Pay-to-bid Auctions

Brennan C. Platt, Joseph Price, and Henry Tappen*

March 31, 2009

Abstract

We analyze an auction format in which bidders pay a fee each time they increase the auction price. The bidding fees are the primary source of revenue for the seller, but result in an expected revenue equivalent to standard auctions. Our model predicts a particular distribution of ending prices, which we are able to test against observed auction data. Our model fits the data well for over three-fourths of the items which are routinely auctioned. The notable exceptions are items related to video game systems; these result in more aggressive bidding and higher expected revenue. By incorporating mild risk-loving preferences in the model, we can explain nearly all of the auctions.

1 Introduction

The relatively minor setup cost of internet websites has facilitated the creation of many varieties of auction formats. Most of these have close analogs to auctions that have existed for centuries prior, but occasionally a site develops a novel approach. Such is the case with Swoopo, a German company founded in 2005 now operating distinct sites in the United Kingdom, Spain, the United States, and Austria (with

*Department of Economics, Brigham Young University, 149 FOB, Provo, UT 84602, (801) 422-8904, brennan_platt@byu.edu

the last two being introduced in 2008).[1] Swoopo's novel feature is that participants must pay a fee each time they place a bid.

Each auction begins at a price of zero and with a specified amount of time on a countdown clock. A participant places a bid simply by clicking a button. Each submitted bid increments the current price by a fixed amount ($0.15 in the standard auction); hence, participants do not choose how much to bid, only whether to bid at any given moment. When the bid is placed, the bidder is immediately charged a bid fee ($0.75) and the clock is reset (typically to 15 seconds remaining), extending the ending time of the auction.

If the time expires before another bid is placed, the last bidder pays the current price (on top of any bid fees he incurred) and wins the object. This framework posses a fundamental tradeoff for participants: a person who bids early in the auction could win the item for a tiny fraction of the retail price, yet there is a high probability that someone else will bid after him, in which case he has wasted the bid fee. In essence, placing a bid in Swoopo is like betting that no one will bid in the next 15 seconds.

A pay-to-bid auction creates two sources of revenue from the auction; in Swoopo's case, the bid fees are far more important than the final price. The average item's final price is only 31% of the retail value; yet for every dollar earned in final price, Swoopo earns five dollars in bidding fees. On its face, this auction format appears to be highly lucrative to its owners and leads one to question the rationality of the bidders.

We provide a simple model that rationalizes the behavior of bidders. The key insight of this model is that the aggregate probability of being outbid is pinned down by the requirement that prior bidders had to be willing to bid ex-ante. As a consequence, we can derive a probability density function describing the probability that an auction will end at any particular number of bids. We also show that the pay-to-bid auction raises an expected revenue that is nearly equal to the bidders' valuation of the item. One can consider this as a revenue equivalence result, since a

[1] Swoopo is the exclusive seller on their site, rather than allowing other to auction items as on eBay.

large number of risk neutral bidders in a standard auction setting would produce the same expected revenue.

We then test this model against auction data collected from Swoopo's website, obtaining a remarkable fit for many of the routinely auctioned items. In many cases, this fit is achieved by using a valuation distinct from the retail value, which is determined by maximum likelihood estimation. Indeed, for 78% of the routinely auctioned items, we cannot reject the hypothesis (at $p = 0.05$) that the observed distribution comes from our model. The items which routinely do not fit the model appears to be video game systems and related accessories, which produce a distinct bidding pattern and higher revenue for Swoopo.

We assume that the valuation of the item is known and the same across all potential bidders. This assumption is quite plausible for the types of items regularly auctioned on Swoopo. All items are new and unopened; most are readily available from traditional or internet retailers. Indeed, it is more appropriate to consider the valuation to be the lowest price for which the item may readily be obtained elsewhere (rather than as the individual's literal reservation price). As such, the assumption of a common valuation would be less plausible when there is a shortage of the item (such as a new popular holiday gift) or when there is significant price dispersion in the outside market.

In our model, placing a bid is essentially a gamble on whether someone else will bid afterwards. The bidder weighs the expected payoff from winning versus the certain bid cost. If he knew he would always be outbid, the bidder would not waste the bid fee; thus, in equilibrium, those who follow this bidder must play mixed strategies. That, in turn, requires bidders to be indifferent about bidding at each stage. This indifference condition provides the key in determining the probability of any particular bidding outcome and deriving the expected revenue.

The indifference condition pins down the *aggregate* probability of being outbid; the relationship between this probability and individual strategies depends on the particular environment considered. For instance, there could be a fixed pool of bidders watching the auction, each of whom plays a mixed strategy as to whether (and when) to place the next bid. Alternatively, new customers may arrive according to some Poisson process, and on arrival, they use a mixed strategy to decide whether or not to bid. There may even be heterogeneity in the valuation of the object, so there is some probability that a person with a high enough valuation will arrive before the time expires. But in the end, the process that generates the aggregate probability of being outbid is not particularly important, so long as it keeps preceding bidders active.

At first glance, one might consider the pay-to-bid auction as a variation on all-pay auctions. There, every person pays exactly what he or she bid, even though only the highest bidder wins the object. But the strategic properties of the two formats differ greatly. If we were to model an all-pay auction in the same sequential format as the pay-to-bid auction, bidders would face a choice of how long to remain active, with no option to reenter if they drop out. Each time *any* active bidder placed a bid, *every* active bidder would pay the bid fee to remain active. If no one placed an additional bid within the time limit, the auction would close, awarding the prize to the last person who raised the bid. That winner would not be charged anything beyond his bid fees.

The strategic environment of a pay-to-bid auction differs significantly. The price increment after a bid is distinct from the bid fee; hence, the highest bidder pays the final price in addition to his bidding fees. Second, a bidder only incurs a bid fee each time he increases the price, not when *anyone* increases the price.

Pay-to-bid auctions have more in common with previous work on bidding costs. These are typically modeled as exogenous transaction costs that the bidder incurs but are not received as part of the sellers revenues, as in our model. These could be interpreted as legal costs for preparing a bid, or an increasing cost of financing due to capital market imperfections, for instance. In early work such as Che and Gale (1998), only the winner of the auction incurred the bid cost, which depends on the size of the bid. Gavious, Moldovanu, and Sela (2002) analyze an all-pay auction (or *contest*) in which all bidders incur the bid cost.

On one level, the strategic impact of the bidding costs on bidders is similar to bid fees in our model. A person only chooses to bid if his expected payoff is sufficient to cover the certain expense. Since valuation is homogenous in our model, this imposes our indifference condition; in the bidding cost models, the cost will discourage those with low valuations from participating. Also, in our model, the same person may place multiple bids, instead of a single bid as in the bidding cost models. Indeed, when bidding costs are modeled in an ascending auction, participants conserve on bid revisions by making jumps in the bids; here, the auctioneer has denied them that option. Most importantly, bidding costs do not contribute to the seller's revenue as in our model.

The paper proceeds as follows: in section 2, we present the model and identify its unique symmetric equilibrium. We also discuss several variations on the model. In section 3, we discuss the data we have extracted from Swoopo's website. Section 4 presents evidence on the model's fit to the data and section 5 concludes.

2 Model

An item being sold has a known, objective value of v to all potential bidders. This is initially offered at a price of 0. A customer immediately pays b dollars each time he places a bid, which also increments the price by s dollars. Moreover, the auction clock is reset to T minutes. The auction concludes if no one places the next bid within T minutes, resulting in the last bidder paying the current price and receiving the item.

Initially, let us assume that potential bidders' decisions of when (if at all) to act are distributed throughout the T minute interval without atoms; thus ties never occur. This then constitutes a full information, extensive form game. We examine symmetric subgame perfect equilibria.

When a customer contemplates placing a bid, he is essentially placing a bet of b dollars that no one else will bid after him. Consider a customer who observes the current price at $s(q-1)$ dollars. If he were to place the q^{th} bid, his expected payoff would be $(v - s \cdot q)(1 - \mu_{q+1}) - b$, where μ_{q+1} denotes the probability that anyone else will place the $q + 1^{\text{th}}$ bid.

First, note that if $q > Q \equiv \left\lfloor \frac{v-b}{s} \right\rfloor$, then this customer will strictly prefer not to bid, regardless of μ_{q+1}. For the earlier stages of the game, the customer must be indifferent about bidding. If the q^{th} bidder strictly preferred bidding, then $\mu_q = 1$ and therefore the $q-1^{\text{th}}$ bidder would strictly prefer not to bid. This would continue until the first or the second bidder would refuse to bid.

This indifference condition then determines the unique probabilities that define equilibrium behavior. In particular, for all $q > Q$, $\mu_q = 0$, and for all $q \leq Q$, $\mu_q = 1 - \frac{b}{v - s\cdot(q-1)}$.

2.1 Individual behavior: fixed pool

The probabilities μ represent the aggregate consequence of individual mixed strategies. By specifying the process by which bidders participate in the auction, we can analyze these individual decisions and interpret the strategies better. However, the aggregate is sufficient to analyze the outcome of the auction, which we do in the next section.

First, consider an environment in which a fixed pool of n agents are watching the auction. Their individual mixed strategies have two components: a probability of bidding β_q, and a probability density function $g_q(t)$ on $[0, T]$ denoting when to bid. The latter has no restrictions — since agents are indifferent about bidding, there is no particular reason to do so earlier or later. The former, however, will bear direct relation to the aggregate probabilities.

Consider the decision when $q-1$ bids have been placed. The $n-1$ customers who currently are not winning must decide whether to place a bid. With probability $(1-\beta_q)^{n-1}$, none of them do. In order to satisfy the indifference condition for the $q-^{\text{th}}$ bid, this must equal $1-\mu_q$. So $\beta_q = 1 - (1-\mu_q)^{\frac{1}{n-1}}$.

2.2 Individual behavior: exogenous arrival

Suppose instead that bidders arrive to the auction according to a Poisson process, with rate λ per minute. Upon arrival, the bidder instantaneously decides whether to bid (employing mixed strategy β_q). Whether he bids and loses, or declines to bid, he then departs the auction.

We again consider when the $q-1^{\text{th}}$ bid has just been placed. During the next T minutes, the probability that n customers arrive is $\frac{e^{-\lambda T}(\lambda T)^n}{n!}$, and the probability that all of them decline to bid would be:

$$\sum_{n=0}^{\infty} \frac{e^{-\lambda T}(\lambda T)^n}{n!} \cdot (1-\beta_q)^n = e^{-\lambda T \beta_q}$$

In order to satisfy the indifference condition for the $q-1^{\text{th}}$ bid, this must equal $1-\mu_q$; however, it may not be able to do so. In particular, if λ is sufficiently small, for any β_q, we can have $e^{-\lambda T \beta_q} > 1-\mu_q$. Even so, equilibrium still exists, with minor adaptation. Individual mixed strategies would be $\beta_q = \min\left\{-\frac{\ln(1-\mu_q)}{\lambda T}, 1\right\}$, with aggregate probability that someone bids of $1-e^{-\lambda T}$ when the latter case holds. In such an equilibrium, all bidders who arrive when q is low will strictly prefer to bid even knowing that those who follow will do the same, because there is a large enough chance that no one will show up at all. For high enough q, eventually the indifference condition will hold and μ_q becomes the aggregate probability again.

2.3 Expected revenue: fixed price auction

A variant of the pay-to-bid auction keeps the final price of the good fixed at some price p, while $s=0$ effectively. That is, the winner will pay p in addition to his bidding fees. In practice, Swoopo still increments the price after each bid as if $s=0.15$ (which allows us to determine the number of bids which occurred), but ignores sq in computing the final price. A special case of this when $p=0$ is referred to as a 100% off auction.

The advantage of examining this special case is that it produces a clean probability distribution. Since $(v-p)(1-\mu_q)=b$, we derive that $\mu_q = 1-\frac{b}{v-p}$ for all q. From this conditional probability (of the q^{th} bid occurring, given that the $q-1^{\text{th}}$ already has), we can construct the probability $f(q) = \frac{b}{v-p}\left(1-\frac{b}{v-p}\right)^q$ that the bidding ends at exactly q bids.

We then examine the expected revenue:

$$\sum_{q=1}^{\infty}(bq+p)f(q) = v\left(1-\frac{b}{v-p}\right) \tag{1}$$

and variance of revenue:

$$\left(v\left(1-\tfrac{b}{v-p}\right)\right)^2 f(0) + \textstyle\sum_{q=1}^{\infty}\left(bq+p-v\left(1-\tfrac{b}{v-p}\right)\right)^2 f(q) \tag{2}$$

$$= \tfrac{(v-p-b)\left((2v-p)bp+(v-p)^3\right)}{(v-p)^2} \tag{3}$$

Thus we note that this auction format does not extract more revenue than any other. Indeed, a typical 1st or 2nd price sealed-bid auction among identical buyers would raise a revenue of v.

2.4 Expected revenue: Standard pay-to-bid auction

In the standard auction, μ_q varies with q, which significantly complicates the probability density function of ending bids, $f(q)$. In particular, we obtain

$$f(q) \equiv (1-\mu_{q+1})\prod_{j=1}^{q}\mu_j = \frac{b}{v-s\cdot q}\prod_{j=1}^{q}\left(1-\frac{b}{v-s\cdot(j-1)}\right). \tag{4}$$

Unfortunately, this density function does not lend itself to analytical evaluations of expected revenue. However, note that $f(q+1) = \frac{v-b-s\cdot q}{v-s\cdot(q+1)}f(q)$. We can use this fact to construct a continuous approximation of f that is far more tractable. We observe that the slope at $q+\frac{1}{2}$ is approximately $f(q+1)-f(q)$, while $f\left(q+\frac{1}{2}\right)$ is approximately $\frac{f(q+1)+f(q)}{2}$. Thus we obtain:

$$\frac{f'\left(q+\frac{1}{2}\right)}{f\left(q+\frac{1}{2}\right)} \approx \frac{f(q+1)-f(q)}{\frac{f(q+1)+f(q)}{2}} = \frac{2(s-b)}{2v-b-s(1+2q)},$$

or with rearrangement, $f'(q) \approx \frac{2(s-b)}{2v-b-2sq} f(q)$. This differential equation has the solution $f(q) = K(b + 2qs - 2v)^{\frac{b-s}{s}}$. To determine the constant, we impose a boundary condition that

$$\int_0^{\frac{2v-b}{2s}} f(q) = 1,$$

which is to say, we ensure that this is a probability density function over all positive q that have non-negative $f(q)$. The resulting continuous approximation is:

$$f(q) \equiv \frac{2b}{2v-b} \left(\frac{2v - b - 2qs}{2v-b} \right)^{\frac{b-s}{s}}.$$

The expected revenue can then be readily computed as:

$$\int_0^{\frac{2v-b}{2s}} (b+s)qf(q)dq = v - \frac{b}{2} \tag{5}$$

with variance:

$$\int_0^{\frac{2v-b}{2s}} \left((b+s)q - v + \frac{b}{2} \right)^2 f(q)dq = \frac{b(2v-b)^2}{4(b+2s)} \tag{6}$$

When v is large relative to b, the approximation is very accurate. Numerical evaluations of expected revenue using the true distribution always produced revenue of $v - b$.

3 Data

From the inception of their US website in September 2008 through January 2009, Swoopo.com has auctioned over 34 thousand items, all via pay-to-bid auctions. Most of these auctions are repetitions of identical objects, with only 623 unique items. Furthermore, a small handful of items are auctioned much more frequently than the others; only 47 items were auctioned more than 100 times, with the most popular (a Nintendo Wii Sports system) being auctioned 2,728 times.

Swoopo lists all ended auctions on their website. For each auction, the site provides the suggested retail price, the final auction price, the bid fees paid by the winner, and the end time. Also listed is the bid fee and the price increment that occurs with each new bid. Most auctions increase by \$0.15, but penny auctions increase by \$0.01 with each new bid. We divide the final price by the bid increment to get the total number of bids in each auction. We multiply the number of bids by the cost per bid and add this to the amount the winner pays to get the item to get Swoopos final revenue for that auction.

The price the winner pays differs across auctions. The most common auction is where the bidder pays the final auction price (66% of all auctions). The other auctions involve the players paying a predetermined amount for the item (with the number of bidders or final auction price having no influence on this amount). The most common case is a 100% off auction where the winner of the auction pays nothing for the item except for the money already paid to place bids (12% of auctions).

We find that in 54% of the auctions, Swoopo's revenue exceeds the suggested retail price (with an average revenue that is 176% of the retail price). Revenue varies substantially with the type of item being auctioned, though; the majority of revenue comes from auctions of video game systems or accessories, where average revenue is 259% of the retail price. Among standard auctions (where bidders pay the final auction price), the final price is 31% of the suggested retail price on average; the bid fees generate an additional 156% of retail price.

In many cases the suggested retail price is significantly higher than what is readily available from other online retailers. To provide an alternative measure of retail price, we found each item on Amazon.com. By this measure, Swoopo's suggested retail price is about 28% higher than the price on Amazon.com (weighting items by the number of auctions for that item).

Swoopo also provides the usernames of the winner and last 10 bidders of each auction. Across all of the auctions in our sample, there were 15,540,820 bids placed, of which 2% occurs as one of the last ten bids that we can observe. Among the last ten bidders across all auctions, there are 69,519 unique usernames with 13,001 of these users win at least one item. Over 63% of the winners only win once and a very small contingent (less than 1%) have won more than 25 auctions.

A common thought upon observing a pay-to-bid auction is why anyone would be the first bidder. While it is unlikely that the first bidder will win the auction, we do have 134 auctions in which the first bidder wins the auction, 299 where the 2nd bidder does, and 8,667 that are won by one of the first 9 bidders. In fact, two of the surprising predictions of the model are that a large fraction of the auctions will end during the early bids and that the probability of an auction ending at a given bid number is decreasing in the number of previous bids.

Swoopo provides bidders the option to use a bid butler which will place bids for them based on preset parameters. To use the bid butler, the bidder decides at what price to start bidding on an item, at what price to stop bidding on an item, and how many bids they are willing to make. The bid butlers are designed to be as frugal as possible and will wait until there is 10 seconds left on the auction clock to bid. We find that 32% of the last ten bids across auctions are made by bidders using a bid butler but 64% of auctions are won by someone using a bid butler.

Swoopo is not the only website in the US to offer this type of pay-to-bid auction. There are several other sites providing a similar service (bidstick.com, fantasticbid.com, gobid.com, and grababid.com). However, using data from a web traffic monitoring company (Alexa.com) we found that Swoopo.com accounts for 87% of the web traffic going to this group of sites.

4 Evidence

Our primary objective is to determine if our model appropriately depicts the pay-to-bid environment, using Swoopo data as our evidence. We test this hypothesis by comparing the observed distribution of the number of bids to the theoretical predictions produced by our model. As mentioned earlier, the auctions for video game systems behave very differently from the other items and the discussion that follows with separate the results based on the type of item.

We start by providing some illustrative examples in Figure 1, showing the degree to which the distribution of bidding behavior we observe in the swoopo data matches the predictions of our model. In each example, the numbers along the x-axis are the number of bids that occur in an individual auction (which maps one-on-one with the revenue for the auction). The y-axis is the probability (either theoretical or actual) that the auction ends at that number of bids.

The model predicts a distribution of ending bids based on the valuation of the item, the bid fee, and the increment by which the final price rises with each bid. The latter two are clearly specified for each auction. For the valuation, we initially use the retail price reported by Swoopo; however, it is quite possible that bidders may value an item more or less than its retail price indicates.[2] To account for this possibility, we compute a maximum likelihood estimate of the valuation (described in detail in the appendix), and repeat our statistical tests using this estimated valuation.

The top row of figures displays the results from three of the most common non-video game items and the bottom row has three of the most common video game items. The upper figures provide visual evidence that the actual distribution of revenues matches very closely with that predicted by the model. We have a large number of auctions ending with few bids (and hence low revenue), and the probability of the bidding ending at q is declines as q increases.

As a consequence, many of these auctions conclude in a net loss for Swoopo — for items similar to those depicted in the top row, as many as 69% of the auctions of a given item will not generate enough revenue (in bid fees and final price combined) to cover the retail price. However, the long right tail generates enough compensating revenue so that, on average, these auctions raise revenue equal to the valuation. These results also indicate that (for this set of items) bidders may in fact be rational in participating in the pay-to-bid auction, even early on in the auction, because of

[2]For instance, if the same item is available from internet retailers at a discount, bidders would behave as if their true valuation is the discount price. One would expect Swoopo bidders to be sufficiently internet-saavy to have checked Amazon.com or similar sites for a price comparison. On the other hand, if the item were difficult to find through other outlets, such as the Wii video game console, the true valuation could be well above retail.

the relatively high probability of obtaining the item at a price well below its retail price.

Items depicted in lower row of Figure 1 show that auctions for video game systems and accessories do not match the predictions of our model very well. In particular, we find there is an initial range for which the probability is increasing in q, which cannot be generated by this model. In addition, we find that the density of ending bids in the low end is much smaller than predicted, and the right tail of the distribution is much thicker. Using the MLE estimate of the valuation improves the fit of our model but we still find too few auctions ending with low revenue.[3]

We next turn to some statistical tests comparing the match between the prediction of our model and what we observe in the Swoopo data. We restrict our analysis to items that are sold in 45 or more standard auctions. Out of 505 distinct items (sold in the standard pay-to-bid auction), 65 meet this criterion. For each of these items, we perform a Pearsons chi-square goodness of fit test comparing the observed distribution of ending bids to the theoretical distribution. We also perform a t-test comparing the theoretical and actual mean revenue. In Table 1, we report the fraction of items for which the tests fail to reject the null hypothesis that the observed data differs from the prediction of the model (separately based on whether the test statistic corresponds to a p-value greater than .01, .05, and .10).

In Panel A of Table 1, we use the reported retail value of the item as our measure of bidders common valuation for the item. We group items into five categories: computers, computer accessories, home electronics, apparel, and video games. For instance, considering computer-related items and using the .05 p-value as a threshold for our test, we find that when comparing the predicted and actual distribution, the majority of the time we fail to reject that the two distribution are statistically different. The degree to which our model matches the observed distribution of the data is much lower for video games and home electronics were we fail to reject that the two distributions are different only about 25% of the time.

[3] These items that poorly fit are among the most frequently auctioned items. This is not surprising, because in each case the item generates greater revenue than the retail v; they are undoubtedly the most profitable items.

In the far right columns of Table 1 we look specifically at mean revenue. Our model does a worse job at predicting the mean revenue for an item, with us failing to reject that the two means are different 42% of time for computers and even less often for the other items. However, in this panel the valuation of the item is determined by the reported retail price, which may be a poor proxy for the true valuation.

Some suggestive evidence that the retail price may be a poor proxy comes by looking at the apparel category, which includes watches (an item for which the reported retail price is much higher than what is actually paid for an item). In Figure 1, graph (c) is for an Invicta Diver Watch (which has a reported retail price of \$370). Our MLE estimate suggests that the true valuation is actually more like \$256. When we use the MLE-based valuation, the plot of our predicted distribution matches much more closely to the data. In fact, when we check the price of this watch at other online retailers we find that the posted price at other internet sites is much closer to our MLE-based valuation than to the reported retail price.

In Panel B of Table 1, we use our MLE-based valuations and report the fraction of the items in each category for which we fail to reject at each of the three common statistical cut-offs. In fact, when using the MLE-based valuation we always fail to reject that the actual distribution for any of the items in the apparel category differs from the predictions of our model. The goodness of fit of our model goes up for all categories, though we still only fail to reject that our predictions differ from the data 46% of the time for the video game items. In aggregate, auctions of 51 of the 65 items (or 78%) can be explained by our theory (at the .05 p-value threshold).

As a final piece of evidence, in Figure 2 we reported the distance (in standard deviation units) between the actual and expected revenue based on our model. We find that most of the non-video game items clump very close to zero with 90% of these items having an actual revenue within one standard deviation of the prediction of our model. For the video-game items, we also see the mode occurring near zero but there is a large right tail to this distribution with 40% of the items having an actual revenue with more than a standard deviation of revenue above the prediction of the model.

5 Risk preferences

To this point, we have assumed bidders to be risk neutral. We now make a minor adaptation of our model by incorporating preferences towards risk. Since wealth is unobserved, we sidestep this by assuming CARA utility $u(w) = \frac{1-e^{-\alpha w}}{\alpha}$. From here, we follow the same process as before, imposing the indifference condition to solve for μ:

$$(1-\mu_{q+1})\frac{1-e^{-\alpha(w+v-sq-b)}}{\alpha} + \mu_{q+1}\frac{1-e^{-\alpha(w-b)}}{\alpha} = \frac{1-e^{-\alpha w}}{\alpha}, \tag{7}$$

which yields:

$$\mu_q = \frac{1-e^{\alpha(b+s(q-1)-v)}}{e^{\alpha b} - e^{\alpha(b+s(q-1)-v)}}. \tag{8}$$

5.1 100% off auctions

As before we can construct a probability density function from these μ. In the special case of 100% off auctions, this becomes:

$$f(q) = \left(\frac{1-e^{\alpha(v-b)}}{1-e^{\alpha v}}\right)^q \left(\frac{e^{\alpha v} - e^{\alpha(v-b)}}{1-e^{\alpha v}}\right) \tag{9}$$

which generates an expected revenue of $b\left(\frac{e^{\alpha(b-v)}-1}{1-e^{\alpha b}}\right)$. Note that expected revenue is decreasing in α; the auction generates more revenue as agents become less risk averse.

We can perform a maximum likelihood estimation on $f(q)$ as before; however, α is not separately identifiable from v or b. This is because the distribution has the form $f(q) = (1-\mu)\mu^q$. Thus, the MLE procedure can only identify μ.

Fortunately, all of the items where Swoopo has routinely used 100% off auctions have an obvious objective value for v. These items include cash ($1,000 or $80) and vouchers for free Swoopo bids (either 300 or 50, worth $0.75 per bid). The bid fee b is also exogenously set, so we can use the MLE procedure to identify α.

For these items, we find that α to be slightly negative, between -0.01 and -0.001. This implies that Swoopo participants are risk loving. The resulting theoretical distribution is depicted in Figure 3; with the MLE risk preferences, the fit is remarkable.

Using risk neutral bidders (and the objective value of v), the theoretical distribution would predict too few bids at every point and hence too high a probability of ending early. With the addition of risk-loving preferences, bidders become a bit more aggressive at every point of the auction.

5.2 Standard auctions

In the case of standard auctions, this environment no longer lends itself to creating a continuous approximation to the density function. However, we can still generate the distribution numerically and perform maximum likelihood estimations of both the valuation and the risk parameter.

The resulting risk parameter α indicates that bidders are mildly risk-loving, primarily in the range of -0.001 to -0.006. These values are in the same range as those found among bettors at horse tracks (Golec and Tamarkin, 1998). The resulting MLE valuation v are also plausible for each item, generally below the retail price with the exception of items that were in short supply during this data period.

The improvement in fit is remarkable. Figure 4 illustrates the new theoretical distribution, demonstrated on the same three items from the lower panel of Figure 1. The risk neutral model was unable to match the hump shape of the distribution because for any choice of the parameter v, $f(q)$ was decreasing in q. In the altered model, $f(q)$ will first rise then fall if α is sufficiently large.

Intuitively, this occurs because a low q implies a greater variance in the outcome (because the payoff if no one else bids, $v - sq$, is large). A risk loving participant will be willing to bid in spite of unfair odds when the gamble has a highly skewed payoff. Thus μ_q can be larger than in the risk neutral model (more aggressive bidding) while keeping bidders indifferent, and as a result, it is less likely that the auction ends with low q. As q increases, though, the variance in outcome diminishes, and the same risk loving participant will require closer to fair odds in order to bid. Indeed, as the final price sq approaches v, the resulting μ_q approximate the risk neutral result.

Panel C of Table 1 reports the degree of fit for all auctions. Note that nearly all of the auctions are explained by the model once risk preferences are included. In particular, nearly twice as many video game items fit the new model. t-Tests are not available because the distribution does not generate an analytical solution for expected revenue.

5.3 Risk-loving preferences

It is not particularly surprising that Swoopo participants would have a preference for risk; after all, this auction is essentially a form of gambling. Like a slot machine, the bidder deposits a small fee to play, aspiring to a big payoff (of obtaining the item well below its value). The only difference is that the probabilities of winning are endogenously determined.

The idea that a gambler would voluntarily take on risk, paying more than the expected payout to play, is puzzling. Economists have tried to rationalize such behavior in one of two routes: by assuming some intrinsic utility from placing the gamble or by assuming convex preferences with respect to wealth.[4] Regarding the former, it is noteworthy that Swoopo prominently advertises itself as "entertainment shopping." While many economists have referred to intrinsic utility as motivation for gambling, very few have provided formal modeling of the concept. Diecidue, Schmidt, and Wakker (2004) provide a brief but useful survey of these works and offer their own formal model. They also conclude that utility from gambling would necessarily contradict stochastic dominance.

We favor the latter approach: assuming convex preferences (that is, an increasing marginal utility of wealth within some range). Friedman and Savage (1948) provided an early formal model of this phenomenon. Their justification for increasing marginal utility was that major increases in wealth would promote a person into a new social and economic status, while minor increases in wealth would maintain a person within the same class and hence exhibit diminishing marginal utility. This view was supported by Gregory (1980) and Brunk (1981).

[4] In addition, much of behavioral economics is inspired by these and similar puzzles. These lines of work offer alternatives to expected utility, such as prospect theory, ambiguity aversion, or bounded rationality, to name a few. We omit discussion of these theories only due because because their application in a game theoretic setting is not a settled matter.

More recent work by Golec and Tamarkin (1998) and Garrett and Sobel (1999) proposes that what appears as "risk-loving" in the Friedman-Savage utility should be more aptly named "skewness-loving." In an empirical study of horse-track betting and US lottery games, respectively, both papers conclude that gamblers experience disutility from variance in payoffs (*i.e.* are still risk averse), but gain utility from skewness in payoffs. Golec and Tamarkin show that this will appear as risk-loving preferences when estimating in a typical CRRA utility. Our estimated risk parameters are within the same range as their estimates.

This view of risk-loving behavior seems appropriate for our model, particularly for items such as video games. The bidders may not be able to justify (to their spouse, their parent, or themselves) outlaying \$300 to buy the Wii at retail; however, the potential to win the Wii early in the auction and spend only a tenth of that makes it worth the \$0.75 gamble, even at unfair odds. As the number of bids (and hence the final price) increase, the skewness in outcome decreases and these gamblers are less tolerant of unfair odds. Indeed, as the final price approaches the retail price, μ_q eventually adjust so that bidding becomes a fair bet.

For more practical items such as computers or apparel, bidders behave as if they are risk neutral. This probably reflects that they intend to purchase the item at the (MLE esitmate) retail price if unsuccessful at Swoopo.

6 Conclusion

Our model presents a simple model of bidder choices in a pay-to-bid auction. In the unique symmetric subgame perfect equilibrium, potential bidders are indifferent about participating, and the exact mixed strategy is determined by this indifference condition. Using these mixed strategies we can establish that expected revenue will be near the value which bidders place on the auctioned item.

Even when restricted to risk neutral bidders, the model is able to explain a strong majority of the observed data. However, items associated with video game systems routinely produce more aggressive bidding and higher revenue than can be explained by this model, even if bidders valued the item far above its retail price. These exceptions are these are the most profitable of Swoopo's auctions, and are among the most frequently run.

Once we allow bidders to have preferences towards risk, the model can explain almost all of the observed auctions. Buyers of video game systems have small preferences for risk (which may be interpreted as a love of skewness), which makes them aggressive in the early phase of the bidding. This emphasizes the fact that pay-to-bid auctions are a form of gambling, though the odds on some items are fairer than on others; this seems related to the practicality of the item.

Appendix: Maximum Likelihood procedure

MLE for 100% off auctions

In the case of 100% off auctions, the maximum likelihood function produces an analytical solution. For a set I of auctions of the same item, let $\{q_i\}_{i\in I}$ denote the ending number of bids for each auction. The log-likelihood function would be $\sum_{i\in I} q_i \ln\left(1-\frac{b}{v}\right)+\ln b-\ln v$. If we take the first-order condition with respect to v, we obtain the solution (where m denote the number of auctions in I):

$$v = b\left(1+\frac{1}{n}\sum_{i\in I} q_i\right) \tag{10}$$

If we were to take the first-order condition with respect to b, we obtain the same equation, which is to say that we cannot separately identify v and b in the model. However, b has an objective dollar value set by Swoopo which we feel is reasonable to use. On the other hand, the retail price listed by Swoopo does not have to reflect the true valuation v.

MLE for standard pay-to-bid auctions

In the standard auction, the maximum likelihood function must be numerically solved, but the procedure is otherwise identical. The probability density function leads to a log-likelihood function of: $\sum_{i\in I}\frac{b-s}{s}\ln(2v-b-2sq_i)+\ln 2b-\frac{b}{s}\ln(2v-b)$. Taking the first-order condition with respect to v, we find:

$$\frac{1}{n}\sum_{i\in I}\frac{2sq_i}{2v-b-2sq_i}=\frac{s}{b-s} \tag{11}$$

Since the left hand side has a positive partial derivative with respect to v, there will be a unique v which solves this equation, which can be numerically determined.

Figure 1: Representative examples: comparison of the theoretical and actual distribution of ending bids for each item.

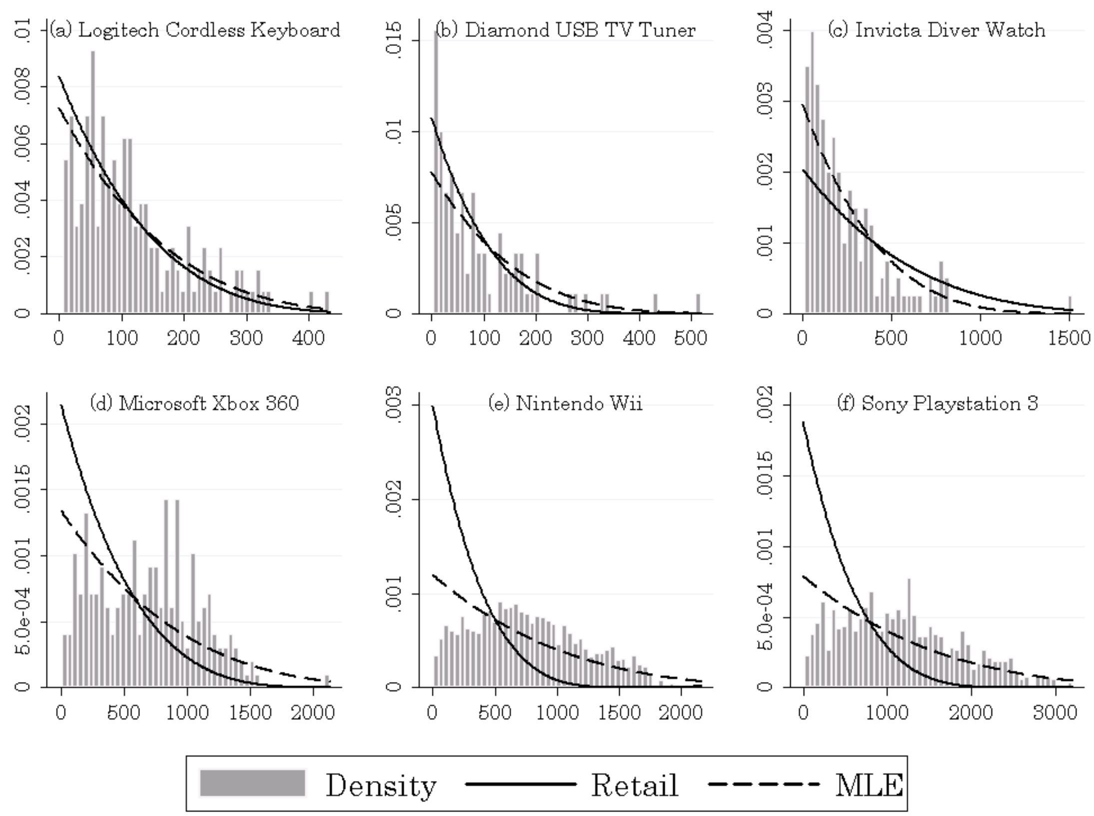

Notes: The unit of observation is the individual auctions for each item. In each figure, the x-axis denotes the final number of bids placed for the item. The y-axis is the probability that the auction concludes at that number of bids. Bars denote the observed frequencies; the solid line gives the theoretical frequencies using the retail value, while the dashed line does the same using the MLE estimated value.

Table 1: Statistical tests comparing theoretical distribution with observed data.

A. Using retail price as item valuation.

		Pearsons χ^2 Test (compares distributions)			t-Test (compares means)		
Item Type	N	$p \geq .10$	$p \geq .05$	$p \geq .01$	$p \geq .10$	$p \geq .05$	$p \geq .01$
Computers	7	0.714	0.714	0.714	0.286	0.429	0.429
Computer Accessories	24	0.609	0.609	0.696	0.292	0.333	0.417
Apparel	12	0.417	0.417	0.417	0.083	0.083	0.167
Home Electronics	7	0.286	0.286	0.286	0.143	0.143	0.286
Video Games	15	0.222	0.222	0.333	0.000	0.067	0.133

B. Using MLE-based valuation of item

		Pearsons χ^2 Test			t-Test		
Item Type	N	$p \geq .10$	$p \geq .05$	$p \geq .01$	$p \geq .10$	$p \geq .05$	$p \geq .01$
Computers	7	1.000	1.000	1.000	1.000	1.000	1.000
Computer Accessories	24	0.792	0.833	0.875	0.833	0.833	0.917
Apparel	12	1.000	1.000	1.000	0.833	1.000	1.000
Home Electronics	7	0.714	0.714	0.857	0.857	0.857	0.857
Video Games	15	0.467	0.467	0.533	0.400	0.533	0.600

C. Using MLE-based risk preferences and valuation of item

		Pearsons χ^2 Test		
Item Type	N	$p \geq .10$	$p \geq .05$	$p \geq .01$
Computers	7	1.000	1.000	1.000
Computer Accessories	24	0.875	0.875	0.917
Apparel	12	1.000	1.000	1.000
Home Electronics	7	0.875	0.875	1.000
Video Games	15	0.733	0.800	0.800

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

Figure 2: Deviation from Expected Revenue in risk neutral model (separately for video games and other items)

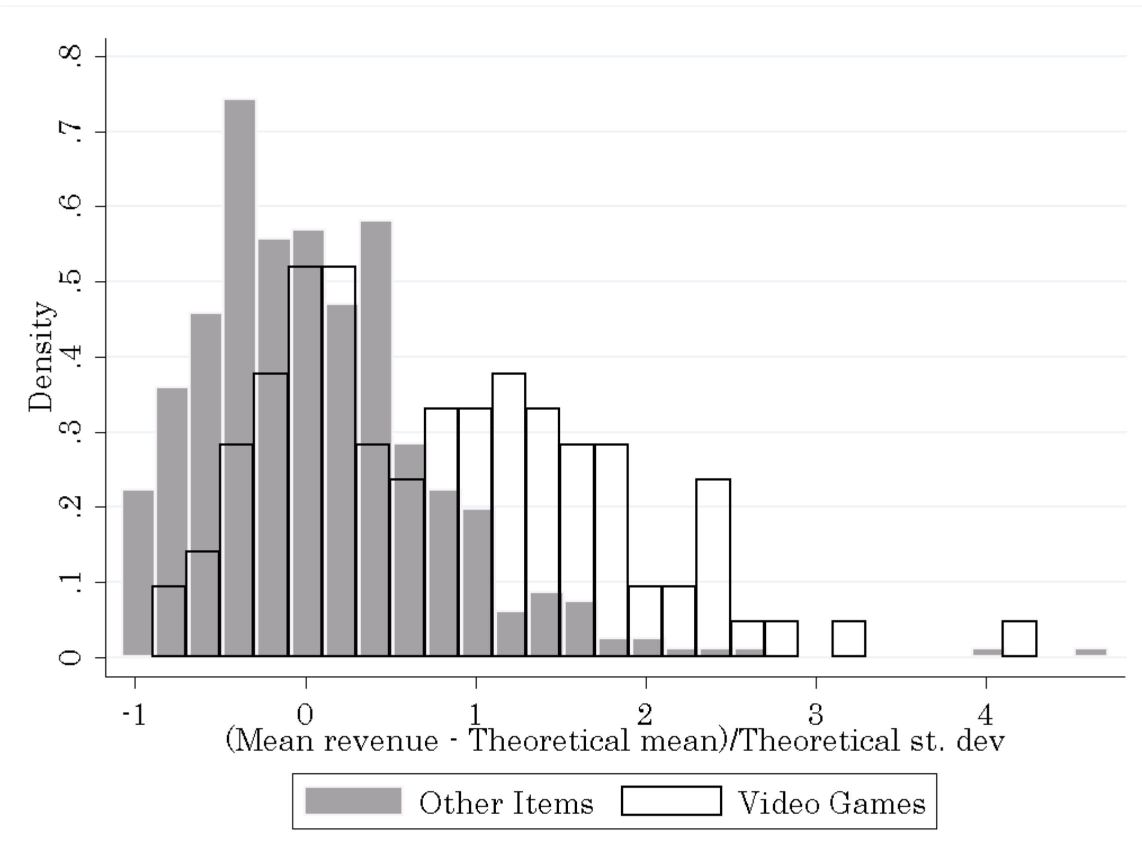

Notes: The unit of observation is each unique item. The x-axis indicates by how many standard deviations the observed average revenue differs from the theoretical expected revenue. The y-axis indicates the fraction of items auctioned with that observe outcome.

Figure 3: 100% Off Auctions with risk-loving bidders: comparison of the theoretical and actual distribution of ending bids for each item.

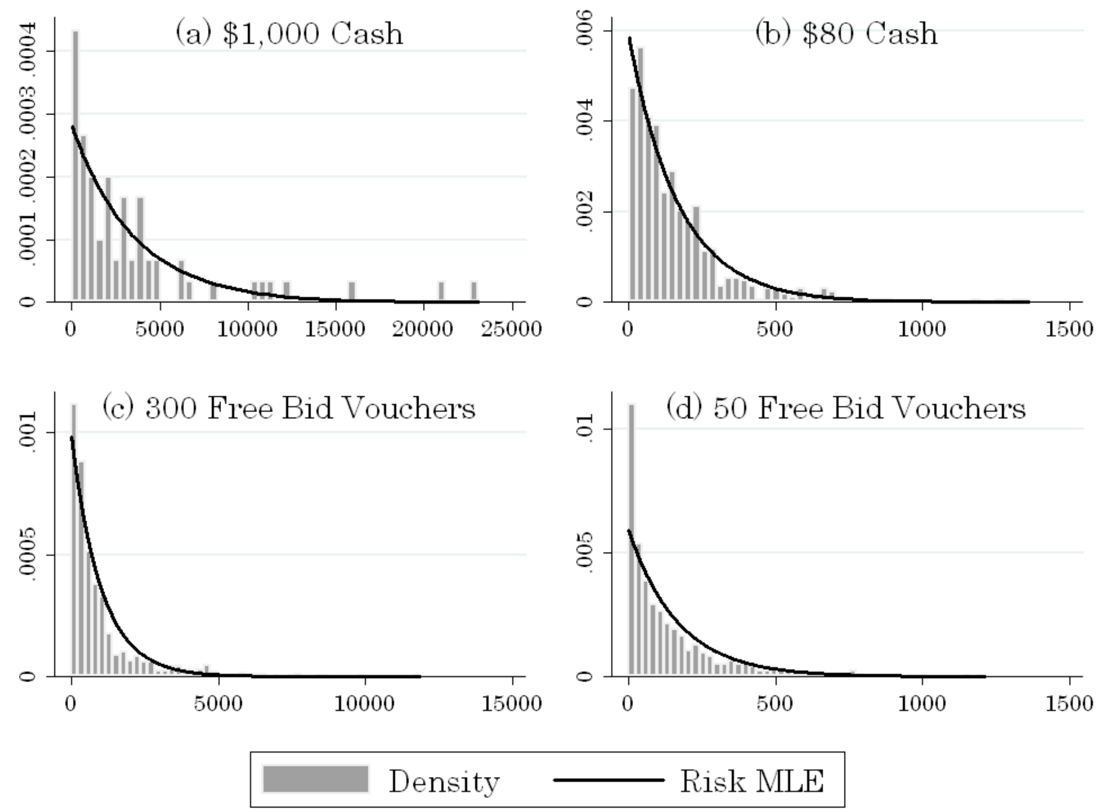

Notes: Risk MLE denotes the theoretical distribution of ending bids when MLE-estimated bidder risk preferences is used; item value is assumed to be retail value.

Figure 4: Standard auctions with risk-loving bidders: comparison of the theoretical and actual distribution of ending bids for each item.

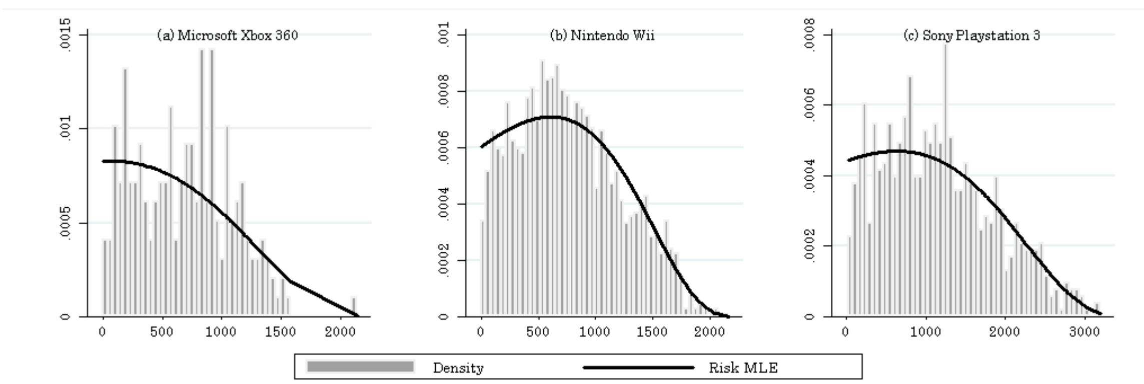

Notes: Risk MLE denotes the theoretical distribution of ending bids when MLE-estimated item value and bidder risk preferences are used.

References

Brunk, Gregory. (1981) "A Test of the Friedman-Savage Gambling Model," *Quarterly Journal of Economics* **96**, 341-348.

Che, Yeon-Koo, and I. Gale. (1998) "Standard Auctions with Financially Constrained Bidders," *Review of Economic Studies* **65**, 1-21.

Diecidue, Enrico, U. Schmidt, and P. Wakker. (2002) "The Utility of Gambling Reconsidered," *Journal of Risk and Uncertainty* **29**, 241-259.

Friedman, Milton, and L. Savage. (1948) "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy* **56**, 279-304.

Gavious, Arieh, B. Moldovanu, and A. Sela. (2002) "Bid Costs and Endogenous Bid Caps," *RAND Journal of Economics* **33**, 709-722.

Golec, Joseph, and M. Tamarkin. (1998) "Bettors Love Skewness, Not Risk, at the Horse Track," *Journal of Political Economy* **106**, 205-225.

Garrett, Thomas, and R. Sobel. (1999) "Gamblers favor skewness, not risk: Further evidence from United States lottery games," *Economics Letters* **63**, 85-90.

Gregory, Nathaniel. (1980) "Relative Wealth and Risk Taking: A Short Note on the Friedman-Savage Utility Function," *Journal of Political Economy* **88**, 1226-1230.

Rapoport, Amnon, H. Otsubo, B. Kim, and W. Stien. (2009) "Unique Bid Auction Games," *Jena Economic Research Papers* 2009-005.