced bidders is 3; in each auction of the M2T-I&SE subset, the number of super-experienced bidders is 3; and in each auction of the M3T subset, the sum number of ordinary-experienced and super-experienced bidders is also 3.

The average seller's profit margin of the M2T-I&OE auctions is highest at 72.02%, followed by 65.49% of the average seller's profit margin of the M3T auctions, and 59.09% of the average seller's profit margin of the M2T-I&SE auctions. In the findings of the M3T auctions, we have alpha estimators $\alpha^e < \alpha^s < \alpha^i$ as -0.0823 < -0.0507 < -0.0149, which also match our previous findings.

Table 14. Example of auctions with two of three types of bidders VS all three types of bidders

Visa $50 Gift Card, N=6	**M2T- I &OE Ne=3**	**M2T- I &SE Ns=3**	**M3T Ne + Ns = 3**
Auctions with Inexperienced and Ordinary Experienced bidders only			
Average Seller Profit Margin	72.02%	Not Applicable	Not Applicable
Alpha_i Estimate	-0.0207		
Alpha_oe Estimate	-0.0941		
Auctions with Inexperienced and Super Experienced bidders only			
Average Seller Profit Margin	Not Applicable	59.09%	Not Applicable
Alpha_i Estimate		-0.0237	
Alpha_se Estimate		-0.0651	
Auctions with Three Type of bidders			
Average Seller Profit Margin	Not Applicable	Not Applicable	65.49%
Mean (Ne)			1.7396
Mean (Ns)			1.2604
Alpha_i Estimate			-0.0149
Alpha_oe Estimate			-0.0823
Alpha_se Estimate			-0.0507

Table 15 shows an example of the Walmart $100 Gift Card M3T auctions by studying auctions with a fixed number of bidders $\boldsymbol{N}$, and a fixed sum of the number of ordinary-experienced and super-experienced bidders $\boldsymbol{Ne+Ns}$. As we can see, the average seller's profit margin increases in $\boldsymbol{N}$; and for fixed $\boldsymbol{Ne+Ns}$, the average seller's profit margin increases in $\boldsymbol{Ne}$, which matches our model predication in assuming ordinary-experienced bidders are more aggressive in auctions with rivals of other types as found previously.

Table 15. Example of M3T auctions with different bidder distribution

N	Ne+Ns	Ne	Ns	Average Seller's Profit Margin in Auctions with Mixed Types of Bidders
4	**3**	**1.72**	**1.28**	**-43.51%**
		1	2	-44.72%
		2	1	-43.04%
5	**3**	**1.68**	**1.32**	**-36.84%**
		1	2	-45.00%
		2	1	-33.05%
	4	**2.44**	**1.56**	**-34.94%**
		1	3	-54.63%
		2	2	-42.44%
		3	1	-20.24%
6	**3**	**1.77**	**1.23**	**-21.06%**
		1	2	-22.03%
		2	1	-20.77%
	4	**2.57**	**1.43**	**-20.23%**
		1	3	-23.29%
		2	2	-21.17%
		3	1	-16.38%
	5	**3.27**	**1.73**	**-20.01%**
		1	4	-28.72%
		2	3	-20.37%
		3	2	-18.49%
		4	1	-12.33%
7	**3**	**1.86**	**1.14**	**-19.47%**
		1	2	-21.80%
		2	1	-17.28%
	4	**2.61**	**1.39**	**-0.48%**
		1	3	-8.73%
		2	2	-2.68%
		3	1	3.89%
	5	**3.22**	**1.78**	**2.37%**
		1	4	-3.15%
		2	3	-0.63%
		3	2	3.56%
		4	1	5.80%
	6	**3.88**	**2.13**	**14.39%**
		1	5	2.63%
		2	4	5.13%
		3	3	9.80%
		4	2	5.25%
		5	1	18.22%

I studied all 15 item types that have 100 or more observations in each of the subsets, and the findings are similar to the above examples. As expected, revenue increases in the alpha one-type-bidder model; expected revenue increases in both alphas in the straightly mixed, two-type-bidder model; and expected revenue increases in all three alphas in the straightly mixed, two-type-bidder model. The empirical findings match our model that:

1. Alpha MLE:

   - When there are only one type of bidders in an auction (1T):

   Alpha MLE of inexperienced bidders is highest, indicating they are the most risk-averse, followed by ordinary-experienced bidders, and super-experienced bidders, i.e. $\alpha^{s,1t-s} < \alpha^{e,1t-e} < \alpha^{i,1t-i}$.

   - When there are two of the three types of bidders in an auction (M2T):

   In M2T-I&OE auctions, alpha MLE $\alpha^{e,m2t-i\&oe} < \alpha^{i,m2t-i\&oe}$.

   In M2T-I&SE auctions, alpha MLE $\alpha^{s,m2t-i\&se} < \alpha^{i,m2t-i\&se}$.

   In M2T-OE&SE auctions, alpha MLE $\alpha^{e,m2t-oe\&se} < \alpha^{s,m2t-oe\&se}$.

   And alpha MLE $\alpha^{s,m2t-i\&se} < \alpha^{e,m2t-i\&oe}$.

   - When there are three types of bidders in an auction (M3T):

   In M3T Auctions, alpha MLE $\alpha^{e,m3t} < \alpha^{s,m3t} < \alpha^{i,m3t}$, indicating that ordinary-experienced bidders are the most risk-seeking or least risk-averse when all three types of bidders are playing against each other, followed by super-experienced bidders.

2. Expected Seller's Profit Margin:

   - When there is only one type of bidders in an auction (1T):

   $\boldsymbol{E}[\boldsymbol{R}]^{1t-i} < \boldsymbol{E}[\boldsymbol{R}]^{1t-e} < \boldsymbol{E}[\boldsymbol{R}]^{1t-s}$

   - When there are two of the three types of bidders in an auction (M2T):

   For $\boldsymbol{N}^{m2t-i\&oe} = \boldsymbol{N}^{m2t-i\&se} = \boldsymbol{N}^{m2t-oe\&se}$,

   when $\dfrac{\boldsymbol{N}^{e,m2t-i\&oe}}{\boldsymbol{N}^{m2t-i\&oe}} = \dfrac{\boldsymbol{N}^{s,m2t-i\&se}}{\boldsymbol{N}^{m2t-i\&se}}$, $\boldsymbol{E}[\boldsymbol{R}]^{m2t-i\&se} < \boldsymbol{E}[\boldsymbol{R}]^{m2t-i\&oe}$;

when $\dfrac{\boldsymbol{N}^{s,m2t-i\&se}}{\boldsymbol{N}^{m2t-i\&se}} = \dfrac{\boldsymbol{N}^{s,m2t-oe\&se}}{\boldsymbol{N}^{m2t-oe\&se}}$, $\boldsymbol{E}[\boldsymbol{R}]^{m2t-i\&se} < \boldsymbol{E}[\boldsymbol{R}]^{m2t-oe\&se}$;

when $\dfrac{\boldsymbol{N}^{e,m2t-i\&oe}}{\boldsymbol{N}^{m2t-i\&oe}} = \dfrac{\boldsymbol{N}^{s,m2t-oe\&se}}{\boldsymbol{N}^{m2t-oe\&se}}$, $\boldsymbol{E}[\boldsymbol{R}]^{m2t-i\&oe} < \boldsymbol{E}[\boldsymbol{R}]^{m2t-oe\&se}$.

- When there are three types of bidders in an auction (M3T):

If $\boldsymbol{N}^{m2t-i\&oe} = \boldsymbol{N}^{m2t-i\&se} = \boldsymbol{N}^{m3t}$ and $\dfrac{\boldsymbol{N}^{s,m2t-i\&se}}{\boldsymbol{N}^{m2t-i\&se}} = \dfrac{\boldsymbol{N}^{e,m2t-i\&oe}}{\boldsymbol{N}^{m2t-i\&oe}} = \dfrac{\boldsymbol{N}^{e,m3t} + \boldsymbol{N}^{s,m3t}}{\boldsymbol{N}^{m3t}}$

then $\boldsymbol{E}[\boldsymbol{R}]^{m2t-i\&se} < \boldsymbol{E}[\boldsymbol{R}]^{m3t} < \boldsymbol{E}[\boldsymbol{R}]^{m2t-i\&oe}$

The Observed Number of Bidders May Not Equal to N

Another problem we consider is that the number of bidders in our observations are proxy for *N*, which may not always equal to the real *N*. Let us say auctions with at least two bids are informative, and if an auction ends with one bidder, we will not be able to know whether real *N=1 or N>1*. We use conditional probability to solve this problem.

For simplicity, we look at auctions with ***N=2*** that we are able to solve for ***f(t)*** without adding restriction on betas. Since all conditional probabilities sum up to 1, probability of an auction ends with two or more bids equal to the probability that the auction does not end earlier, i.e. $1-\left(1-\beta_1\right)=\beta_1$, so that for auctions with only one type of bidder, we have the conditional probability of an auction ending with two or more bidders

$$f^c(t)=\frac{1}{\beta_1}f(t)=\frac{1}{\beta_1}\prod_{j=1}^{t-1}\beta_j\left(1-\beta_t\right)\ \ \forall\ t\geq 2$$

Similarly, for auctions with two types of bidders, the conditional probability of an auction ending with two or more bidders

$$f^c(t,l_t)=\frac{1}{\frac{1}{2}\left(\beta_1^{k^m}+\beta_1^{k^j}\right)}f(t,l_t)\ \ \forall\ t\geq 2,\ m\neq j, k=\left\{i,oe,se\right\}$$

Thus, taking into consideration of conditional probability, we now compute Maximum likelihood estimates of α by choosing the parameter α to maximise $\sum_j \ln f^c(t_j;\alpha)$.

Table 16 shows examples of Alpha MLE results using unconditional and conditional probability density functions. In the analysis of auctions with only one type of bidder, alpha estimators using conditional probability are slightly higher; in analysis of auctions with two types of bidders, alpha estimators of all bidders groups using conditional probability are all slightly higher; both findings indicate that the conditional probability model recognises bidders in a less risk-seeking or more risk-averse way.

Table 16. Examples of alpha MLE in N=2 model with conditional probability

Auction Item	**Shell $50**		**CVS $25**		**Wal-Mart & Sams Club $25**	
Total Auction Count	282		257		221	
Average Seller Profit Margin	-75.25%		-69.42%		-71.86%	
Treat All Bidder Identical	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**
Alpha Estimate	0.0840	0.0901	0.1390	0.1577	0.1527	0.1748
Auctions with Inexperienced Participants Only	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**
Alpha_i Estimate	0.2638	0.3292	0.2826	0.3594	0.3537	0.4830
Average Seller Profit Margin	-91.10%		-82.86%		-85.63%	
Auctions with Ordinary Experienced Participants Only	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**
Alpha_oe Estimate	0.0497	0.0524	0.0771	0.0855	0.0856	0.0951
Average Seller Profit Margin	-61.70%		-53.37%		-56.50%	
Auctions with Super Experienced Participants Only	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**
Alpha_se Estimate	0.1468	0.1651	0.1356	0.1536	0.1176	0.1321
Average Seller Profit Margin	-85.31%		-68.59%		-65.46%	
Auctions with Inexperienced & all Experienced Participants	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**
Average Seller PM	-75.25%		-69.42%		-71.86%	

Auction Item	Shell $50		CVS $25		Wal-Mart & Sams Club $25	
Alpha_i Estimate	0.1217	0.1426	0.2095	0.2529	0.2362	0.2930
Alpha_e Estimate	0.0808	0.0859	0.1160	0.1295	0.1325	0.1491
Likelihood Ratio Test Statistics	16.3040	20.01	22.6378	29.2844	27.6076	35.4214
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0
Auctions with Inexperienced & Ordinary Experienced Participants Only (Excluding Super Experienced Participants)	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**	**Unconditional**	**Conditional**
Average Seller PM	-72.94%		-69.55%		-74.11%	
Alpha_i Estimate	0.1380	0.1625	0.2120	0.2552	0.2973	0.3829
Alpha_oe Estimate	0.0690	0.0728	0.0995	0.1104	0.1072	0.1199
Likelihood Ratio Test Statistics	23.0022	28.21	22.3774	29.0552	42.3008	54.8306
Likelihood Ratio Test Result	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0	Reject H_0

Another source of measurement error for the real number of bidders is that auctions have real $\boldsymbol{N}>2$ but only two bidders actively bid and auctions are considered as having real $\boldsymbol{N}=2$. I attempted to estimate $\boldsymbol{N}$ together with α by MLE, but the results do not fall in reasonable ranges, indicating some possible computation limitation in this approach.

Discussion

Through applying our two-type-bidder models and three-type-bidder model on observation, we conclude that bidders in different experience groups have different risk attitudes, which affect their bidding strategies and seller's profitability. When separating bidders into two types by prior auction experience of 20, we are able to explain the increasing trend in a seller's average profit margin when the number of experienced bidders increases at a given fixed number of bidders in an auction. When separating bidders into three types by prior auction experience of 20 and 1,000, we notice that the super-experienced bidders compete less aggressively compared to ordinary-experienced bidders when competing in auctions with existence of rivals of different types. This matches the trend that the average seller's profit margin grows in bidders' average prior auctions from the start and gradually drops as the average prior auction becomes very large, indicating there are more super-experienced bidders in an auction when allowing multiple bidder types in models, even if some bidders, e.g. the inexperienced group, are risk-averse, the seller may obtain a profit if the other bidders are risk-seeking.

One limitation of our models based on the real number of bidders is an explanation of the influence of the observed number of bidders on a seller's profitability. The differences between the real and observed number of bidders have not been fully captured in our models. For instance, bidders who do not actively bid are not observed or counted in the number of bidders, and how to estimate the number of these unobserved bidders is to be investigated in future research.

References

Augenblick, N. (2012). Consumer and producer behavior in the market for penny auctions: A theoretical and empirical analysis. *Unpublished manuscript.*

Augenblick, N. (2015). The sunk-cost fallacy in penny auctions. *The Review of Economic Studies*, *83*(1), 58-86.

Byers, J. W., Mitzenmacher, M., & Zervas, G. (2010, June). Information asymmetries in pay-per-bid auctions. In *Proceedings of the 11th ACM conference on Electronic commerce* (pp. 1-12). ACM.

Eyster, E. (2002). Rationalizing the past: A taste for consistency. *Nuffield College Mimeograph.*

Goodman, J. (2012). Reputations in bidding fee auctions.

Hinnosaar, T. (2016). Penny auctions. International Journal of Industrial Organization, 48, 59-87.

Jullien, B., & Salanié, B. (2000). Estimating preferences under risk: The case of racetrack bettors. *Journal of Political Economy*, *108*(3), 503-530.

Platt, B. C., Price, J., & Tappen, H. (2013). The role of risk preferences in pay-to-bid auctions. *Management Science*, *59*(9), 2117-2134.

Stix, E. (2012). An empirical study of online penny auctions. *Brown University: Providence, RI.*

Stone, B. (2010). Penny Auction Sites Hurt by Glut of Competitors. *Bloomberg Businessweek.*

Wang, Z., & Xu, M. (2012). Learning and strategic sophistication in games: The case of penny auctions on the internet. *Unpublished manuscript.*

Wang, Z., & Xu, M. (2016). Selling a dollar for more than a dollar? Evidence from online penny auctions. *Information Economics and Policy*, *36*, 53-68.

Xu, M., Li, S., & Yan, J. (2012). ALL–PAY AUCTIONS WITH A BUY–PRICE OPTION. *Economic Inquiry.*

Zimmerman, A. (2011). Penny Auctions Draw Bidders with Bargains, Suspense. *The Wall Street Journal.*

Zheng, H., Goh, K. Y., & Huang, K. W. (2011). Evaluating the Impacts of Auction Bidding Restrictions on Consumer Surplus and Behaviors—An Empirical Study of Penny Auctions.

Appendix 1. Data Collection & Empirical Facts

Data were collected by Ruby script and stored in the database with the following fields

Table 1. Database structure

Field	Description	Example
Auction ID	A number ID assigned by the auctioneer in ascending order	395881
Auction Item	The item being auctioned	Papa John's $10 Gift Card
Auction Ending Time	Date and time that the auction ended at	9/06/2011 4:58 PM
RRP	Recommended Retail Price, provided by the auctioneer	$10.00
Winning Price	The price that the auction ended at	$0.13
Winner	Username of the winning bidder	loviebug
Other Bidders	Username of the all bidders that placed one or more bids in the auction	Bidding4winning, MaylwinPlease, house9er
Number of Bids Placed by Each Bidder	Total number of bids placed by each bidder is observed, with no information on who placed which bid, except the winning bid	6 bids placed by loviebug, 5 bids placed by Bidding4winning, 1 bid placed by MaylwinPlease, 1 bid placed by house9er

Table 2. Key facts of data

Total Number of Auctions	**572,400**
Total Number of Bidders	**253,464**
Total Bids Placed	**60,795,092**
Total Unique Items	**1,743**
Total Seller Revenue	**50,635,419**
Seller Profit Margin by RRP	**53.38%**

Table 3. Top 10 Most Repeatedly Listed Item Types of Each Category

Item Type	RRP	Total Auctions Count	Average Profit Margin	Total Profit in Dollars
Category - Bidpack (123 types)		287,372	98.64%	$13,277,806.65
50 Bids Bid on any Auction	$37.50	30,619	113.68%	$1,305,287.97
10 Bids Bid on any Auction	$7.50	24,297	56.15%	$102,320.74
25 Bids Bid on any Auction	$18.75	23,495	95.53%	$420,839.50
75 Bids Bid on any Auction	$56.25	23,402	103.86%	$1,367,174.09
15 Bids Bid on any Auction	$11.25	23,286	40.87%	$107,066.12
30 Bids Bid on any Auction	$22.50	22,887	77.76%	$400,430.95
20 Bids Bid on any Auction	$15.00	19,317	47.65%	$138,068.26
100 Bids Bid on any Auction	$75.00	18,719	118.71%	$1,666,599.37
20 Bids Quickfire Auction	$15.00	9,313	104.70%	$146,260.67
15 Bids Quickfire Auction	$11.25	8,137	81.22%	$74,349.80

Item Type	RRP	Total Auctions Count	Average Profit Margin	Total Profit in Dollars
Category - Gift Card (172 types)		243,619	36.0908%	$13,277,806.65
Shell $50 Gift Card	$50.00	9,100	45.52%	$207,116.00
Best Buy $50 Gift Card	$50.00	9,084	22.50%	$102,195.00
Wal-Mart & Sams Club $100 Gift Card	$100.00	8,481	66.79%	$566,445.99
Target $50 Gift Card	$50.00	7,719	50.57%	$195,174.92
The Home Depot $50 Gift Card	$50.00	7,229	27.29%	$98,639.71
$50 Gift Card Red Lobster/Olive Garden	$50.00	7,145	36.27%	$129,574.58
Kohl's $50 Gift Card	$50.00	6,918	15.92%	$55,067.28
Visa $50 Gift Card	$50.00	6,671	104.74%	$349,360.27
Macy's $50 Gift Card	$50.00	6,420	5.26%	$16,884.60
Bed Bath & Beyond $50 Gift Card	$50.00	6,160	7.20%	$22,176.00

Table 3. Top 10 Most Repeatedly Listed Item Types of Each Category - continued

Item Type	RRP	Total Auctions Count	Average Profit Margin	Total Profit in Dollars
Category - Small Goods Under $100 (654 types)		23,660	-38.2136%	-$564,097.90
Keurig Mini Plus Single Cup Brewer	$99.95	472	-20.22%	-$9,539.07
Polaroid POGO Instant Mobile Printer	$49.99	323	-46.55%	-$7,516.32
SanDisk Ultra Backup 16GB Flash Drive	$59.99	314	-49.79%	-$9,378.87
Apple TV 1080p (3rd Generation)	$99.00	311	30.12%	$9,273.65
Sony Clock Radio for iPod/iPhone/ MP3	$69.99	300	-55.31%	-$11,613.44
Apple TV	$99.00	293	-28.52%	-$8,272.80
Apple iPod shuffle 2GB Silver	$49.00	286	-22.80%	-$3,195.19
Cuisinart Centro Griddle / Grill	$99.99	264	-15.98%	-$4,218.30
Apple 2GB iPod Shuffle (Choice of Blue or Slate)	$49.99	255	-15.18%	-$1,935.06
Leap Frog Leapster2 System	$69.99	228	-40.99%	-$6,541.07

Item Type	RRP	Total Auctions Count	Average Profit Margin	Total Profit in Dollars
Category - General Items (692 types)		18,799	-0.6618%	-$31,202.62
iPod touch 8GB 4th Generation	$199.00	719	86.11%	$123,207.05
Kindle Fire w/ Wi-Fi	$199.00	445	146.02%	$129,308.01
Nintendo 3DS 3-D Game Console	$169.99	309	42.64%	$22,397.47
Nintendo DSi Black	$169.99	297	12.75%	$6,437.10
Nintendo DSi XL Midnight Blue	$169.99	275	0.50%	$233.74
Kindle Fire HD 16 GB Tablet	$199.00	239	129.59%	$61,634.30
CHI Ceramic 1\ Hairstyling Iron"	$189.99	222	-51.16%	-$21,578.15
KitchenAid 5-Qt. Stand Mixer	$375.00	219	-20.31%	-$16,679.59
Beats by Dr. Dre Solo Headphones	$199.95	215	-9.69%	-$4,165.66
Weber Wood Burning Outdoor Fireplace	$159.99	204	-36.02%	-$11,756.19

Table 3. Top 10 Most Repeatedly Listed Item Types of Each Category - continued

Item Type	RRP	Total Auctions Count	Average Profit Margin	Total Profit in Dollars
Category - Expensive Consumer Electronics (102 types)		2,742	119.192%	$1,910,711.33
Apple iPad 2 16GB WiFi + AT&T 3G	$529.99	312	209.45%	$346,339.99
Apple iPhone 4S 16GB (Unlocked)	$549.00	256	176.52%	$248,088.27
Apple iPad 2 16GB (WiFi)	$399.00	250	261.28%	$260,626.80
Apple iPad 32GB (WiFi + 3G)	$649.99	177	89.38%	$102,830.11
Xbox 360 250GB Kinect Bundle	$399.99	126	30.01%	$15,124.66
Apple iPhone 5 16GB Unlocked	$649.00	119	222.66%	$171,962.54
Nintendo Wii Mega Bundle	$449.99	105	3.11%	$1,469.44
Canon EOS Rebel XS Digital SLR Camera	$599.99	88	108.73%	$57,408.48
Apple 11-inch MacBook Air	$999.00	80	65.16%	$52,075.87
Nikon D3100 Digital SLR Camera	$649.95	75	1.25%	$609.33

Table 4. Distribution of Bidders with Different Prior Auctions Experiences

Total Auctions Participated	Ratio in Bidder Population	Ratio in Total Bids Placed
< 20	86.91%	18.64%
20+	13.09%	81.36%

Table 5. Distribution of Bidders with Different Prior Auctions Experience

Total Auctions Participated	Numbers of Bidders	Ratio in Bidder Population	Average Total Auctions Participated by Each Bidder	Average Number of Auctions Won	Average Total Bids Placed by Each Bidder	Average Number of Bids Placed per Auction	Winning Percentage
1	30,016	11.76%	1.0	0.01	17.1	17.14	1.06%
2	23,309	9.13%	2.0	0.02	27.2	13.61	1.23%
<7	134,103	52.53%	3.2	0.05	32.4	10.08	1.56%
<25	230,761	90.39%	6.9	0.20	55.7	8.12	2.89%
<200	252,696	98.98%	10.7	0.53	96.8	9.02	4.96%
200+	2,607	1.02%	919.4	169.35	13,937.1	15.16	18.42%
Total	255,304	100.00%	20.0	2.26	238.1	11.90	11.28%

Table 6. Distribution of Bidders in Different Prior Auctions Experience Ranges

Total Auctions Participated	Numbers of Bidders	Ratio in Bidder Population	Average Total Auctions Participated by Each Bidder	Average Number of Auctions Won	Average Total Bids Placed by Each Bidder	Average Number of Bids Placed per Auction	Winning Percentage
<5	97,514	38.20%	2.4	0.03	27.7	11.72	1.39%
[5,7)	36,589	14.33%	5.5	0.10	44.8	8.19	1.76%
[7,10)	38,174	14.95%	7.9	0.17	57.2	7.26	2.17%
[10,20)	49,600	19.43%	13.3	0.46	96.9	7.30	3.48%
[20,25)	8,884	3.48%	21.8	1.07	172.4	7.91	4.92%
[25,30)	5,326	2.09%	26.8	1.42	212.7	7.93	5.29%
[30,50)	9,428	3.69%	37.5	2.31	320.8	8.57	6.18%
[50,100)	5,083	1.99%	67.7	5.37	706.0	10.43	7.94%
[100,200)	2,098	0.82%	137.2	15.30	1,836.0	13.38	11.15%
[200,500)	1,443	0.57%	308.2	43.26	4,811.3	15.61	14.04%
[500,750)	378	0.15%	603.4	98.23	9,828.9	16.29	16.28%
[750,1000)	587	0.23%	696.4	119.70	11,342.1	16.29	17.19%
[1000,2000)	331	0.13%	1,357.1	283.67	25,407.8	18.72	20.90%
[2000,4000)	147	0.06%	2,789.3	527.69	43,438.8	15.57	18.92%
[4000,5000)	31	0.01%	4,414.2	905.03	57,622.1	13.05	20.50%
[5000,10000)	53	0.02%	6,724.6	1,369.66	88,487.6	13.16	20.37%
10000+	15	0.01%	12,722.2	2,445.93	97,466.8	7.66	19.23%
Total	255,304	100.00%	20.0	2.26	238.1	11.90	11.28%

Table 7. Top 100 Most Experienced Bidders by Number of Auctions Participated

Username	Total Auctions Participated	Total Bids Placed	Total Auctions Won	Winning Probability	Average Bid Cost Estimated
Tuffenough	16,864	128,946	2,844	16.86%	$0.43
BETTERQUIT	16,751	142,696	3,876	23.14%	$0.37
DONT_WASTE_YOUR_MON	16,278	73,891	2,236	13.74%	$0.32
URBIDSDONE	14,857	79,043	2,805	18.88%	$0.20
ToughButFair	13,771	61,065	2,860	20.77%	$0.04
YouWillWasteUrBids	12,327	69,986	1,811	14.69%	$0.13
recovery9000	11,734	80,932	2,862	24.39%	$0.40
bruizer	11,630	196,005	2,275	19.56%	$0.45
SaveYours_Not_Stopping	11,523	50,302	2,282	19.80%	$0.22
JustGoAway	11,302	74,508	3,526	31.20%	$0.13
burnboy	11,144	164,131	1,973	17.70%	$0.49
ronishida	10,867	78,602	1,946	17.91%	$0.47
BACKOUTNOW	10,848	117,379	1,643	15.15%	$0.48
bunky123	10,630	72,769	1,597	15.02%	$0.27
InfiniteBidding	10,307	71,747	2,153	20.89%	$0.42
ferangi	9,640	214,223	1,930	20.02%	$0.57
I_WONT_EVER_QUIT	9,215	99,649	1,817	19.72%	$0.38
tonyl587	9,086	80,349	988	10.87%	$0.63
GoAhead_ThrowUr_BidAway	8,892	122,769	2,012	22.63%	$0.45
Pounder	8,731	53,882	1,735	19.87%	$0.39
mirbel	8,579	148,807	934	10.89%	$0.58
LETS-DANCE	8,534	115,754	2,168	25.40%	$0.50
You_Are_Gonna_Regret_It	8,411	67,999	1,211	14.40%	$0.54
NotStopping	8,223	88,843	2,842	34.56%	$0.40
wackpack14	8,150	220,928	2,588	31.75%	$0.51
A-NONY-MOUSE	7,872	78,335	489	6.21%	$0.64
jetranger	7,736	137,113	2,245	29.02%	$0.46
tinabrit	7,691	100,081	1,785	23.21%	$0.59
camort	7,630	191,692	1,506	19.74%	$0.57
bjwhite3114	7,416	42,985	2,242	30.23%	$0.27

NO_WIN_FOR_YOU_HERE	7,357	62,976	1,172	15.93%	$0.47
vita4ever	7,323	57,952	1,505	20.55%	$0.49
BID_4EVER	7,293	37,586	1,063	14.58%	$0.30
xavierMommy	7,241	67,155	1,386	19.14%	$0.56
iSTINKatLOSING	7,101	127,051	1,303	18.35%	$0.44
bidforever11	6,946	110,980	1,573	22.65%	$0.45
Stone_Cold	6,937	70,874	1,420	20.47%	$0.36
puzzledrex	6,836	68,978	1,586	23.20%	$0.49
collegeboy	6,772	236,845	2,113	31.20%	$0.55
SketcheeWon	6,614	93,782	1,183	17.89%	$0.58
donnav7777	6,588	46,540	992	15.06%	$0.38
BID_BLACK_HOLE	6,433	105,961	890	13.83%	$0.60
ItchyFingers	6,416	76,099	1,045	16.29%	$0.47
Craigdcole	6,325	64,392	1,220	19.29%	$0.58
shaveubear	6,224	29,325	322	5.17%	$0.62
Kkatgironde	6,223	97,472	684	10.99%	$0.63
insane8	6,128	49,636	1,489	24.30%	$0.26
wonteverstop1	6,108	85,703	1,028	16.83%	$0.49
BiddingManiac	6,056	70,516	1,282	21.17%	$0.41
DontWasteYerBids	5,981	69,406	1,385	23.16%	$0.45
a8u2g0u6s0t	5,922	34,090	1,057	17.85%	$0.36
NeuroSurgeon	5,917	31,070	1,274	21.53%	$0.25
62loghome	5,776	135,715	1,107	19.17%	$0.63
danado	5,572	58,869	1,083	19.44%	$0.51
gottaget1	5,556	46,134	1,362	24.51%	$0.44
bre7857	5,552	194,931	379	6.83%	$0.69
molder	5,533	36,399	1,410	25.48%	$0.28
jbeans	5,517	44,928	503	9.12%	$0.51
glassmanz	5,462	43,289	1,078	19.74%	$0.44
nyledger	5,424	60,510	1,743	32.14%	$0.46
GoldTrigger	5,338	148,801	1,426	26.71%	$0.53
rollthedice	5,335	156,138	1,364	25.57%	$0.47
BIDiot	5,295	68,658	1,376	25.99%	$0.43
UBETTERQUIT	5,206	38,251	2,008	38.57%	$0.19
BIDCRAZY1964	5,163	37,087	1,226	23.75%	$0.33

nanaalcorn	5,083	58,304	1,011	19.89%	$0.51
UnknownMember	5,042	72,212	1,130	22.41%	$0.51
ONCEIBIDIWILLNOTSTOP	5,002	31,820	922	18.43%	$0.45
zacksdad	4,929	24,413	689	13.98%	$0.33
ItsGoingToBeExpensive	4,887	21,543	1,108	22.67%	$0.20
wontstopiforever	4,874	77,251	713	14.63%	$0.56
BidRidder	4,790	16,462	1,001	20.90%	$0.03
BETTERQUITNOW	4,778	41,479	1,010	21.14%	$0.29
PummelWeed	4,713	60,264	1,441	30.58%	$0.42
BIDSTOTHEMAX	4,674	37,164	1,321	28.26%	$0.28
HereComesTrouble	4,644	21,058	375	8.07%	$0.59
chazown	4,639	154,057	1,412	30.44%	$0.63
vegasgirl926	4,630	75,894	1,482	32.01%	$0.52
BidSomeWhereElse	4,604	30,103	605	13.14%	$0.52
KUSHKUSH	4,588	118,262	1,065	23.21%	$0.57
peacockrancher	4,588	38,108	402	8.76%	$0.51
torilee	4,397	115,444	854	19.42%	$0.59
janets	4,369	44,757	371	8.49%	$0.64
Flinko	4,358	43,444	1,197	27.47%	$0.56
JDONTCARE	4,310	64,407	1,240	28.77%	$0.44
bluemoons	4,303	93,334	588	13.66%	$0.65
suze46	4,297	94,661	219	5.10%	$0.70
bigcuz	4,290	41,518	1,128	26.29%	$0.67
WILLBANKRUPTU	4,235	29,535	626	14.78%	$0.24
theendisnear	4,192	24,612	971	23.16%	$0.11
HALJR16	4,180	13,580	670	16.03%	$0.46
DudeGetOffMyItem	4,166	48,479	847	20.33%	$0.48
Buffy1	4,149	85,132	1,267	30.54%	$0.51
CS110	4,107	29,898	1,393	33.92%	$0.11
huntelkhard	4,075	66,463	992	24.34%	$0.60
ciderguy1970	4,029	21,856	921	22.86%	$0.24
i wanna win	4,020	203,877	533	13.26%	$0.71
FURY_OF_BIDS	4,013	23,466	692	17.24%	$0.04
LIMIT_LESS	4,013	25,764	923	23.00%	$0.30
pinky860	3,907	52,739	114	2.92%	$0.72

Figure 1. Full image of Figure 6 in Chapter 2

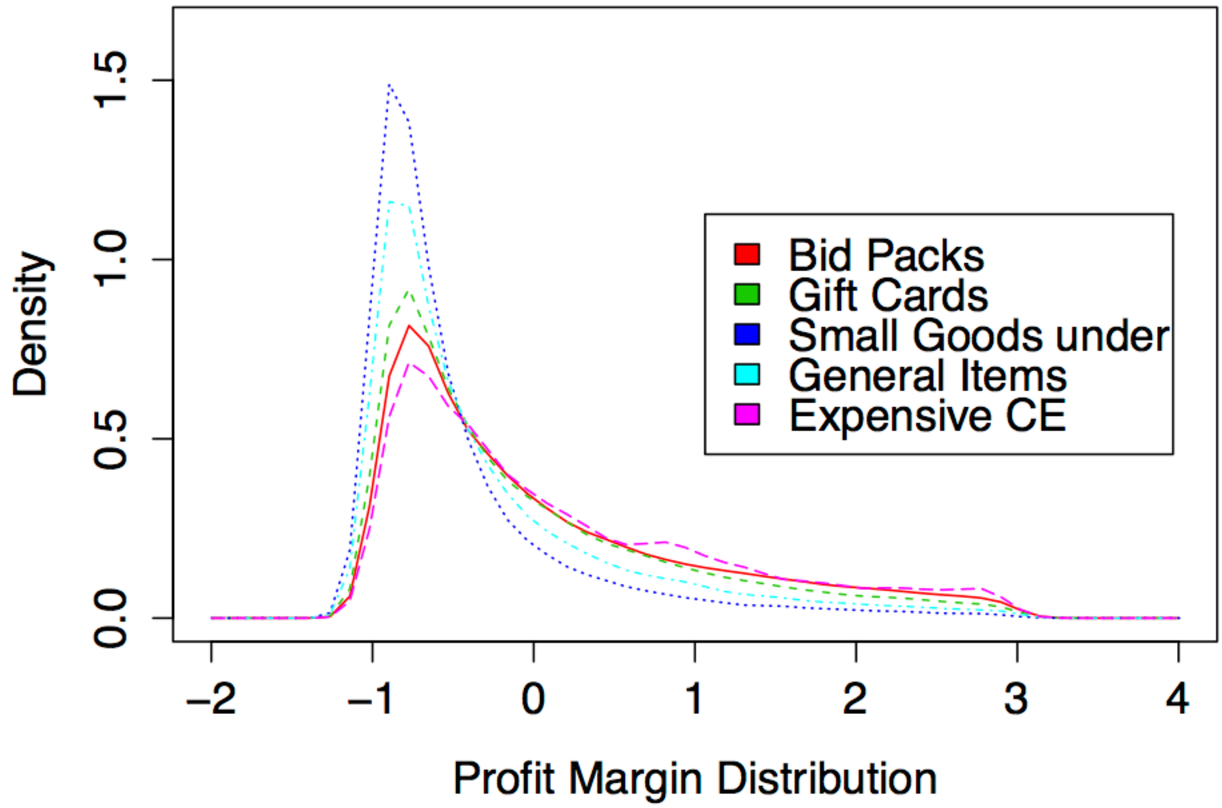

Figure 2. Full image of Figure 10 in Chapter 2

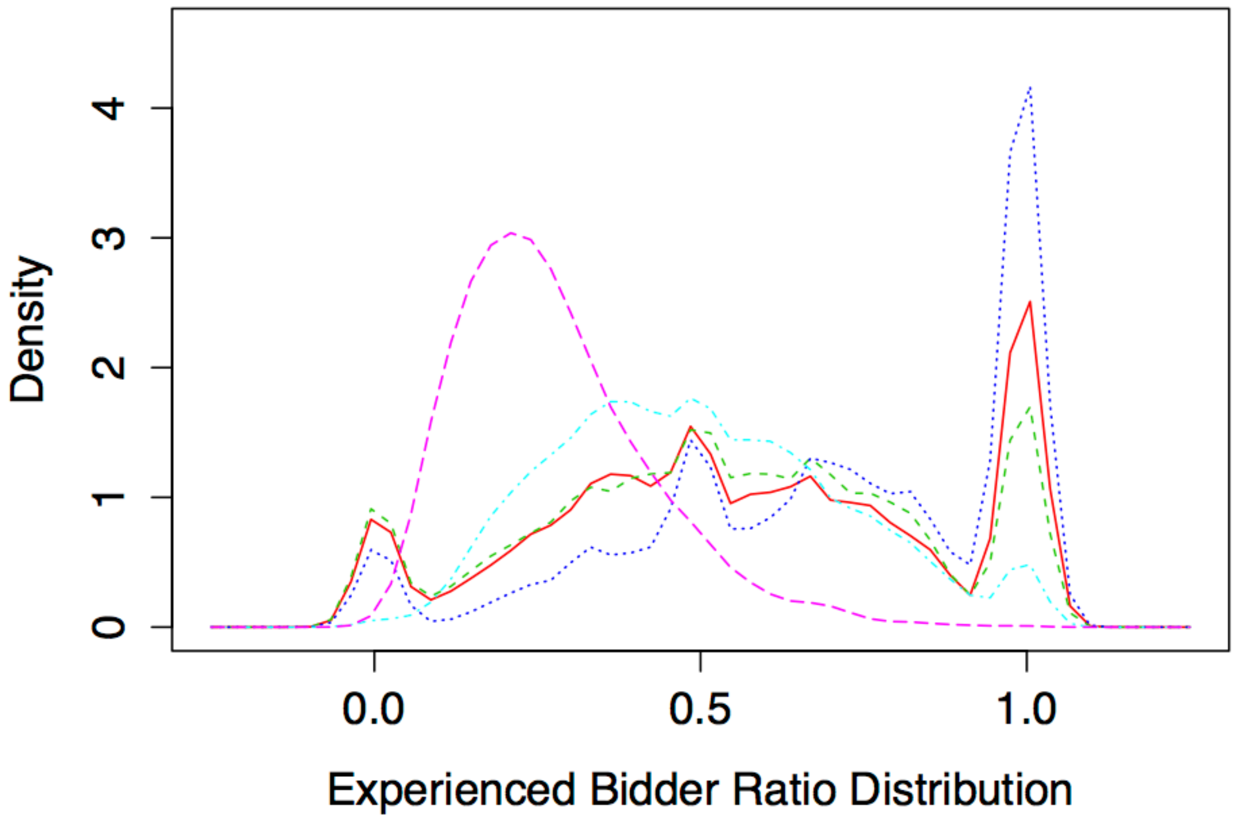

Appendix 2. Results of Alpha MLE

These 42 items are chosen out of all 653 items that have been repeatedly auctioned over 30 times, for their extreme average seller profit margins (highest tail, lowest tail, and those close to zero).

Table 8. Alpha MLE in Single-Bidder-Type model

Item Name	RRP	Average Winning Bid	Average Seller Profit Margin	Alpha MLE	Standard Error	Observation Count
Boston Market $10 Gift Card	$10	$0.07	-46.98%	0.1243	0.0141	341
Barnes & Noble $10 Gift Card	$10	$0.09	-29.15%	0.0570	0.0105	487
Papa John's $10 Gift Card	$10	$0.09	-32.58%	0.0835	0.0162	224
Panera Bread $10 Gift Card	$10	$0.09	-29.02%	0.0567	0.0114	409
Cold Stone Creamery $10 Gift Card	$10	$0.10	-24.97%	0.0446	0.0173	172
iTunes $15 Gift Card	$15	$0.12	-38.1%	0.0628	0.0065	620
Starbucks $15 Gift Card	$15	$0.12	-36.91%	0.0597	0.0097	277
Payless $20 Gift Card	$20	$0.21	-19.14%	0.0185	0.0035	976
Blockbuster Video $25 Gift Card	$25	$0.10	-68.82%	0.1171	0.0227	39
Marshalls $25 Gift Card	$25	$0.34	2.46%	-0.0045	0.0027	907
Cheesecake Factory $25 Gift Card	$25	$0.35	5.79%	-0.0071	0.0026	947
L.L.Bean $25 Gift Card	$25	$0.24	-28.36%	0.0262	0.0029	992
Miracle Blade 11-pc. Cutlery Set	$30	$0.38	-3.36%	0.0006	0.0113	37
Walmart / Sams Club $30 Gift Card	$30	$0.40	0.83%	-0.0023	0.0029	563

Item Name	RRP	Average Winning Bid	Average Seller Profit Margin	Alpha MLE	Standard Error	Observation Count
Altec Lansing Portable Speaker	$40	$0.16	-69.73%	0.0775	0.0154	35
Iron Man 2 (Blu-ray + DVD)	$40	$0.17	-68.63%	0.0743	0.0144	38
Diamond Mini Rockers Mobile Speake	$40	$0.14	-74.18%	0.0922	0.0151	46
True Grit (Blu-ray + DVD)	$40	$0.52	-1.42%	-0.0002	0.0070	54
\Up\" (Blu-ray + DVD)"	$46	$0.16	-73.51%	0.0785	0.0083	110
Foot Locker $50 Gift Card	$50	$0.64	-2.84%	0.0006	0.0026	246
Avatar (Blu-ray) Collector's Edition	$55	$0.23	-68.39%	0.0541	0.0081	63
Logitech Z305 Laptop Speaker	$60	$0.21	-73.19%	0.0598	0.0056	144
Altec Lansing Expressionist Speakers	$80	$0.18	-82.74%	0.0712	0.0105	48
Diamond Accent Heart Earrings	$80	$0.22	-79.44%	0.0597	0.0104	36
WORX 3-in-1 Blower/Vac/Mulcher	$90	$1.20	1.12%	-0.0004	0.0024	85
Deluxe 12-inch diameter Calpha	$160	$0.66	-68.76%	0.0301	0.0053	35
Nintendo DS Lite w/bonus 50 Bi	$168	$6.72	204.95%	-0.0020	0.0019	36
Nintendo DSi XL Midnight Blue	$170	$2.25	0.5%	-0.0001	0.0007	275
Nintendo Wii Console Plus Sports Pac	$250	$12.61	283.34%	-0.0017	0.0012	39
Royal Caribbean $250 Gift Certificate	$250	$3.25	-1.07%	0.0001	0.0008	103
Kindle DX Wireless Reader	$379	$5.08	1.82%	-0.0001	0.0008	45
Apple iPad 2 16GB (WiFi)	$399	$18.97	261.28%	-0.0059	0.0002	250
*Acer Aspire One Netbook	$400	$26.67	406.79%	0.0014	0.0010	34
Apple iPad 2 16GB WiFi + AT&T 3G	$530	$21.58	209.45%	-0.0039	0.0002	312

Item Name	**RRP**	**Average Winning Bid**	**Average Seller Profit Margin**	**Alpha MLE**	**Standard Error**	**Observatio n Count**
Nikon D3100 Digital SLR Camera	$650	$8.66	1.25%	-0.0000	0.0004	75
10 Bidpack + $10 Boston Market Car	$18	$0.12	-46.52%	0.0754	0.0061	587
75 Bidpack+ Breast Cancer Fund	$56	$2.41	225.51%	-0.0391	0.0021	174
75 Bidpack + Food Bank Donation	$56	$2.32	213.78%	-0.0378	0.0026	116
Nintendo DS Lite & 25 Bidpack!	$149	$6.93	254.28%	-0.0011	0.0024	30
50 Bidpack + Nintendo DSi	$207	$2.64	-3.29%	0.0003	0.0017	34
375 Bids Bid on any Auction	$281	$19.67	431.64%	-0.0024	0.0002	806
500 Bids Bid on any Auction	$375	$14.97	203.35%	-0.0036	0.0002	548

12 of the above items have over 30 auctions participated by inexperienced bidders or experienced bidder only and MLE of alpha of each subsets of auctions are estimated.

Table 9. Alpha MLE in subsets of auctions of different types of participants

Item	**RRP**	**Average Seller Profit Margin**	**Alpha MLE (All Auctions)**	**Alpha MLE (Inexperienced Bidder Only Auctions)**	**Alpha MLE (Experienced Bidder Only Auctions)**	**Total Auction Count**	**IE Bidder Only Auction Count**	**EX Bidder Only Auction Count**
Boston Market $10 Gift Card	$10	-46.98%	0.1243	0.6329	0.1441	341	92	168
10 Bidpack + $10 Boston Market Card	$18	-46.52%	0.0754	0.3520	0.0673	587	156	233
iTunes $15 Gift Card	$15	-38.10%	0.0628	0.3199	0.0616	620	114	281
Starbucks $15 Gift Card	$15	-36.91%	0.0597	0.2108	0.1046	277	58	97
Papa John's $10 Gift Card	$10	-32.58%	0.0835	0.5313	0.0217	224	73	97
Barnes & Noble $10 Gift Card	$10	-29.15%	0.0570	0.4246	0.0128	487	143	223
Panera Bread $10 Gift Card	$10	-29.02%	0.0567	0.4851	0.0012	409	138	174
L.L.Bean $25 Gift Card	$25	-28.36%	0.0262	0.3306	0.0348	992	47	613
Cold Stone Creamery $10 Gift Card	$10	-24.97%	0.0446	0.5279	0.1139	172	47	74
Payless $20 Gift Card	$20	-19.14%	0.0185	0.2404	0.0397	976	76	527
Marshalls $25 Gift Card	$25	2.46%	-0.0045	0.2138	0.0239	907	32	401
Cheesecake Factory $25 Gift Card	$25	5.79%	-0.0071	0.1878	0.0062	947	35	469

Appendix 3. Empirical Motivation

Table 9. Distribution of number of bidders in an auction in all observations

Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations
1	14,250	**31**	970	**61**	123
2	37,018	**32**	863	**62**	95
3	53,712	**33**	737	**63**	84
4	61,687	**34**	648	**64**	91
5	61,813	**35**	615	**65**	72
6	57,474	**36**	565	**66**	86
7	49,973	**37**	481	**67**	74
8	41,612	**38**	457	**68**	79
9	33,612	**39**	430	**69**	68
10	27,334	**40**	433	**70**	66
11	22,152	**41**	381	**71**	60
12	17,541	**42**	371	**72**	56
13	14,362	**43**	331	**73**	67
14	11,734	**44**	309	**74**	67
15	9,586	**45**	246	**75**	47
16	7,923	**46**	295	**76**	40
17	6,568	**47**	224	**77**	50
18	5,588	**48**	212	**78**	52
19	4,725	**49**	197	**79**	43
20	3,880	**50**	189	**80**	46
21	3,416	**51**	179	**81**	46
22	2,961	**52**	172	**82**	54
23	2,566	**53**	151	**83**	36
24	2,246	**54**	154	**84**	38
25	1,952	**55**	146	**85**	47
26	1,743	**56**	122	**86**	38
27	1,574	**57**	115	**87**	29
28	1,266	**58**	127	**88**	35
29	1,249	**59**	103	**89**	25
30	1,021	**60**	106	**90**	35

Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations
91	32	**121**	19	**151**	6
92	32	**122**	17	**152**	5
93	16	**123**	17	**153**	7
94	28	**124**	15	**154**	6
95	25	**125**	14	**155**	9
96	25	**126**	27	**156**	4
97	25	**127**	8	**157**	10
98	18	**128**	6	**158**	3
99	18	**129**	10	**159**	4
100	20	**130**	5	**160**	8
101	16	**131**	9	**161**	2
102	24	**132**	11	**162**	4
103	27	**133**	8	**163**	7
104	18	**134**	10	**164**	8
105	15	**135**	13	**165**	5
106	17	**136**	9	**166**	5
107	21	**137**	11	**167**	6
108	14	**138**	4	**168**	4
109	25	**139**	10	**169**	3
110	16	**140**	8	**170**	9
111	26	**141**	9	**171**	7
112	16	**142**	7	**172**	4
113	16	**143**	11	**173**	2
114	15	**144**	7	**174**	6
115	15	**145**	6	**175**	3
116	10	**146**	8	**176**	7
117	11	**147**	9	**177**	6
118	23	**148**	3	**178**	6
119	11	**149**	10	**179**	3
120	13	**150**	9	**180**	2

Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations
181	6	**211**	2	**241**	2
182	3	**212**	4	**242**	5
183	4	**213**	5	**243**	1
184	7	**214**	5	**244**	6
186	8	**215**	4	**245**	5
187	5	**216**	6	**246**	2
188	2	**217**	4	**247**	2
189	3	**218**	2	**248**	3
190	5	**219**	6	**249**	4
191	5	**220**	2	**250**	2
192	4	**221**	1	**251**	2
193	4	**222**	4	**252**	1
194	3	**223**	5	**253**	3
195	5	**224**	5	**254**	2
196	3	**225**	5	**255**	1
197	6	**226**	3	**256**	2
198	7	**227**	2	**257**	4
199	4	**228**	4	**258**	2
200	1	**229**	2	**259**	3
201	4	**230**	1	**260**	1
202	1	**231**	4	**261**	5
203	3	**232**	3	**263**	2
204	7	**233**	1	**265**	1
205	4	**234**	4	**266**	4
206	5	**235**	5	**267**	2
207	3	**236**	2	**268**	2
208	2	**237**	4	**269**	1
209	2	**238**	2	**271**	3
210	4	**240**	1	**272**	2

Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations	Number of Bidders in an Auction	Number of Observations
274	1	**291**	3	**326**	1
275	1	**293**	3	**327**	1
276	1	**295**	3	**328**	1
277	4	**296**	2	**338**	1
278	1	**297**	1	**342**	1
279	2	**298**	1	**346**	2
280	3	**307**	1	**349**	1
281	1	**308**	2	**352**	1
282	1	**309**	2	**356**	1
283	1	**310**	1	**360**	1
284	4	**315**	1	**364**	1
286	4	**317**	1	**372**	1
287	1	**318**	1	**375**	1
288	1	**323**	2	**377**	1
289	1	**324**	1		
290	1	**325**	2		

Figure 3. Distribution of number of bidders in an auction (Full Database)

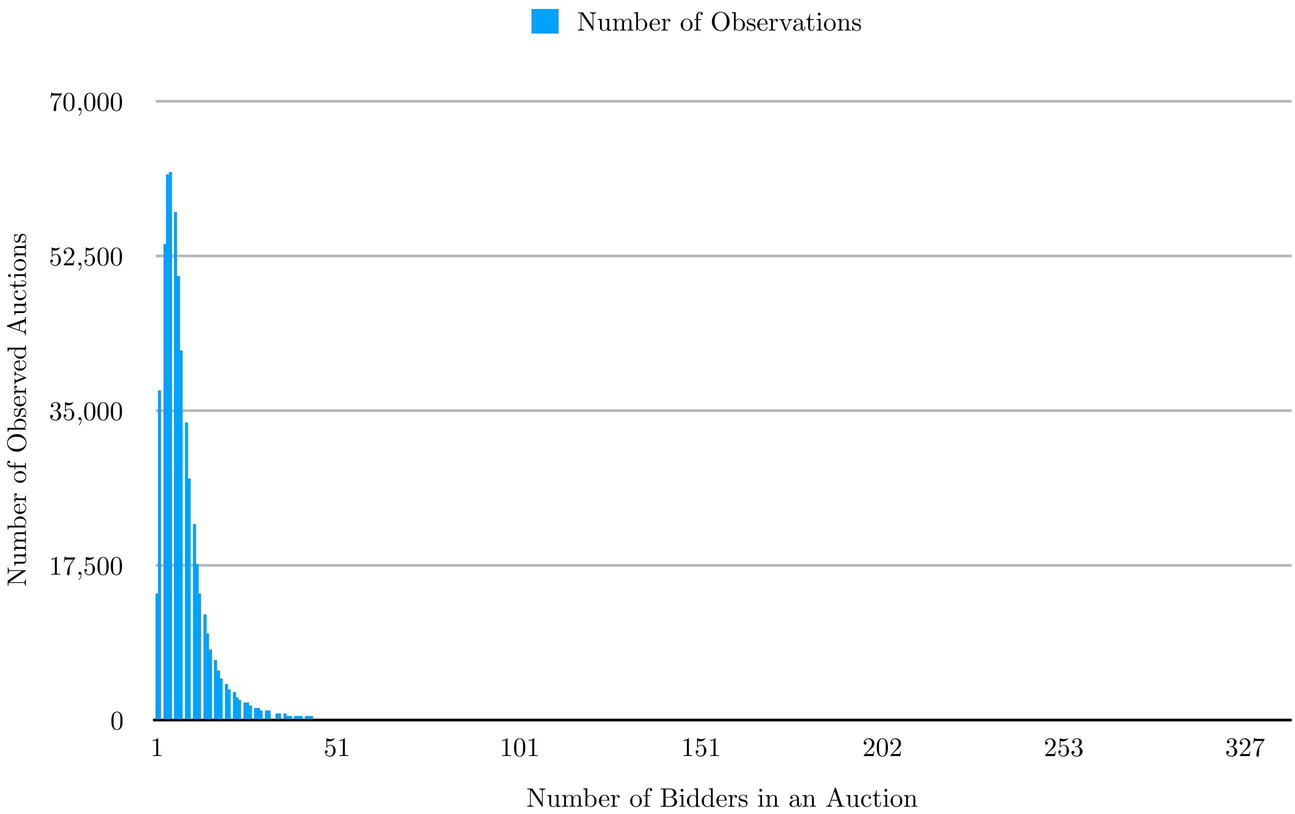

Figure 3. Distribution of number of bidders in an auction (N<50)

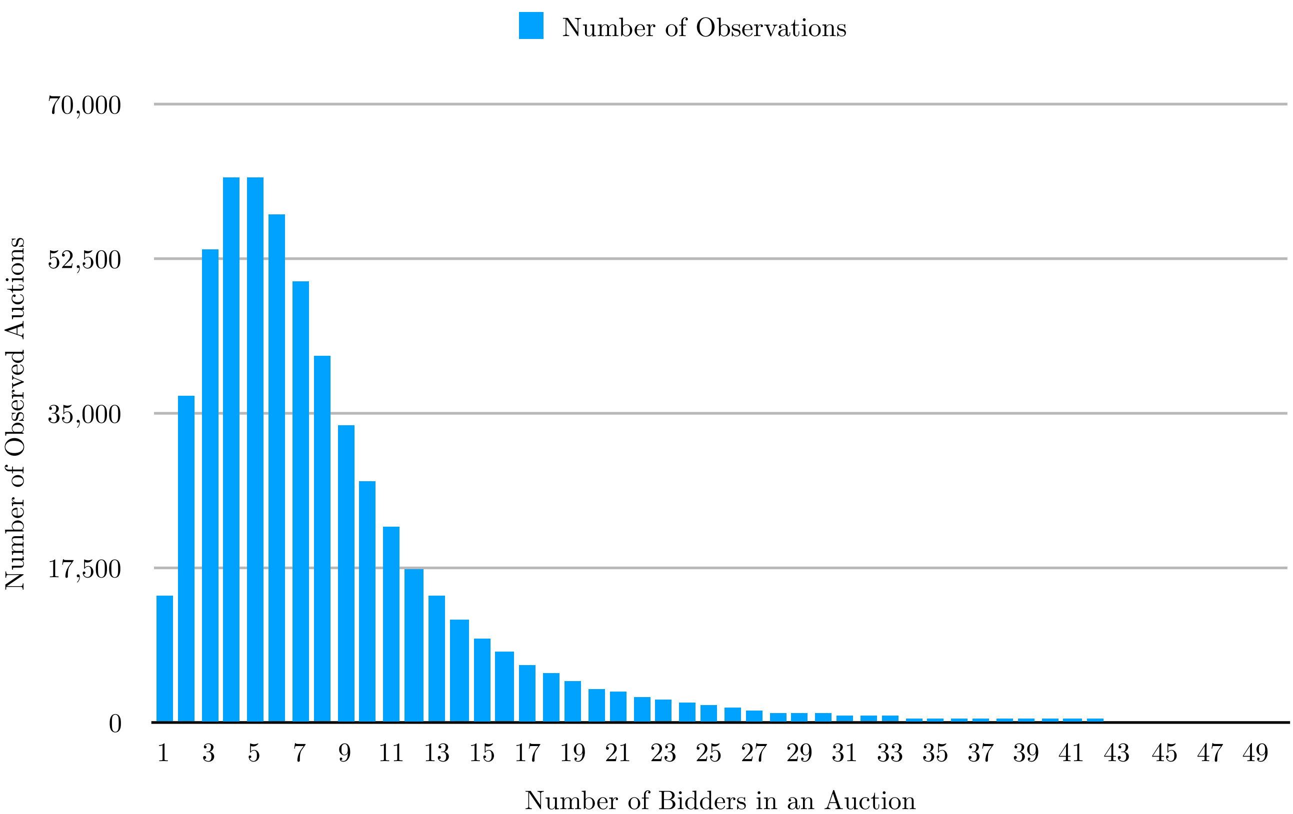

Table 10. Distribution of number of bidders of different experience groups with N up to 24

Number of Bidders in an Auction	All Auctions	Auctions with Inexperienced Bidders Only	Auctions with Experienced Bidders Only	Auctions with Mixed Types of Bidders
1	14,250	4,822	9,428	0
2	37,018	4,317	24,118	8,583
3	53,712	3,115	27,819	22,778
4	61,687	2,017	24,148	35,522
5	61,813	1,304	17,900	42,609
6	57,474	844	12,261	44,369
7	49,973	531	7,792	41,650
8	41,612	366	4,779	36,467
9	33,612	265	2,667	30,680
10	27,334	186	1,457	25,691
11	22,152	99	785	21,268
12	17,541	63	417	17,061
13	14,362	43	225	14,094
14	11,734	32	130	11,572
15	9,586	23	53	9,510
16	7,923	11	35	7,877
17	6,568	11	19	6,538
18	5,588	8	7	5,573
19	4,725	7	1	4,717
20	3,880	2	3	3,875
21	3,416	2	3	3,411
22	2,961	4	1	2,956
23	2,566	1	1	2,564
24	2,246	2	1	2,243