

The Role of Risk Preferences in Pay-to-Bid Auctions

Brennan C. Platt, Joseph Price

Department of Economics, Brigham Young University, 135 FOB, Provo, UT 84602, (801) 422-8904, brennan_platt@byu.edu

Henry Tappen

Microsoft Corporation

We analyze a new auction format in which bidders pay a fee each time they increase the auction price. Bidding fees are the primary source of revenue for the seller, but produce the same expected revenue as standard auctions (assuming risk-neutral bidders). If risk-loving preferences are incorporated in the model, expected revenue increases. Our model predicts a particular distribution of ending prices, which we test against observed auction data. The degree of fit depends on how unobserved parameters are chosen; in particular, a slight preference for risk has the biggest impact in explaining auction behavior, suggesting that pay-to-bid auctions are a mild form of gambling.

Key words: penny auction, bid fees, risk-loving preferences

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, but occasionally a site develops a novel approach. Such is the case with Swoopo, a German company founded in 2005 which has operated websites in the United Kingdom, Spain, the United States, and Austria. Swoopo's distinctive 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. When a participant places a bid, the current price increases by a fixed amount (\$0.15 in the US site's typical auction in early 2009), the bidder is immediately charged a bid fee (\$0.75), and the auction is extended by a set amount of time (15 seconds).¹ If the time expires before another bid is placed, the last bidder pays the current price (on top of any bid fees incurred) and wins the object.

An observer's first experience with this auction format typically follows a predictable course. First, the newcomer notes that the current price on her particular item is remarkably low—indeed,

¹ In an English auction, this last feature is known as a *soft close*. Extending the clock ensures that participants always have an opportunity to respond after being outbid, often expressed in oral auctions as “Going once, going twice, sold!” Amazon.com used to offer English auctions with a soft close, and uBid.com currently does, while e-Bay uses a hard close where the ending time is fixed (See Roth and Ockenfels 2002).

the median auction closes with a final price that is 10% of the retail price. The clock is within seconds of expiring, yet our observer quickly realizes that the deal is elusive, as additional bids are placed every five to ten seconds. Indeed, her past bids (had she placed any) are sunk, having no bearing on her likelihood of winning once she is outbid. Ideally, she would wait and place the last bid, but she immediately recognizes this to be the fundamental dilemma of the auction: while it is highly likely that new bids will be submitted, there is a small chance that none will and the auction will end.

On appreciating the low probability of winning, thoughts turn to how much revenue the auction generates. For every dollar increase in the final price, Swoopo collects five dollars in additional revenue through bid fees. To make a profit, the final price only needs to exceed one-sixth of the item's cost. This thought process has led many to question the rationality of auction participants as well as the ethics of the auctioneers. Blogs and news articles (Kato 2009, Reklaitis 2009, Gimein 2009) have vented their frustrations, referring to Swoopo as gambling or as an outright scam.

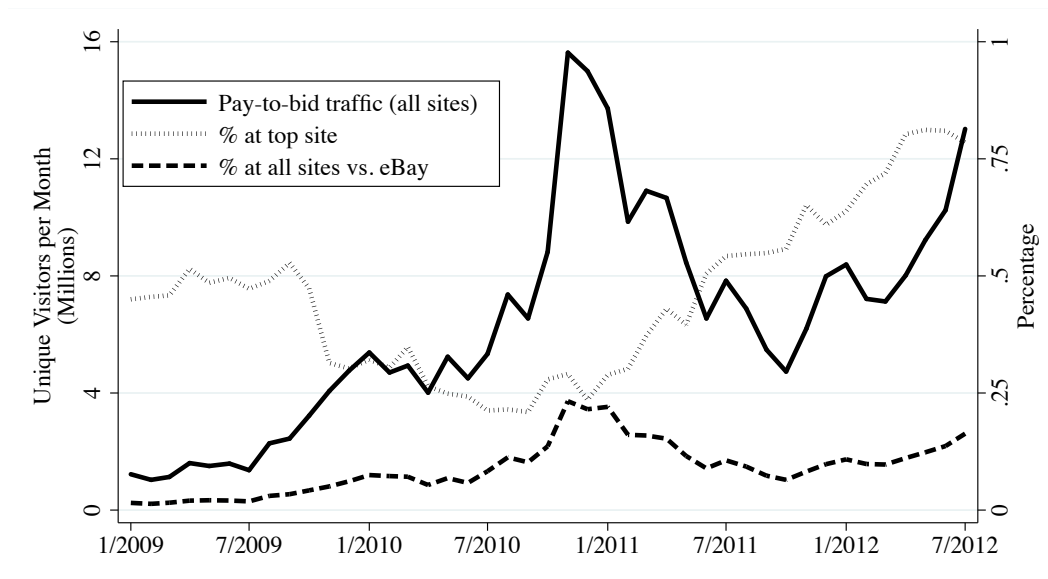
Even so, pay-to-bid auctions are growing in popularity. In January 2009, 10 websites conducted such auctions; within one year, the format had proliferated to 112 websites. This competitive pressure led to reduced bid fees and profit, with many entrants exiting over the course of 2011 — including Swoopo, which filed for bankruptcy in March of that year. Even after this consolidation, traffic among remaining sites (see Figure 1) has grown to 13 million visitors per month. For comparison, traffic at eBay fluctuated around 75 million unique visitors per month throughout this period. In other words, pay-to-bid auctions have garnered 16% of the traffic held by the undisputed leader in online auctions.

This paper proposes a parsimonious model of the pay-to-bid auction format and tests the extent to which observed bidding is consistent with the model's predictions. The key insight of this model is that in equilibrium, the probability of being outbid must be consistent with prior bidders being willing to bid *ex-ante*. This necessary condition generates a density function for the probability that the auction ends at any given number of bids.

We examine the predictions of our model using information on 49,000 auctions collected from Swoopo's website. Our goal is to assess which features are needed to replicate the observed distribution of ending bids. The theoretical distribution depends on two parameters that may not be directly observed: the buyers' valuation of the item and their risk preferences regarding the item. We first assume risk neutrality and use average retail prices from Amazon.com as a proxy for buyer valuations; however, this results in a poor fit, explaining only 10% of routinely auctioned items.

Alternatively, we use maximum likelihood to estimate one or both parameters (separately for each auctioned item). Using estimated buyer valuations improves the fit to explain 57% of items. Using Amazon prices with estimated risk preferences does even better, explaining 75% of items.

Figure 1 Monthly traffic on pay-to-bid websites.



Notes: Data from Compete.com, a web traffic monitoring company. The solid line indicates the number of unique visitors per month across all pay-to-bid websites. The dotted line indicates the percentage of those visits held by the top pay-to-bid website (Swoopo through 7/2009, BidCactus through 2/2011, and QuiBids thereafter). The dashed line provides the ratio of all pay-to-bid traffic versus eBay traffic.

With both degrees of freedom, the model explains 86% of items. In particular, bidders appear to be mildly risk loving; the estimated level of absolute risk preference is typically an order of magnitude smaller than that of racetrack bettors. The auction also generates above-normal profits when bidders are risk loving. This analysis suggests that gambling is an essential element of pay-to-bid auctions, driven largely by risk preferences rather than an intrinsic joy of winning.

We assume that the valuation of the item is known and the same across all potential bidders. While only Alcalde and Dahm (2011) analyze a first-price auction in a complete-information environment, this assumption is quite common in all-pay or war-of-attrition auctions (*e.g.* Baye *et al.* 1996, Clark and Riis 1998, Barut and Kovenock 1998, Konrad and Leininger 2007). The common value allows us to isolate a key aspect of pay-to-bid auctions: a bidder is gambling that others who value the item *at more than the current price* may still abstain from bidding. We believe that this would still be the driving force (though less visibly) in a private valuation framework.

Furthermore, a common value is quite plausible for the types of items regularly auctioned on Swoopo. All items are new, unopened, and frequently available from traditional or internet retailers. Unlike rare art or collectibles, the market prices of these items are well established.

1.1. Related Literature

Pay-to-bid auctions are a new phenomenon, but have quickly garnered academic scrutiny. Our paper is most closely related to Augenblick (2009), who independently develops the same risk-neutral model we present in Section 2. Our work differs in two important ways.² First, in the empirical analysis, Augenblick aggregates auction results across different items using a normalization that makes the probability of the next bid comparable across items of different value. This provides a large number of observations for the empirical test, in which he concludes that the risk-neutral model cannot explain these auctions because the hazard rates are lower than predicted.

The disadvantage of that approach is that it cannot detect differences in fit among the various items. When evaluated item-by-item (as we do), some of the goods sold on Swoopo produce a distribution similar to that of the risk-neutral model. At the same time, others (mostly home electronic and video-game related items) have a very different shape in their distribution of ending prices (see Figure 2). These exceptional items produce the greatest profit, and thus are the most frequently auctioned, constituting 58% of all observed auctions. As a consequence, the aggregate test fully rejects the risk-neutral model.

The key difference in our theoretical work is in the adaption of the baseline model, in an attempt to explain excess profits. Augenblick incorporates a behavioral model of errors in judging sunk costs. An unsuccessful bidder experiences regret for past bid fees if he leaves empty handed, which makes him more eager to continue bidding. Bidders also underestimate their future regret (*i.e.*, bidders are not time consistent). This produces excess profits, but is not capable of producing the hump-shaped distribution observed for video game systems. Indeed, as long as the bid fee is greater than the bid increment, the density function resulting from the sunk cost model will be strictly decreasing, whereas our risk-loving model produces an excellent fit (see Figure 4).

Alternative explanations for excess profits are offered by Byers *et al.* (2010). They introduce asymmetries into our standard model, where some bidders are either better informed regarding the number of active bidders or the intent of other bidders, have different valuations on the item, or have access to cheaper bid fees. They show that this informational disparity can create both more aggressive bidding and higher expected revenue than the standard model. The key to this analysis, though, is that bidders with wrong beliefs think that all other players have the same wrong beliefs and, indeed, are unaware that their beliefs could be wrong. This produces a thicker tail in the distribution of ending bids, but cannot produce the hump shape.

Both Augenblick (2009) and Byers *et al.* (2010) invoke potentially interesting behavioral assumptions to explain bidder behavior, and formalize many popular assertions regarding pay-to-bid auctions. While such an approach can prove fruitful, we find it valuable to first determine what can

²Beyond the scope of our work, Augenblick (2009) also empirically investigates when Swoopo should initiate new auctions to maximize profit, concluding that they supply more than optimal, perhaps for entry-deterrence.

be explained in a model of fully rational, utility-maximizing bidders. This establishes a baseline from which to judge the value added of a behavioral model. In Section 5.3, we empirically evaluate these alternative models, granting them the same latitude in parameter selection.

At first glance, one might consider the pay-to-bid auction as a mere reformulation of the all-pay auction. In an all-pay auction, each participant pays what he bids, even though only the highest bidder wins the item. A second-price all-pay auction (or war of attrition) does the same except that the winner pays the second highest bid. As in ours, these models often assume a common value on the item being sought, and often reach a mixed strategy equilibrium with bidders having an expected payoff of zero. These are often modeled in a static, sealed-bid environment (*e.g.* Maynard Smith 1974, Amann and Leininger 1996, Baye *et al.* 1996, Krishna and Morgan 1997), but their strategic properties differ from the pay-to-bid auction even if set up in a dynamic format.

The analogy arises because each time an all-pay participant raises his bid, he commits to pay that increase regardless of whether he wins. This commitment is hence like a bid fee. However, in a pay-to-bid auction, the bid fee is distinct from the price increment, and the winning bidder must pay the final price on top of his bid fees. Moreover, in an all-pay auction, if *any* active bidder increases the bid, *every* active bidder would have to likewise increase their own bids to remain active (much like calling a bet in poker). A pay-to-bid participant only incurs a bid fee each time *he* increases the price. Even so, as with the war of attrition, our model has ready application to rent-seeking or competition for mates, which are discussed in the conclusion.

2. Model

We begin by formalizing the auction rules. An item being sold has a known, objective value of v to n potential bidders (or *customers*), who enjoy utility $u(w)$ from a dollar payoff w . The state of the auction is described by the number of elapsed periods, q , and the current winning bidder, $i \in \{1, \dots, n\}$.³ The price begins at $p_0 = 0$. Each time someone bids, a new period begins and the price is raised by exactly s dollars. Thus, the price in period q is $p_q = s \cdot q$.

During each period, the $n - 1$ customers (or n customers in the initial period) who are not currently winning simultaneously choose whether to place a bid. If no one places a bid, the auction closes and the bidder currently winning pays the current price p_q and receives the item. If $k > 0$ customers place a bid, one of them is randomly selected with probability $\frac{1}{k}$; that customer becomes

³ Note that the individual bidding histories are not included in the state, as past bids are sunk. Some variations of pay-to-bid auctions would require the full history, leading to a much more complex solution. For instance, in Summer 2010, Swoopo introduced a *Swoop-it-now* option which allowed unsuccessful bidders to purchase the item at retail price minus the bid fees they had already paid. Swoopo does not report when bidders exercise this option, so we are unable to perform any empirical investigation of it; instead, we restrict our data to auctions before May 2010.

the new current winning bidder, and must immediately pay b dollars as a bid fee.⁴ A new period then begins. Thus, if a customer has initial wealth w and places the q^{th} bid, he either obtains either $u(w - b)$ if someone else places the $q + 1^{\text{th}}$ bid, or $u(w + v - b - s \cdot q)$ otherwise. Not bidding leaves him with $u(w)$.

This constitutes a complete-information, extensive-form game; we examine the symmetric subgame perfect equilibria. Here, symmetry requires that at period q , all customers who are not currently winning employ the same mixed strategy $\beta_{q+1} \in [0, 1]$ of attempting to place the $q + 1^{\text{th}}$ bid. Let $Q \equiv \frac{v-b}{s}$, or if $s = 0$ then $Q = \infty$. When $s > 0$, we assume that Q is an integer; we comment on this assumption after presenting the equilibrium solution.

2.1. Risk-Neutral Bidders

As is typical in the auction literature, we first consider risk-neutral bidders, *i.e.* $u(w) = w$. This leads to a very clean characterization of equilibrium which does not depend on initial wealth w .

PROPOSITION 1. *Assume $u(w) = w$. Let $\mu_1 \in [0, 1]$, $\mu_q = 1 - \frac{b}{v-s \cdot (q-1)}$ for $1 < q \leq Q$, and $\mu_q = 0$ for $q > Q$. Then the strategy profile $\beta_1 = 1 - (1 - \mu_1)^{\frac{1}{n}}$ and $\beta_q = 1 - (1 - \mu_q)^{\frac{1}{n-1}}$ for $q > 1$ constitutes a symmetric subgame perfect equilibrium.*

The proof is provided in the appendix. The intuition is simply that the hazard rate of the next bid occurring, μ_q , must take its stated value to make customers indifferent about bidding in period $q - 1$. As in any mixed strategy equilibrium, customers are indifferent, but follow the individual strategy β_q since anything else would break that indifference. Only the probability that anyone places the first bid, μ_1 , is non-unique, since there is no prior bidder who must be made indifferent.

Indeed, while other symmetric equilibria exist, they must break indifference in some period, and this inevitably leads to a degenerate outcome, as characterized in the following:

PROPOSITION 2. *In any symmetric subgame perfect equilibrium besides that proposed in Proposition 1, the auction will end either in period 0 with no bidders, or in period 1 with one bidder.*

The war-of-attrition game in Maynard Smith (1974) produces similar degenerate equilibria. One can obtain this degenerate outcome through a variety of equilibrium strategies (mentioned in our proof, and given further attention in Augenblick 2009); yet the outcome is trivial, and does not appear to be empirically relevant. Although some auctions do conclude without any bids or with only one, they occur with low frequency, consistent with our equilibrium in Proposition 1.

⁴For actual pay-to-bid websites, a new period begins once any bid is received; thus, ties never technically occur. However, two bidders could submit their bids so close together that the one with the slightly later bid could be unaware that the earlier bidder had initiated a new period. Hinosaar (2010) models ties by charging all k tied players the bid fee, randomly selecting one as the current winner. The equilibrium outcome coincides with ours when $s = 0$ or $n = 2$; more generally, it retains similar properties though the analysis is more complicated.

We do not consider asymmetric equilibria, in which otherwise homogeneous bidders employ different mixed strategies at a particular q . Since these auctions are essentially anonymous, it is difficult to see how one would credibly communicate his asymmetric strategy. In any case, if these asymmetric strategies still produced the same aggregate probability μ_{q+1} of the $q + 1^{\text{th}}$ bid occurring, then they would still be payoff equivalent to our equilibrium in Proposition 1. On the other hand, if the aggregate probability differs from μ_{q+1} , a bidder in period $q - 1$ would no longer be indifferent about placing the q^{th} bid; thus, the mixed strategy equilibrium would unravel, producing a degenerate outcome just as in Proposition 2.

In our analysis, Q played an important role as the period in which the current price is equal to the item's value minus the bid fee. Because of this, no one is willing to bid after period Q . We have assumed that Q is an integer, and in our view this is not a very drastic assumption. For instance, when $s = \$0.01$ (a penny auction), then this only requires that v be expressed in whole dollars and cents, since b already is. Augenblick (2009) shows that even if this did not literally hold, an ϵ -perfect equilibrium would approximate the equilibrium in our Proposition 1.

Note also that when $s = 0$, $Q = \infty$, meaning that there is no upper bound on the potential number of periods. Moreover, $\mu_q = 1 - \frac{b}{v}$ is constant for all q . This is not surprising, since the payoff from winning is constant in q , rather than falling as it does when $s > 0$.

2.2. Expected Revenue

The aggregate probability of the next bid occurring, μ_{q+1} , is of particular importance in establishing the expected outcome of the auction. This conditional probability (that the $q + 1^{\text{th}}$ bid occurs, given that the q^{th} already has) allows us to construct the probability density that the auction ends at exactly q bids:

$$f(q) \equiv (1 - \mu_{q+1}) \prod_{j=1}^q \mu_j = \begin{cases} 1 - \mu_1 & \text{if } q = 0 \\ \frac{b}{v-s \cdot q} \mu_1 \prod_{j=2}^q \left(1 - \frac{b}{v-s \cdot (j-1)}\right) & \text{if } 1 \leq q \leq Q. \end{cases} \quad (1)$$

This density function is decreasing in q ; that is, the unconditional probability of ending at a given number of bids decreases as bids increase. When b is small relative to v , $f(q)$ can be approximated by a generalized Beta distribution of the first kind (treating q as continuous rather than discrete), which is demonstrated in the appendix.

Some auctions are conducted with $s = 0$, so the winner pays nothing beyond his own bid fees; this is called a *100% off auction*. In that case, the distribution simplifies to:

$$f(q) = \begin{cases} 1 - \mu_1 & \text{if } q = 0 \\ \mu_1 \frac{b}{v} \left(1 - \frac{b}{v}\right)^{q-1} & \text{if } q \geq 1. \end{cases} \quad (2)$$

Indeed, aside from the adjustment due to μ_1 , this is just a geometric distribution.

The first benchmarks in any auction model are expected revenue and efficiency. The latter is not particularly relevant here, since all customers value the item equally; however, the outcome is inefficient if the item is not sold, which happens with probability $1 - \mu_1$. Of course, if the item being sold is durable, the seller is able to immediately initiate a new auction at practically no cost, repeating until the item is sold. Thus, one can reasonably set $\mu_1 = 1$, which is equivalent to considering the expected revenue conditional on the item being sold.

We can calculate the expected revenue of the seller directly using the probability density f .

$$E(\text{Rev}) = \sum_{q=1}^Q (b+s)qf(q). \quad (3)$$

When $s = 0$, it is straightforward to directly compute that $E(\text{Rev}) = \mu_1 v$; when $s > 0$, the direct calculation can be done, but is cumbersome to present. We 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 someone wins the auction, or in other words, the probability that at least one bid is placed. This computation yields $\mu_1 \cdot v$. This expected surplus is split between the seller and buyers, yet by construction, the expected surplus of the buyers is zero; hence, the seller's expected revenue is $\mu_1 v$ even when $s > 0$.

We note that under risk neutrality, this auction is not any more lucrative than standard auction formats. A typical 1st or 2nd price sealed-bid auction among (nearly) identical buyers would raise a revenue close to v . If we assume $\mu_1 = 1$ as described above, we obtain the same revenue v from the pay-to-bid auction, independent of the size of bid fee, bid increment, or initial price.

The variance of revenue, however, is dependent on these parameters. The direct computation of variance is cumbersome; instead, we use the Beta distribution approximation of $f(q)$ and obtain a variance of: $b \cdot (v - s)^2 / (b + 2s)$. Variance is increasing in b and v , and decreasing in s .

2.3. Risk Preferences

In the preceding analysis, we assumed that bidders are risk neutral. A natural extension is to incorporate preferences towards risk, since placing a bid is inherently a gamble on whether anyone else will place the next bid. Bidders still maximize expected utility, only now with functional form $u(w) = \frac{1 - e^{-\alpha w}}{\alpha}$. The virtue of using CARA utility is that decisions are independent of initial wealth, which is unobserved in our empirical setting. We follow the same process as before, requiring bidders to be indifferent between placing a risky bid versus not participating, then solving for μ :

$$(1 - \mu_{q+1}) \frac{1 - e^{-\alpha(w+v-sq-b)}}{\alpha} + \mu_{q+1} \frac{1 - e^{-\alpha(w-b)}}{\alpha} = \frac{1 - e^{-\alpha w}}{\alpha}, \quad (4)$$

which yields:

$$\mu_q = \frac{1 - e^{\alpha(b+s(q-1)-v)}}{e^{\alpha b} - e^{\alpha(b+s(q-1)-v)}}. \quad (5)$$

As before, individual strategies $\beta_q = 1 - (1 - \mu_q)^{\frac{1}{n-1}}$ will constitute a symmetric subgame perfect equilibrium; and indeed, the reasoning is exactly as in the proof of Proposition 1. Similarly, the same degenerate equilibria exist as in Proposition 2.

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

$$f(q) = \left(\frac{1 - e^{\alpha(v-b)}}{1 - e^{\alpha v}} \right)^q \left(\frac{e^{\alpha(v-b)} - e^{\alpha v}}{1 - e^{\alpha v}} \right). \quad (6)$$

As in the risk-neutral model, this is a geometric distribution, though risk preferences have altered its parameters. It 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. Indeed, if customers are risk loving ($\alpha < 0$), then expected revenue is greater than the customers' valuation of the item, v .

After incorporating risk preferences in a standard pay-to-bid auction with $s > 0$, the resulting probability density is far less tractable:

$$f(q) \equiv (1 - \mu_{q+1}) \prod_{j=1}^q \mu_j = \frac{1 - e^{\alpha b}}{e^{\alpha b} - e^{\alpha(b+s q - v)}} \prod_{j=1}^q \left(\frac{1 - e^{\alpha(b+s(j-1)-v)}}{e^{\alpha b} - e^{\alpha(b+s(j-1)-v)}} \right). \quad (7)$$

When b is small relative to v , $f(q)$ can be approximated by an exponential generalized beta distribution of the first kind, truncated to $q \geq 0$, as shown in the appendix. As $\alpha \rightarrow 0$, this density approaches the risk-neutral solution in Equation 1. Direct analytic solutions for expected revenue are no longer possible. Indeed, even the indirect method used in the risk-neutral case is no longer applicable, since the auction not only creates value by transferring an item to someone who values it, but also creates risk, whether for good ($\alpha < 0$) or ill ($\alpha > 0$). However, numerical computation is relatively simple, revealing a few key features of f .

First, the support of f is the same as before, placing positive probability everywhere from $0 \leq q \leq Q$. Increases in v have essentially the same effect that they did in the risk-neutral case: it will increase the support and flatten the distribution. When $\alpha > 0$ (*i.e.* risk averse), f has a similar convex shape to the risk-neutral density function, only with greater curvature as α rises.

The distribution behaves quite differently for $\alpha < 0$ (*i.e.* risk loving). When α is very close to zero, an inflection point \hat{q} is introduced near zero such $f'' < 0$ below \hat{q} ; thus f is no longer strictly convex. As α decreases, this inflection point takes on higher values, and eventually, creates a hump-shaped distribution. Indeed, the density function is maximized at $q = \max \left\{ 0, \frac{v-b}{s} + \frac{1}{\alpha s} \ln \left(\frac{1 - e^{\alpha b}}{1 - e^{\alpha s}} \right) \right\}$, which increases as α becomes more negative or v becomes larger.

While we cannot analytically solve for expected revenue in this environment, we can derive the effect of α on the expected revenue.

PROPOSITION 3. *In the equilibrium expressed by Equation 5, the average final bid is decreasing in α , s , or b . This implies that expected revenue is decreasing in α .*

When b or s increase, the average final bid is decreasing, but the revenue per bid is increasing. Thus, the comparative static on average revenue must be numerically evaluated. Across numerous parameterizations, we consistently find that revenue is increasing in b and decreasing in s iff $\alpha < 0$.

These results are quite intuitive. As customers become more risk loving, they enjoy more utility from this risky auction, which the auctioneer fully extracts due to the indifference condition. On the other hand, an increase in b or decrease in s creates larger variance in outcome, which increases the utility of risk-loving customers but decreases that of risk-averse customers.

3. Data

From the inception of their US website in September 2008 through May 2010, Swoopo.com auctioned over 126,000 items, all via pay-to-bid auctions. Most of these auctions are repetitions of identical objects, with 1,958 unique items.⁵ Furthermore, a small handful of items are auctioned much more frequently than others. We focus on the 172 items that were auctioned more than 100 times, accounting for 49,000 auctions, with the most popular (a Nintendo Wii system) being auctioned 3,307 times. Excluded are 23 different “bidpacks,” in which the winner receives a certain number of bid fee vouchers. Through January 2009, Swoopo also experimented with the 100% off auction, where the winner of the auction pays nothing for the item except for the money already paid in bid fees. These constituted 12% of all auctions in that date range, but only four items were repeatedly auctioned in this format.

Swoopo lists all of their ended auctions on their website. For each auction, the site provides the final auction price, the bid fees paid by the winner, and the end time. Also listed are the bid fee and the price increment that occurs with each new bid. Most auctions increase by either \$0.15 (prior to July 2009) or \$0.12 (after that date), but penny auctions increase by \$0.01. 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 bid fee and add this to the final price to determine Swoopo’s total revenue for that auction. While Swoopo also lists a suggested retail price for each item, these are significantly higher than prices available at other online retailers. Thus, we replace these with prices found on Amazon.com over the same time period,⁶ and use this throughout as our measure of retail price.

⁵ There are 2,218 unique names among these items, but we combine some items together, such as products that differ only in color or movies that have the same price but different titles. On the other hand, when the same product was auctioned under different bid fees or price increments, we treat these as distinct items.

⁶ Amazon prices were obtained from myPriceTrack.com, which lists historical prices offered by Amazon and its third-party affiliates. For each item, we use the mean price at which it was offered over the timeframe that the item was auctioned. Among items priced under \$1,000, Swoopo’s listed retail price is roughly 10% higher than Amazon’s average price; for higher priced items, Swoopo overstates the price by only 2%. If we consider the maximum Amazon price over the same period, Swoopo and Amazon nearly agree.

Swoopo also provides the usernames of the winner and last 10 bidders of each auction. We do not observe the full history of bids in our data—that is, the identity of each bidder for each period. This is of little consequence, though, since our model predicts bidder indifference about bidding at any point in time, and thus has little to say regarding individual strategies.

A common question 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, there are 2,086 auctions in which the first bidder wins the auction and 13,308 that are won by one of the first 9 bidders. In fact, two 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 is decreasing in the number of bids.

Swoopo is not the only website in the US to offer pay-to-bid auctions, but it attracted half a million unique visitors per month throughout most of our sample period, which consistently placed it among the top five pay-to-bid sites. One advantage of studying Swoopo is that, unlike many competitors, it provided information on all past auctions. Also, as the creators of this auction format, their rules were the most transparent. Later entrants began to differentiate themselves with more exotic bid fee pricing and other features that stretch beyond the scope of our theory.

4. Evidence

Our empirical objective is to assess under what conditions our model of rational pay-to-bid participants can explain observed bidding on Swoopo. Our key theoretical prediction is that the final number of bids in a given auction is a random variable with distribution $f(q)$. If a given item is repeatedly auctioned, we can test whether this sample distribution is consistent with its theoretical counterpart. In particular, we examine the role played by the bidders' common valuation and risk preference in replicating the observed behavior with the theoretical model. This process is somewhat challenging since the correct value for model parameters v and α may not be obvious a priori, as both relate to the motivation for gambling.

Regarding the item valuation, a natural approach is to set the valuation equal to the retail price; yet there are good reasons why even Amazon's price might not reflect the true valuation. For instance, the euphoria of winning an item may provide extra utility beyond what a standard purchase would, and would result in a *joy-of-winning premium*. In this vein, Swoopo prominently advertised itself as "entertainment shopping." In addition, the Amazon price is itself a noisy sample that may not always reflect the broader market price of other retailers or the item's general availability.

An alternative approach to explain gambling behavior is to assume that participants are *risk loving*. In our functional form, this means $\alpha < 0$; yet this does not suggest a precise value for α . Several papers have studied risk preferences of race track bettors, but their estimates are often

sensitive to whether long shots are included and the size of the bet. Using our same functional form and assuming a bet (*i.e.* our bid fee) of £1, Jullien and Salanié (2000) estimate $\alpha = -0.055$. Ali (1977) uses a functional form $u(w) = aw^\gamma$ and estimates an absolute risk aversion (comparable to our α) of -0.178. Kanto, *et al* (1992) and Golec and Tamarkin (1998) arrive at similar estimates.

Due to these uncertainties in parameter values, we take a two-step approach in testing our model. In the first step, we compute a maximum likelihood estimate of one or both parameters, separately for each regularly-auctioned item. That is, we chose the parameter(s) so as to maximize $\sum_i \ln f(q_i; v, \alpha)$, where i represents each observed auction of that item, q_i is the ending number of bids in that auction, and f is the theoretical distribution given in Equation 7. If not estimated, the parameter v is set to the item's Amazon price, and the parameter α is set to 0 for risk neutrality (in which case, Equation 1 is used for f). This results in the four specifications listed in Table 1, including a base specification in which neither parameter is estimated. As in the model, we assume all customers seeking a given item share the same v and α .

In the second step, we perform statistical tests to quantify how closely the estimated theoretical distribution matches the observed sample distribution. This can be done in several ways, such as comparing the sample and theoretical mean (via a t-test). However, a richer and more demanding test would compare the full distribution of ending bids (as in a Pearson's χ^2 goodness-of-fit test or a Kolmogorov-Smirnov (K-S) test). Even if the sample and theoretical distributions have similar means, these distributional tests will reject equality if the relative densities (*i.e.* shape of the distributions) differ by too much in any particular region. These tests are performed for each auctioned item under each of the specifications.

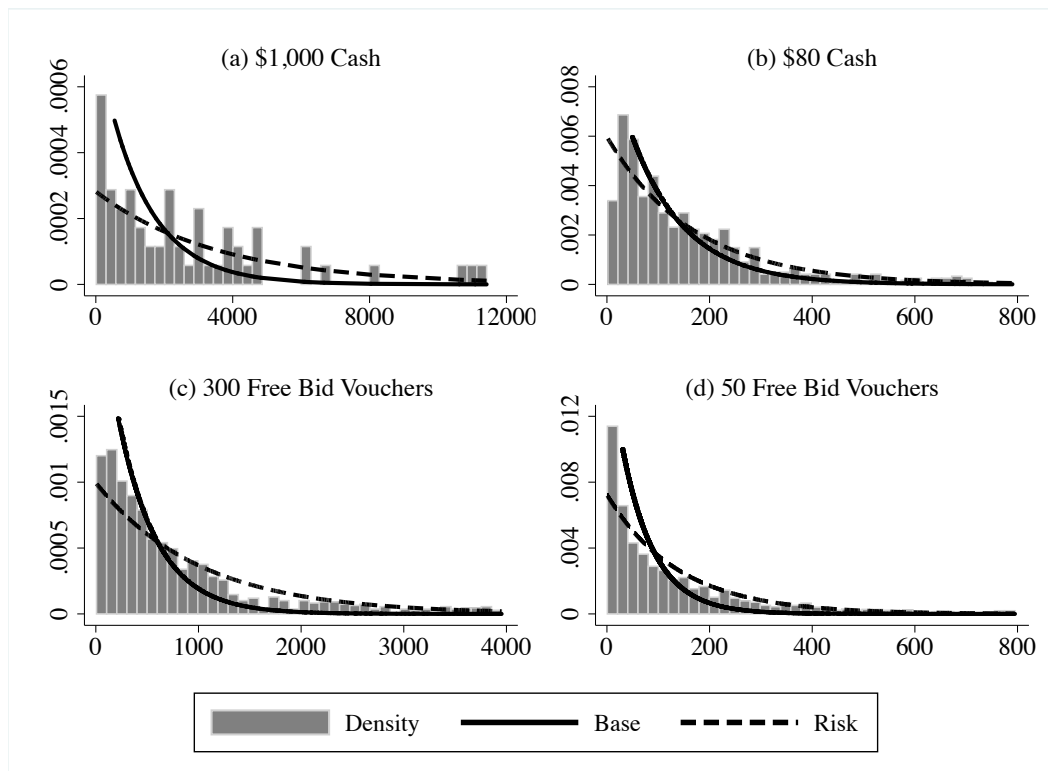
While this two-step process of first selecting parameters and testing fit gives us some flexibility, it does not enable us to shape $f(q)$ at will. For instance, risk neutrality necessarily imposes a decreasing and convex density ($f'(q) < 0$ and $f''(q) > 0$). Introducing risk-loving preferences allows the theoretical density to take a particular hump shape, but is not capable of replicating every unimodal distribution, as we demonstrate in Section 5.3.

This process effectively tests the best-case scenario for the model—its ability to explain the data under the most favorable parameter values. These results are most useful in comparison among the specifications, indicating which parameter is most important in explaining the observed behavior. The remainder of this section examines these comparisons. In addition, one can apply the same approach to other proposed models for pay-to-bid auctions, and again assess their relative ability to explain the observed auctions. We pursue these latter comparisons in Section 5.

4.1. 100% off Auctions

We start by considering the simplest case of auctions in which $s = 0$; that is, the winner pays nothing more than his bid fees. These provide the most tractable theoretical distribution and serve

Figure 2 100% off auctions: theoretical and observed distribution of ending bids, by item.



Notes: The unit of observation is each auction of the given 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. Density (bars) denotes the observed frequencies. Base (solid line) gives the theoretical frequency when $\alpha = 0$, while Risk (dashed line) uses the maximum-likelihood estimate for α . In both cases, the item valuation is set to its retail value.

to illustrate our two-step empirical approach. Four items were frequently auctioned under the 100% off rule, including cash (\$1,000 or \$80) and vouchers for free Swoopo bids (either 300 or 50, worth \$0.75 per bid at the time).

Figure 2 illustrates the observed outcome in a histogram for each item. The numbers along the x-axis are the final number of bids that occurred in an individual auction (which maps proportionally with the revenue from the auction). The y-axis indicates the frequency with which the auction ended at that number of bids.

Note that a large number of auctions end with few bids (and hence low revenue), and the probability of the auction ending at q declines (at a decreasing rate) as q increases. As a consequence, many of these auctions conclude in a net loss for Swoopo—over half of the auctions of a given item will not generate enough revenue (in bid fees) to cover the retail price. However, the long right tail generates enough compensating revenue so that, on average, the bid fees more than cover the cost of providing the cash or free bids.

Next, we quantify how well the model explains this observed behavior. In our *base specification*, we assume risk neutrality and set the valuation to its retail price. However, the theoretical distribution performs rather poorly in each case, placing too much weight on low bids and too little on higher bids. This predicted density function is plotted as a solid line in Figure 2. For each item, it fails all three tests: the observed mean is significantly higher than predicted, and the χ^2 and K-S tests reject equality of the theoretical and observed distributions (at a p-value of 5%).

We next ask if alternative parameterizations can perform any better, chosen via maximum likelihood estimation. However, α is not separately identifiable from v . This is because the distribution has the form $f(q) = (1 - \mu)\mu^q$. Thus, the MLE procedure can only identify μ . Fortunately, all four auctioned items have an obvious objective value, so it seems reasonable to fix v to its retail price. The bid fee b is also exogenously set, so we can use the MLE procedure to identify α . We refer to this as our *risk specification*.

The procedure finds α to be slightly negative, between -0.0017 (for the \$1,000) and -0.03 (for the 50 free bids). This implies that Swoopo participants are risk loving. The resulting theoretical distribution is depicted in Figure 2 by the dashed line, and the improvement in fit is remarkable. With the addition of risk-loving preferences, bidders become a bit more aggressive at every point of the auction. The t-test rejects none of the items, while the χ^2 and K-S tests narrowly reject the two bid vouchers items.

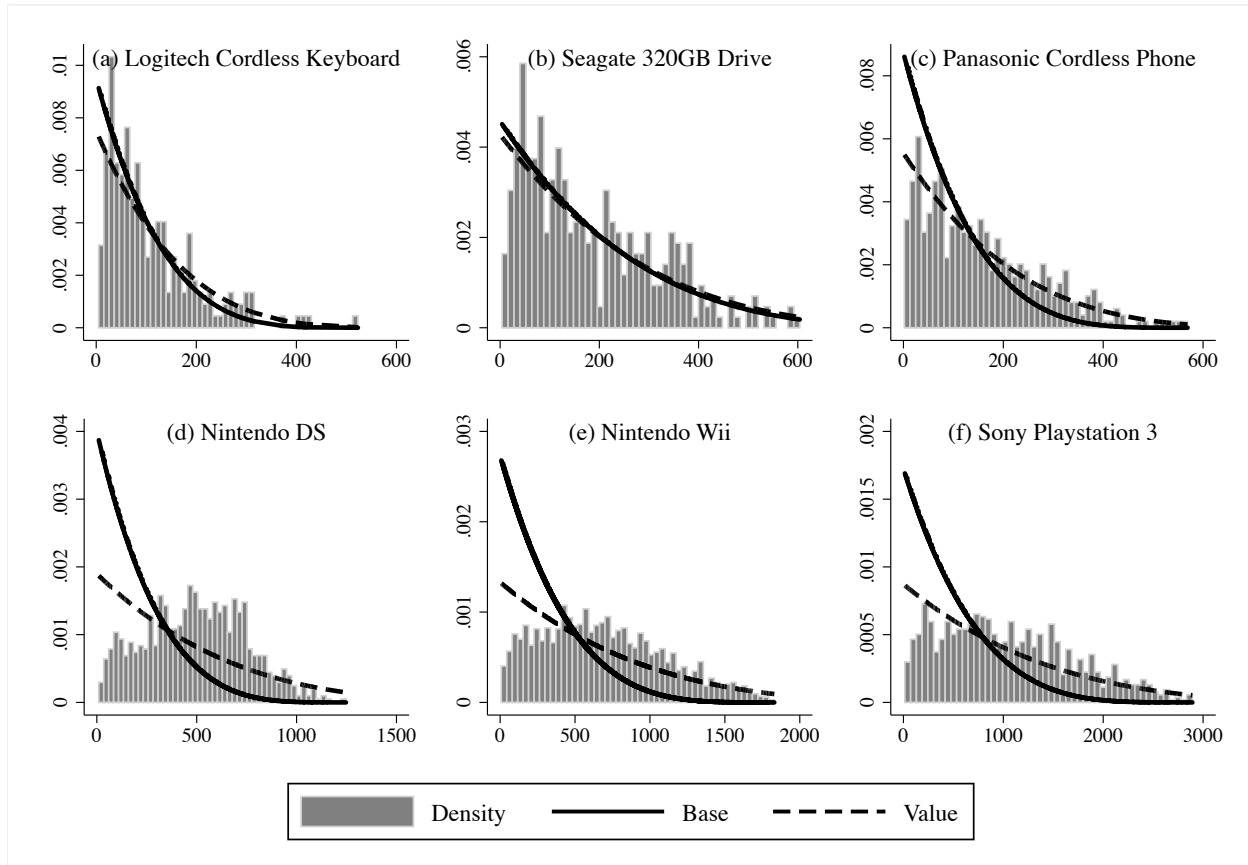
Alternatively, one could assume $\alpha = 0$ and instead identify v from maximum likelihood, which we refer to as the *value specification*. Of course, this estimate produces the same μ and therefore the same degree of fit as the risk specification. However, the estimated valuations are two to three times larger than the retail price. This seems far less convincing than the modest α found in the risk specification, which is much less extreme than estimates for race track bettors.

4.2. Incremental Auctions

We next turn to the far more common auction where $s > 0$, meaning the winner also pays the final price in addition to his bid fees. In this setting, the predicted distribution $f(q)$ depends on four parameters: the valuation of the item, the risk preferences, the bid fee, and the increment by which the final price rises with each bid. The latter two are clearly specified for each auction, while the others will be investigated in our various specifications.

Our data include 172 items which were auctioned more than 100 times. The top row of Figure 3 displays the results from three of the most common non-video game items; the bottom row has three of the most common video game items. Figures (a) and (b) provide visual evidence that the actual distribution of bids (the bar graphs) matches very closely with that predicted by the base specification (the solid lines) of the model.

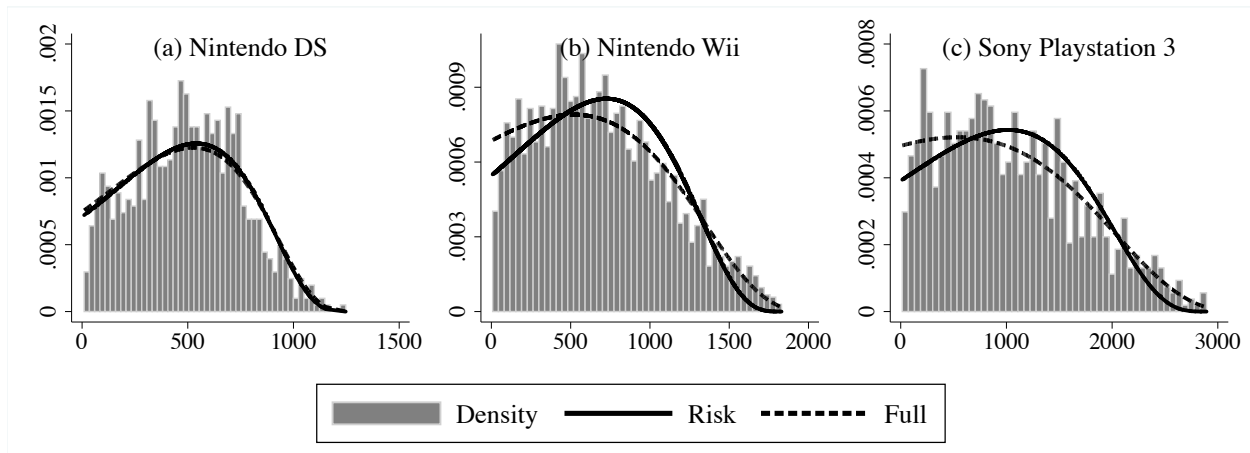
Figure 3 Representative examples: theoretical and observed distribution of ending bids, by item.



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 ($\alpha = 0$).

However, the fit is not as good for other items. For instance, the base-predicted distribution in figure (c) has too much curvature, though the observed density shares much in common with (a) and (b). Items in the lower row of Figure 3 have a much worse fit. In particular, there is an initial range for which the probability is *increasing* in q , which cannot be generated by the risk-neutral model under any parameter values. In addition, the density of ending bids in the low range is much smaller than predicted, and the right tail of the distribution is much thicker.

In the value specification (the dashed lines in Figure 3), risk neutrality is still assumed, but the valuation is found via maximum likelihood. For figures (a) and (b), the valuation, and hence the predicted distribution, is nearly unchanged. However, the fit in figure (c) is greatly improved as a consequence of raising the valuation. At the same time, changes in v offer little improvement for the items in the lower row of Figure 3. Quite simply, the risk-neutral model cannot generate the observed hump shape in the distribution of final prices.

Figure 4 Auctions with risk-loving bidders: theoretical and observed distribution of ending bids, by item.

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 .

However, the introduction of risk preferences offers considerable improvement for these three items, as illustrated in Figure 4. In the risk specification, we fix valuations at the item’s Amazon retail price and estimate α separately for each item. In the *full specification*, α and v are jointly estimated in maximum likelihood. Note that the risk specification (adjusting α alone) provides most of the improvement. Even so, the full specification offers some further improvement in figures (b) and (c), with MLE increasing v to 12% higher than retail.

This goodness of fit plays out similarly among the full set of 172 items. To rigorously quantify the fit and systematically describe the results, we perform three statistical tests for equality of the observed and theoretical distributions of each item. Table 1 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 our four specifications. In effect, this table indicates how often the theory is able to explain observed auctions, depending on how much flexibility is allowed in parameter choices.

We view the t-test in Table 1 as a less demanding standard than the χ^2 and K-S tests, and have mainly included them for comparison purposes. Even so, we note that the maximum likelihood estimation does not guarantee the post-MLE mean will equal the sample mean—the theoretical mean is not a sufficient statistic for the standard model’s predicted distribution, except in a 100% off auction. Also, equality of means indicates that the sample’s average revenue equals theoretical revenue, but outside the base specification, this does not mean the auction earns zero profit. If v exceeds retail price or $\alpha < 0$, predicted revenue will be greater than the retail price.

Table 1 Statistical tests comparing theoretical distribution with observed data.

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

Notes: Each cell reports the percentage of items for which the test does not reject equality, at the p -value indicated in each column. N refers to the number of unique items. K-S refers to the Kolmogorov-Smirnov test.

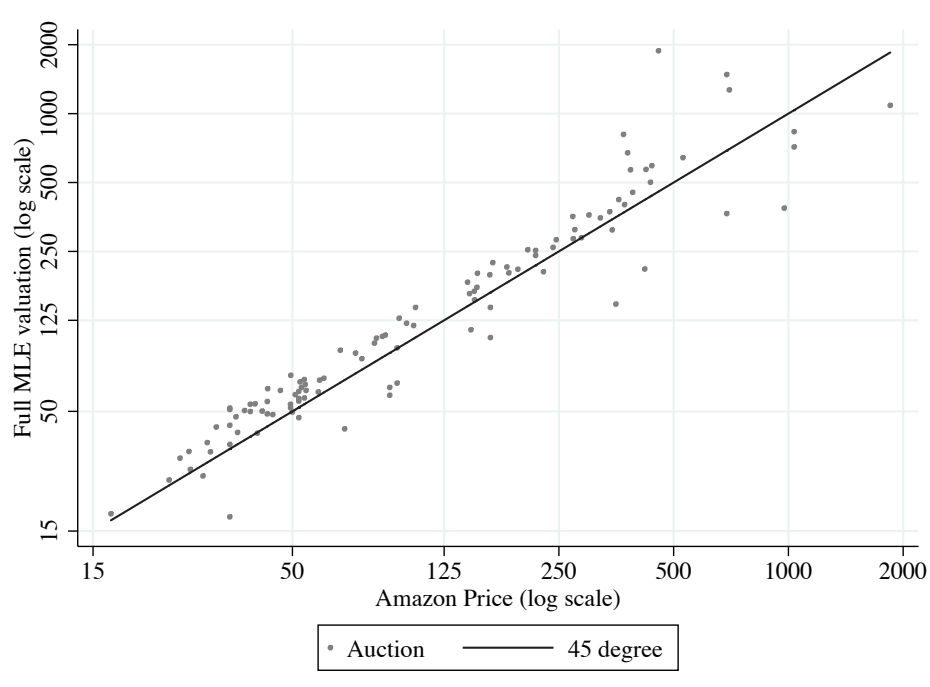
The first row of Table 1 reveals that the base specification performs rather poorly under all tests. For instance, using the .05 p -value as a threshold for a K-S test to compare the theoretical and observed distribution, we reject equality in all but 10.5% of the items. That is, only 10.5% of the sample items look similar to (a) and (b) in Figure 3.

Under the value specification, our model’s goodness of fit dramatically increases. The second row of Table 1 indicates that 93 of the 172 items (or 54%) are consistent with the predictions of value specification.⁷ Relative to the base specification, the value specification can explain an additional 43% of the auctioned items. These include items across many categories and retail prices, though notably absent are the more expensive home electronic and video game items (which mostly resemble the bottom row in Figure 3).

The third row of Table 1 examines the risk specification. The improvement in fit over the base specification is remarkable, as nearly 70% of the items are consistent with the risk specification. Comparing the risk and value specifications, it is noteworthy that incorporating risk preferences (while setting valuations to Amazon prices) explains 15% more of the observed auctions than adjusting valuations (with risk-neutral bidders).

Moving from the risk to full specification, reported in the fourth row of Table 1, an additional 17% of items are not rejected in the K-S test. Of these 33 items, 19 have a retail price under \$100. Thus, having both degrees of freedom allows some fine tuning of the model, particularly for lower priced items; but most of the work is performed by the risk preferences.

⁷ As a robustness check, we also generated the equivalent of Table 1 restricted to items that were auctioned at least 200 times (instead of 100 times). This reduces the number of items by 60%, but the results are nearly identical. On the other hand, if we include all items auctioned at least 45 times (which doubles the number of items), the fraction of items that are rejected in any of the tests falls by 10 percentage points. Of course, as the number of observations per item falls, the tests become less powerful (and thus fail to reject more often). Requiring at least 100 observations seems to be a reasonable (though perhaps conservative) threshold.

Figure 5 Comparison of Amazon price to estimated valuation in the full specification.

Notes: Each point compares an item's Amazon.com price to its MLE-estimated valuation in the full specification. For reference, a 45 degree line is included. On average, v is 15% higher than the Amazon price, with an interquartile range of 4% to 28%. Only 5% of items had v more than 50% of the Amazon price.

In terms of how reasonable their estimated parameters are, the latter three specifications are not equal. In the value specification, the MLE valuation is higher than the Amazon price for all but 8 items, with an average increase of 65%. These valuations seem implausibly high, but they also indicate that Swoopo's average profits were high, though subject to enormous variance. In the full specification, however, the estimated valuations are reasonably close to the Amazon prices (on average, 15% greater), as illustrated in Figure 5.

In both the risk and full specifications, the estimated risk parameter α indicates that bidders are mildly risk loving, primarily in the range of -0.001 to -0.03, with a few estimates as low as -0.09.⁸ Most were not as extreme as the estimated risk preferences of bettors at horse race tracks; only 9 items in the risk specification (and 3 in the full) had an $\alpha < -0.055$.

Another way to interpret α is to compute the expected percentage return from placing the q^{th} bid: $\frac{(1-\mu_{q+1})v-b}{b}$. In a risk-loving environment, the return is lowest (and negative) at the beginning

⁸ Maximum likelihood found $\alpha > 0$ (*i.e.* risk averse) for 23 items in the risk specification. In half of these, the fitted distribution was still rejected as being different from the data. In the full specification, 62 items obtain an $\alpha > 0$, and these are associated with much larger v . In all but 5 of these items, the full specification adds little improvement over the risk specification. Thus, we revert to the risk specification for these items with positive α .

of the auction and draws steadily closer to 0 as q nears Q . To provide a lower bound, we compute the return on the first bid using our estimated risk parameter for each item. The average return was -54%, with estimates across items distributed from -88% to 0%. While these may seem low, Jullien and Salanié (2000) report a -50% return for race track bets with 20-to-1 odds, which decreases with worse odds. The typical first bid on a Swoopo auction has 100-to-1 odds of winning.

Comparing across these four specifications, then, we conclude that the base specification is clearly inadequate, and that risk preferences contribute the most 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, as its estimated parameters are quite reasonable.

5. Discussion

Since the preceding analysis relies heavily on our two-step testing procedure, we provide further discussion of its merits and interpretation in this section. First, we examine the nature of risk-loving preferences, which could explain why α varies across items. Second, we quantify the prevalence of the hump shape. Next, we apply the same methodology to alternative models of pay-to-bid auctions. Finally, we examine some more traditional tests of the model's comparative statics.

5.1. Degrees of Freedom in 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 via one of two routes: by assuming some intrinsic utility from winning via a gamble (expressed here as v greater than retail) or by assuming convex preferences with respect to wealth (expressed using $\alpha < 0$).⁹ From our previous section it appears that the latter is more essential in explaining pay-to-bid auctions: the risk specification was more effective than the value specification, and even the full specification typically used a modest 15% joy-of-winning premium.

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 conclude that utility from gambling would necessarily contradict stochastic dominance.

⁹ Though beyond the scope of this paper, one could replace expected utility with other preferences from behavioral economics. Among these, prospect theory (Kahneman and Tversky 1979) seems most applicable since it can generate risk seeking behavior in cases of low probability wins.

For the alternative, Friedman and Savage (1948) provide an early formal model of increasing marginal utility of wealth within some range. More recently, Golec and Tamarkin (1998) and Garrett and Sobel (1999) propose 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 enjoy utility from skewness in payoffs. Golec and Tamarkin show that this will appear as risk-loving preferences when estimating parameters for CRRA utility.

One could interpret this love of skewness in several ways. For instance, bragging rights are greater if an item is won at a very low price than at close to retail. Alternatively, bidders could see this as a means of overcoming the indivisibility of the item being sold (Kwang 1965, Hartley and Farrell 2002). That is, one cannot buy one tenth of a Wii (or a \$30 substitute that provides one tenth of its enjoyment). Thus, bidders can justify a \$0.75 gamble, even at unfair odds, as it provides access to a valuable lottery.

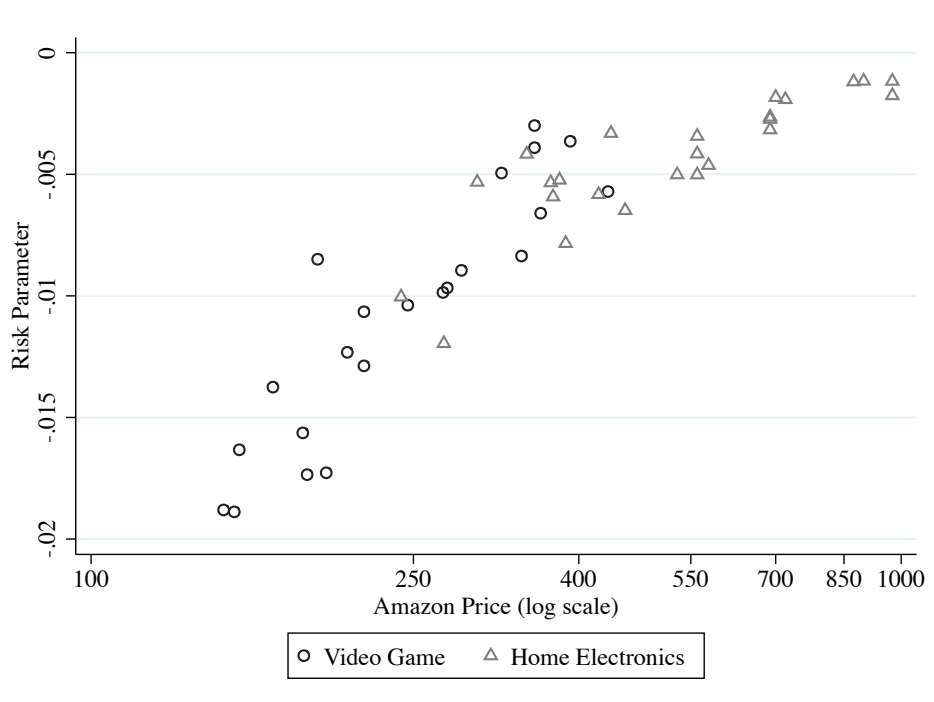
In either case, as the number of bids (and hence the final price) increase, the skewness in outcome decreases and these gamblers are less tolerant of unfair odds—the winner has less to brag about or hasn't reduced the cost of the Wii by much. Indeed, as the final price approaches the retail price, μ_q eventually adjusts so that bidding becomes a fair bet.

Given the interpretations above, it is unsurprising that bidders would have a different risk parameter for different items. One is less likely to brag about great savings on a DVD than on a gaming system. Similarly, estimates of α vary across race track bets; Golec and Tamarkin (1998) finds the magnitude of the risk parameter to be twice as big when estimated for bets on long shots as compared to bets on favorites.

This suggests, however, that the relationship between the risk parameter might be systematically related to the price or category of the item. When restricted to home electronic and video game items whose price lies between \$100 and \$1,000 dollars, the relationship is very nearly log-linear, as illustrated in Figure 6. Indeed, regressing the log of the Amazon price on the risk specification α produces an $R^2 = 0.85$ among video games ($R^2 = 0.77$ among home electronics). Items priced over \$1,000 followed nearly the same relationship, though the risk specification was unable to explain the majority of these. For those under \$100, the relationship is a bit steeper but much more noisy (with an $R^2 = 0.09$). Also, the ten items not included in Figure 6 (in more practical categories of computer accessories and home appliances) had estimated risk preferences near zero.

One potential explanation for a log-linear relationship is that the utility function is misspecified. We have chosen the functional form $u(w) = \frac{1-e^{-\alpha w}}{\alpha}$, which provides an absolute risk aversion of α , precisely the parameter we are estimating. However, if bidders actually have CRRA preferences, $u(w) = aw^\gamma$, their absolute risk aversion is $\frac{1-\gamma}{w}$, which increases with wealth (assuming $\gamma > 1$, as

Figure 6 Comparison of Amazon price to estimated risk parameter in the risk specification.



Notes: Each point compares an item’s Amazon.com price to its MLE-estimated α in the risk specification.

required for risk-loving behavior). Bidders on more expensive items likely have larger budgets (w); our procedure would identify this as an increase in α even if γ remains constant. Of course, one could use this CRRA utility form instead, but this would introduce the average wealth parameter w which is unknown and may not be constant across items.

Thus, allowing α to vary across items can be seen as a means of compensating for many omitted features in our model. Indeed, it is important to have this flexibility in our analysis. If we impose a log-linear restriction between the log price and α while estimating the risk specification on the items included in Figure 6, we explain only half as many items as in the unrestricted model. Some small deviations from the literal log-linear relationship are needed to replicate the observed shape.

5.2. Prevalence of the Hump Shape

Our empirical testing procedure essentially is a question of whether the theory can accommodate certain distributional shapes. Assuming risk neutrality imposes a downward-sloping, convex density function regardless of the parameter v . With risk-loving customers, $f(q)$ is concave for low values of q and can even introduce a hump shape.

Intuitively, the hump-shaped arises because when q is low, there is greater variance in the outcome. In particular, if no one bids, the payoff, $v - sq$, is large. A risk-loving participant will be

willing to bid in spite of unfair odds as long as the gamble has a highly skewed payoff. Thus μ_q can be larger (yielding more aggressive bidding) than in the risk-neutral model while still satisfying our indifference condition, 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 approaches the risk-neutral result.

As discussed in the appendix, $f(q)$ reaches its peak at approximately $\hat{q} \equiv \frac{v-b}{s} + \frac{1}{\alpha s} \ln\left(\frac{1-e^{\alpha b}}{1-e^{\alpha s}}\right)$. If the calculated $\hat{q} \notin [0, Q]$, then the distribution is strictly downward sloping. Using parameters estimated in the full specification, we find 31 items (19% of those in the sample) are predicted to have a hump shape. In the value specification, the K-S test rejected 29 of these items; in the full specification, only three are narrowly rejected. Descriptively, 20 of these are video game items, with the others split between home electronics and computer accessories. Their Amazon prices range from \$30 to \$1,035, though two-thirds of them lie in between \$90 and \$300. These constitute over half of the items of similar price and category.

As an alternative test of whether the density function is increasing in its initial range, we create a histogram of ending bids for each item, and test for a positive initial slope as follows. First, we assign ending bids into 30 equally-spaced bins. Treating each bin as an observation, we let our dependent variable be the number of auctions that concluded in that bin. We then estimate the following regression for each item separately:

$$auctions = \beta_0 + \beta_1 \cdot bin + \beta_2 \cdot post \cdot bin + \beta_3 \cdot post + \epsilon.$$

Here, bin indicates which of bins 1 to 30 the auction ended in, while $post$ is a dummy variable indicating if bin occurs above the peak \hat{q} predicted by the theoretical model.

This regression strongly supports the shape predicted under the MLE parameters. The distribution is upward-sloping (i.e. $\beta_1 > 0$) up to the predicted peak for 28 of the 31 items with a predicted hump shape; this coefficient is statistically significant at the 5% level for 23 of these items. The coefficient is not significant for three items where $\beta_1 < 0$. In addition, the interaction term (β_2) is negative in the same 28 items, indicating a change in the slope as predicted by the model, significantly for both coefficients in 16 items.

In addition to these items with a hump-shaped distribution, another 28 items were rejected in the value specification but not in the full specification. Among those, the improvement in fit occurs because the density function is concave for low q ; this creates a bulge in what would otherwise look much like the exponential distribution. This feature is replicated by $f(q)$ when α is negative but near 0. The change in the sign of the second derivative is subtle, and cannot be tested with any significance in the manner described above.

As a robustness check, we repeated this analysis on the pay-to-bid website StiWin.com, where data was available on 95 items that were auctioned at least 100 times. After maximum likelihood estimation of our full specification, 34 of these were predicted to have a hump shape. These had similar categories but somewhat higher retail prices compared to items in our Swoopo data.

5.3. Comparison to Other Models

Instead of asking the model to fit a broad variety of items with fewer degrees of freedom, an alternative approach is to ask how other models perform when given the same degrees of freedom. Thus, we apply our same methodology to models proposed by Augenblick (2009) and Byers *et al.* (2010). We also apply the same approach to common statistical distributions, even without a theoretic basis for expecting pay-to-bid auctions to follow a particular distributional form.

We start with the latter, proposing six atheoretic distributions shaped by two parameters, listed in Table 2. For each, we follow our two-step process of first choosing both parameters via maximum likelihood, then testing the goodness of fit. For brevity, t-tests are omitted. Note that we include our full specification as a benchmark for comparison.

The Uniform, Laplace, and Normal distributions perform quite poorly. The Log-Normal distribution fares better, with its hump and long right tail, but still fails to replicate the correct curvature in many cases. The final two distributions provide the greatest flexibility, with the ability to have either a hump shape or be strictly decreasing in q . The Gamma distribution performs slightly worse than our full specification, while the Weibull distribution performs slightly better. This is not surprising considering that these two distributions are each members of the GB1 family of distributions (Equation 8, discussed in the appendix), as is our $f(q)$. Relative to our distribution, these have more flexibility in being able to choose a (Weibull) or ρ (Gamma) in Equation 8 arbitrarily, creating the hump shape.

The Gamma or Weibull parameters obtained are quite varied across the items and have no obvious interpretation due to their atheoretic nature. We view it as a success that our more disciplined and intuitive model can perform similarly to these more familiar cousins. Moreover, our theoretical distribution is sufficiently constrained that it cannot replicate every such distribution. To demonstrate this, we created a simulated dataset with 300 observations (the same as our average auctions per item) from a Weibull distribution. Using our two-step process, we attempted to fit our full specification to this data and tested its goodness of fit. This was repeated 180 times with distinct parameters k and λ drawn uniformly from $[0.5, 3] \times [15, 3000]$ (the same parameter range found when fitting the Weibull distribution to the actual auction data). In two-thirds of the cases, the KS test rejected equality. That is, our procedure has reasonable power to avoid Type II errors; random Weibull data is not typically consistent with our theoretical pay-to-bid behavior. The results are stronger for the other atheoretic distributions.

Table 2 Statistical tests comparing alternative distributions with observed data.

Specification	$f(q)$	N	χ^2 test		K-S test	
			$p \geq .10$	$p \geq .05$	$p \geq .10$	$p \geq .05$
Full		169	73.4	76.9	82.2	87.0
Uniform	$\frac{1}{b-a}$	172	0.6	1.7	0.0	0.0
Laplace	$\frac{1}{2b} e^{-\frac{ q-\mu }{b}}$	172	13.4	18.6	10.5	15.7
Normal	$\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(q-\mu)^2}{2\sigma^2}}$	172	15.6	21.0	18.6	25.6
Log-Normal	$\frac{1}{q\sigma\sqrt{2\pi}} e^{-\frac{(\ln q - \mu)^2}{2\sigma^2}}$	172	36.0	42.4	50.0	59.9
Gamma	$\frac{1}{\theta^k \Gamma(k)} q^{k-1} e^{-(q/\theta)}$	172	73.3	79.1	76.2	84.9
Weibull	$\frac{k}{\lambda^k} q^{k-1} e^{-(q/\lambda)^k}$	172	78.5	84.3	88.4	89.5
Bid fee		170	50.0	56.5	50.6	55.9
Extra Bidders		172	65.7	73.3	70.9	75.6
Sunk Costs		168	65.5	73.2	72.0	76.8

With this in mind, we now turn to three other models of pay-to-bid auctions, proposed by Byers *et al.* (2010) and Augenblick (2009). These each share our base specification, but augment it with plausible behavioral errors by bidders. First, Byers *et al.* (2010) propose that perhaps the marginal bidder can obtain bid fees for cheaper than their face value, due to introductory discounts or winning them through the auctions themselves. This maintains the same indifference condition as before, just with the effective bid fee. To test this *bid fee specification* we select the valuation and the bid fee using maximum likelihood. Once tested, this explains only 56% of the items (as reported in Table 2). Note that this is only three more items than our value specification, and far below our full specification. Changing the bid fee offers almost no benefit because $f(q)$ in Equation 1 is downward sloping and concave regardless of b , and cannot produce a hump shape.

Another approach offered in Byers *et al.* (2010) is that bidders systematically underestimate the number of auction participants. Supposing n is the true number of participants, bidders expect that only $n - k$ are participating. As a consequence, $\mu_q = 1 - \left(\frac{b}{v-s \cdot q+s}\right)^{(n-1)/(n-k-1)}$. Using this in $f(q)$, we normalize $n = 1000$ and estimate both v and k via maximum likelihood. This *extra-bidders specification* performs better than the bid fee specification; equality of the observed and theoretical distributions is not rejected for 76% of the items (a strict subset of the items not rejected in our full specification). This is primarily because $f(q)$ can become convex when k is sufficiently large, though not enough to produce a hump shape. This can help in fitting some of the items, with the estimated k reaching as high as 300. On the other hand, k is negative for 63 items, meaning fewer

bidders participate than expected—though for most of these items, estimating v alone is sufficient to explain the bidding.

Another interesting theory, proposed by Augenblick (2009), augments the base specification by making bidders regret their unsuccessful bids, with the level of regret expressed by $\rho \in [0, \infty)$. Thus, after the fact, a customer treats past bids as if they cost $(1 + \rho)b$ instead of b . Customers also have time inconsistency, in that they discount future regret by $\eta \in [0, 1]$, erroneously expecting that tomorrow they will treat past bids as if they cost $(1 + \rho(1 - \eta))b$.

Imposing the typical indifference condition, one obtains $\mu_q = 1 - \frac{b(2+(2-(1+q)\eta)\rho)}{2v-2(q-1)s+b(2+q)(1-\eta)\rho}$ when $q < \frac{2(v-b)}{b+2s}$, and $\mu_q = 1 - \frac{b(2+(2-(1+q)\eta)\rho)}{2v-2(q-1)s}$ otherwise. Again, we substitute these into $f(q)$ and use maximum likelihood to estimate v and ρ . We can also estimate η , though the MLE program was more temperamental when doing so; here, we hold $\eta = 0.0044$, the average value obtained when all three parameters were successfully estimated. At first glance, this *sunk cost specification* performs reasonably well, explaining nearly 77% of the auctioned items (though still a strict subset of items which were consistent with our full specification).

Deeper inspection, however, raises a concern about how this fit is achieved. In 65% of the items (and 66% of those which pass the K-S test), MLE produced a negative value for ρ (averaging -0.4 among these items), even though the theory suggests ρ should be positive. Taken literally, $\rho = -0.4$ would mean that customers treat past bids as if they were 40% *cheaper* than they really were, yet anticipate that if they place another bid, the bid fee will be 39.8% cheaper than it really is. MLE resorts to these negative values in an attempt to replicate the hump shape; with $\rho \geq 0$, $f(q)$ slopes strictly downward (though with a discontinuity at $q = \frac{2(v-b)}{b+2s}$). Also, the resulting valuations are much larger than those from our full specification, averaging 68% above the Amazon price, even when high-priced items are excluded.

In summary, when these plausible theories are given the same flexibility in setting two parameters per item as our full specification, they perform moderately well, but still only explain a strict subset of the items compared to our full specification. Our two-step testing procedure provides latitude by estimating parameters item-by-item, but this flexibility is limited by the underlying theory, and thus has power to distinguish between alternative models.

5.4. Comparative Statics

A traditional empirical test is to examine comparative statics of the model. For our theory, this would mean comparing a given item that was auctioned with different bid fees b or different increments s . Bid fees only changed once, dropping from \$0.75 to \$0.60 on July 6th, 2009; at the same time, the standard increment fell from \$0.15 to \$0.12. In a risk-neutral specification, the model predicts this change would have no effect on average revenue. Once risk preferences are included, the

prediction must be numerically evaluated; but we consistently found that a proportional decrease in b and s will decrease revenue if $\alpha < 0$ and increase revenue if $\alpha > 0$.

Fourteen items were auctioned over 100 times both before and after the change; revenue decreased as predicted in 12 of these items. The other two were not significantly positive (at a p-value of 5%); also the revenue decrease was not significant for one additional item.

Even so, this test is potentially biased because the auctions were spread out over a year on either side. Over that two-year span, the Amazon price dropped by more than 20% in 8 of these items.¹⁰ Thus, v is also likely dropping which would also reduce revenue. To address this, we could restrict ourselves to the four items whose price stayed nearly constant, all of which saw a significant decrease in revenue. Alternatively, we could restrict our timeframe to the two months before and after the change, so prices remain stable. Unfortunately, this dramatically reduces our sample size (to single digits, in some cases), and only one of the items has a significant result.

After the drop in bid fees, Swoopo also experimented with several price increments simultaneously. For instance, one Wii video game was auctioned over 100 times both with $s = 0.01$ and $s = 0.12$; the same occurred for a PS3 video game, as well as with $s = 0.24$. A Nikon camera was auctioned with both $s = 0.02$ and $s = 0.06$. This provides us with five pairs to test (three from the PS3 game). Since these auctions occurred contemporaneously, we have no concern about changing valuations. In numerically-derived comparative statics, an increase in s will decrease average revenue if and only if $\alpha < 0$. This is borne out in the empirical tests; all five show that revenue was lower with a higher price increment, significantly so in all but two (falling just short in one, with a p-value of 7.6%).

6. Conclusion

This paper presents a parsimonious model of rational bidders in a pay-to-bid auction. In the 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 bidders' valuation of the auctioned item; if bidders are risk loving, expected revenue is even higher.

The model's ability to explain the observed data largely depends on how parameter values are chosen for the bidder's risk preferences and item valuation. We begin by assuming risk neutrality and a valuation equal to Amazon's retail price for the item (resulting in a poor fit for most items). We then use maximum likelihood to estimate one or both of these parameters. From this process,

¹⁰ Note that this is much less of a concern in our Section 4 analysis. There, items were counted as distinct if their bid fees differed. Since the fee changed halfway through our sample period, the longest time between the first and last auction of an item was 251 day, with an average of 128 days.

we conclude that risk preferences are the most important factor in explaining bidder behavior, improving the fit much more than adjusted valuations (*i.e.* the joy of winning) can alone.

In sum, pay-to-bid auctions are essentially a form of gambling; thus, it is not surprising that participants bear some resemblance to gamblers from other settings. On a broader level, the pay-to-bid auction describes an incremental king-of-the-hill contest. By incurring a sunk cost (*e.g.* bid fee), anyone can become the current king-of-the-hill; yet that title only becomes permanent if all challengers give up. The contest is incremental because each replacement of a king reduces the hill's value to the eventual winner.

This could describe a particular form of competition among rent seekers. Suppose several firms were seeking the same exclusive license from a bribe-accepting regulator. The regulator could require an up-front bribe each time a firm wishes to make an overture; moreover, to displace the previous overture, the current firm promises to return a greater portion of the future rents to the regulator. The license is awarded once no additional overtures are attempted, and is given to the firm with the last (and hence best) overture. A similar story could be told for competition among suitors, showering gifts or attention on a potential mate.

Applied to these situations, our risk-neutral model would predict that the regulator captures essentially the full value of the license (in expectation). The firms would be indifferent about participation *ex-ante*; yet *ex-post*, the winning firm will typically enjoy large rents, with a final price well below the full value. If the firms are risk loving, the regulator can expect to extract even more than the full value, making this a far more profitable means of allocating the license than other auction formats. Indeed, risk-loving (or skewness-loving) preferences might well be applicable if the residual rents would significantly alter the winner's social class or if, in the case of a suitor, the contest gives him a shot at a mate normally far outside his league.

Appendix

Proof of Proposition 1. Consider a customer in the $q - 1^{\text{th}}$ period. Placing a bid is only useful to her if she wins any tie this period, followed by no bidding in the q^{th} . Taking as given the symmetric strategy β_{q+1} , the probability that none of the other customers places a bid in the next period is $(1 - \beta_{q+1})^{n-1} = 1 - \mu_{q+1}$. Thus, if she bids and wins any ties in period $q - 1$, her expected payoff entering period q is $w - b + (v - s \cdot q)(1 - \mu_{q+1}) = w$.

Her payoff from not bidding is also w ; indeed, becoming the current winner in any future period also offers a payoff of w , so there is no option value of waiting. Thus, she is indifferent between becoming the current winner versus abstaining from bidding, and can employ mixed strategy β_q .

Note that a customer in the $Q - 1^{\text{th}}$ period knows that if she becomes the new highest bidder, the new price $p_Q = s \cdot Q$, so after substituting for Q and $\mu_{Q+1} = 0$, her expected payoff would be $w + b(1 - \mu_{Q+1}) - b = w$. Thus, even though no one will bid after her, she is still indifferent because the value of winning barely covers the bid fee. In the Q^{th} period and beyond, the new current price, p_q , exceeds the value of the good minus the bid fee, so all customers strictly prefer not to bid.

In period 0, customers are also indifferent about becoming the current winner for the same reasons depicted above. Thus, any mixed strategy β_1 may be employed. For $q > 1$, β_{q+1} is pinned down by the need to make customers in period $q - 1$ indifferent about placing a bid; but since there is no period $q = -1$, this does not apply to β_1 .

Proof of Proposition 2. Suppose that there exists a period q where customers randomize with $\beta_q > 1 - \left(\frac{b}{v-s \cdot (q-1)}\right)^{\frac{1}{n-1}}$, with customers following the strategies in Proposition 1 for all periods beyond q . Note that this randomization is a best response, since customers are indifferent about bidding in q , given the strategies β_{q+1} . However, this will make customers in period $q - 1$ strictly prefer not bidding. Their expected utility from doing so would be:

$$\begin{aligned} & (w + v - b - s \cdot (q - 1)) \cdot (1 - \beta_q)^{n-1} + (w - b) \cdot \left(1 - (1 - \beta_q)^{n-1}\right) = \\ & w - b + (v - s \cdot (q - 1)) \cdot (1 - \beta_q)^{n-1} < w - b + (v - s \cdot (q - 1)) \cdot \left(\frac{b}{v - s \cdot (q - 1)}\right) = w, \end{aligned}$$

while not bidding gives them w . Thus, only $\beta_{q-1} = 0$ can occur.

This in turn requires $\beta_{q-2} = 1$. Everyone strictly prefers an attempt to bid in period $q - 2$ because no one will bid in $q - 1$. Thus, winning the tie in period $q - 2$ means winning the auction. Of course, this means everyone strictly prefers not bidding in $q - 3$, and this continues to alternate over earlier periods. That is, for any positive integer t where $0 \leq t \leq q/2$, $\beta_{q-2t} = 1$ and $\beta_{q-2t-1} = 0$. If q is even, the equilibrium outcome concludes with a single bid (since $\beta_1 = 0$). If q is odd, the outcome concludes with no bids (since $\beta_0 = 0$).

Alternatively, customers in period q could randomize with $\beta_q < 1 - \left(\frac{b}{v-s \cdot (q-1)}\right)^{\frac{1}{n-1}}$. Similar logic would require $\beta_{q-2t} = 0$ and $\beta_{q-2t-1} = 1$, and merely reverses the outcome for even or odd q .

Of course, these equilibria can occur starting at any q , provided that customers follow the Proposition 1 strategy for all periods beyond q . This set of strategies exhausts the possibilities for equilibria; if indifference is broken at any period, all earlier periods must follow this alternating strategy.

Continuous Approximations of $f(q)$. In the standard model with risk-neutral customers, $f(q)$ can be approximated by treating q as a continuous rather than discrete variable. Here we show that this results in a particular generalized beta distribution of the first kind. Thus, it is no surprise that our model's predicted distribution bears some resemblance to other distributions in that family.

First, note that $f(q+1) = \frac{v-b-sq}{v-s-sq} f(q)$, since the product term in $f(q)$ also appears in $f(q+1)$. We then approximate $f'(q) \approx f(q+1) - f(q) = \frac{b-s}{v-s-sq} f(q)$, which is a close approximation as long as s is very small relative to v .

This differential equation has the unique solution: $f(q) = b(v-s)^{-\frac{b}{s}} (v-s-qs)^{\frac{b}{s}-1}$, where the constant of integration is determined such that $\int_0^{\frac{v-s}{s}} f(q) dq = 1$. This pdf is a special case of the Generalized Beta distribution of the first kind (*GB1*). The *GB1* distribution has four parameters: a peak parameter a , a range parameter ρ , and two shape parameters γ and δ :

$$GB1(q; a, \gamma, \delta, \rho) = \frac{|a| q^{a\gamma-1} \left(1 - \left(\frac{q}{\rho}\right)^\alpha\right)^{\delta-1}}{\rho^{a\gamma} B(\gamma, \delta)}, \quad (8)$$

where $B(\gamma, \delta)$ is the Beta function. To match our distribution, the parameters are set with $a = 1$, $\gamma = 1$, $\delta = b/s$, and $\rho = (v-s)/s$.

The *GB1* distribution becomes the generalized Gamma distribution as $\delta \rightarrow +\infty$, leaving other parameters free. By setting $\gamma = 1$, one obtains the Weibull distribution. Setting $a = 1$ instead yields the Gamma distribution (McDonald 1984). Under our parameterization, *GB1* becomes the exponential distribution as $s \rightarrow 0$, which is the continuous approximation of the geometric distribution derived in the text for $s = 0$.

As a simple verification, note that using this distribution, the expected revenue computes to be $v-s$. We can also find the variance more easily, which is: $\frac{b}{b+2s}(v-s)^2$.

When the model is augmented with risk-loving preferences, we use the same procedure to approximate $f(q)$. Here, $f'(q) \approx \left(\frac{1-e^{\alpha(b+sq-v)}}{e^{\alpha(b+sq-v)}-e^{\alpha b}} - 1\right) f(q)$. The solution to this differential equation is:

$$f(q) = \frac{e^{\alpha v - q(1-e^{-\alpha b})} (1-e^{-\alpha b}) \left(\frac{e^{\alpha v} - e^{\alpha s(1+q)}}{e^{\alpha v} - e^{\alpha s}}\right)^{\frac{1}{\alpha s} (e^{-\alpha s} - e^{-\alpha b})}}{(e^{\alpha v} - e^{\alpha s}) {}_2F_1\left(1, 1 - \frac{1-e^{-\alpha s}}{\alpha s}, 1 - \frac{1-e^{-\alpha b}}{\alpha s}, \operatorname{sgn}(\alpha) e^{-\alpha(v-s)}\right)}, \quad (9)$$

where ${}_2F_1(\cdot)$ is the hypergeometric function and $\operatorname{sgn}(\cdot)$ is the sign function. This distribution does not allow analytical derivations of its mean and variance, but can be useful for numerical computation.

This is equivalent to an Exponential Generalized Beta distribution of the first kind (*EGB1*), with support truncated to $[0, \frac{v-s}{s}]$ rather than $(-\infty, \frac{v-s}{s}]$. Again, as $s \rightarrow 0$, this approximation of $f(q)$ becomes the exponential distribution. Also, the generalized Gompertz distribution is a special case of *EGB1*, obtained as the δ parameter ($\delta = 1 - \frac{1}{\alpha s} (e^{-\alpha b} - e^{-\alpha s})$, in our case) approaches $+\infty$ (McDonald 1995).

When α is sufficiently negative, $f(q)$ is initially increasing, reaching its peak at $\hat{q} \equiv \frac{v-b}{s} + \frac{1}{\alpha s} \ln\left(\frac{1-e^{\alpha b}}{1-e^{\alpha s}}\right)$, after which it is strictly decreasing.

Proof of Proposition 3. If we compare two distributions, F and \tilde{F} , produced with differing parameters, F is first-order stochastically dominated by \tilde{F} if $F(q) > \tilde{F}(q)$ for all $q < Q$.

The cumulative distribution, $F(q) \equiv \sum_{j=0}^q f(j)$, simplifies due to a telescoping sum. By induction, note that $F(0) = 1 - \mu_1$ and for $q \geq 1$,

$$F(q) = F(q-1) + f(q) = 1 - \prod_{j=1}^q \mu_j + (1 - \mu_{q+1}) \prod_{j=1}^q \mu_j = 1 - \prod_{j=1}^{q+1} \mu_j.$$

If $\mu_j < \tilde{\mu}_j$ for all j , then $\prod_{j=1}^q \mu_j < \prod_{j=1}^q \tilde{\mu}_j$ for all q , and F is first-order stochastically dominated by \tilde{F} .

We then consider the comparative statics of each parameter on μ_q . First, consider α :

$$\frac{\partial \mu_{q+1}}{\partial \alpha} = \frac{e^{\alpha(b-v+qs)}}{(e^{\alpha b} - e^{\alpha(b+qs-v)})^2} (b(1 - e^{\alpha(v-qs)}) - (1 - e^{\alpha b})(v - qs)) < 0.$$

The fractional term is always positive. The last term is 0 at $\alpha = 0$ or $q = Q$, and negative otherwise. To see the latter fact, the derivative of $b(1 - e^{\alpha(v-qs)}) - (1 - e^{\alpha b})(v - qs)$ w.r.t. α is: $-b \cdot (v - qs) \cdot (e^{\alpha(v-qs)} - e^{\alpha b})$, which can only equal 0 if $\alpha = 0$. This is a maximum, since the second-order condition evaluated at $\alpha = 0$ is negative: $-b \cdot (v - qs) \cdot (v - qs - b) < 0$ for $q < Q$. Thus, $\frac{\partial \mu_{q+1}}{\partial \alpha} < 0$ for all $q < Q$ and $\alpha \neq 0$.

The other two derivatives are:

$$\frac{\partial \mu_{q+1}}{\partial s} = \frac{(1 - e^{\alpha b}) q \alpha}{e^{\alpha(b-qs-v)} (e^{\alpha qs} - e^{\alpha v})^2} < 0 \quad \text{and} \quad \frac{\partial \mu_{q+1}}{\partial b} = \frac{\alpha}{e^{\alpha b} (e^{\alpha(qs-v)} - 1)} < 0.$$

For $\frac{\partial \mu_{q+1}}{\partial s}$, the denominator is always positive. In the numerator, one term is negative whenever the other is positive, since $e^{\alpha b} > 1$ iff $\alpha > 0$, making the entire expression negative. For $\frac{\partial \mu_{q+1}}{\partial b}$ in the denominator, $e^{\alpha(qs-v)} > 1$ iff $\alpha < 0$. Since α is in the numerator, the fraction is negative for any α .

We note that an increase in either b or s also decreases Q and hence the support of the distribution. But this simply strengthens the dominance of the distribution with lower s or b .

This establishes that the average final bid $\sum_{q=0}^Q qf(q)$ is increasing in b , and decreasing in α or s . Expected revenue, however, is $(b + s) \sum_{q=0}^Q qf(q)$, but in the case of α , the change in revenue is proportional to the change in average final bid.

Acknowledgments

The authors are grateful for the thoughtful comments of Lars Lefgren, Tim Salmon, Chris Cotton, and Andrzej Skrzypacz, as well as participants at the 2009 Stony Brook Game Theory Festival.

References

- Ali, Mukhtar M. (1977) "Probability and Utility Estimates for Racetrack Bettors," *Journal of Political Economy* **85**, 803-15.
- Alcalde, José, and M. Dahm. (2011) "The Complete Information First-Price Auction and Its Intuitive Solution," *International Game Theory Review* **13**, 353-361.
- Amann, Erwin, and W. Leininger. (1996) "Asymmetric All-Pay Auctions with Incomplete Information," *Games and Economic Behavior* **14**, 1-18.
- Augenblick, Ned. (2009) "Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis." Mimeo, Stanford University.
- Avery, Christopher. (1998) "Strategic Jump Bidding in English Auctions," *Review of Economic Studies* **65**, 185-210.
- Barut, Yasar, and D. Kovenock. (1998) "The Symmetric Multiple Prize All-pay Auction with Complete Information," *European Journal of Political Economy* **14**, 627-644.
- Baye, Michael, D. Kovenock, and C. de Vries. (1996) "The All-Pay Auction with Complete Information," *Economic Theory* **8**, 291-305.
- Byers, John, M. Mitzenmacher, and G. Zervas. (2010) "Information Asymmetries in Pay-Per-Bid Auctions," *Electronic Commerce: Proceedings of the 11th ACM Conference on Electronic Commerce*, ACM:New York. 1-12.
- Clark, Derek, and C. Riis. (1998) "Competition over More than One Prize," *American Economic Review* **88**, 276-289.
- 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.
- Gimein, Mark. (2009) "The Crack Cocaine of Auction Sites," *The Big Money*, July 7, 2009. <http://www.thebigmoney.com/articles/money-trail/2009/07/07/crack-cocaine-auction-sites>.
- 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.
- Hartley, Roger, and L. Farrell. (2002) "Can Expected Utility Theory Explain Gambling?" *American Economic Review* **92**, 613-624.

- Hinnosaar, Toomas. (2010) "Penny Auctions are Unpredictable," Mimeo, Northwestern University.
- Jullien, Bruno, and B. Salanié. (2000) "Estimating Preferences under Risk: The Case of Racetrack Bettors," *Journal of Political Economy* **108**, 503-530.
- Kahneman, Daniel, and A. Tversky. (1979) "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* **47**, 263-291.
- Kanto, Antti, G. Rosenqvist, and A. Suvas. (1992) "On Utility Function Estimation of Racetrack Bettors," *Journal of Economic Psychology* **13**, 491-98.
- Kato, Donna. (2009) "Penny Auction Sites Offer Chance for Big Bargains," *San Jose Mercury News*, May 22, 2009.
- Konrad, Kai, and W. Leininger. (2007) "The Generalized Stackelberg Equilibrium of the All-pay Auction with Complete Information," *Review of Economic Design* **11**, 165-174.
- Krishna, Vijay, and J. Morgan. (1997) "An Analysis of the War of Attrition and the All-Pay Auction," *Journal of Economic Theory* **72**, 343-362.
- Kwang, Ng Yew. (1965) "Why Do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure," *Journal of Political Economy* **73**, 530-535.
- Maynard Smith, John. (1974) "The Theory of Games and Evolution in Animal Conflicts," *Journal of Theoretical Biology* **47**, 209-221.
- McDonald, James. (1984) "Some Generalized Functions for the Size Distribution of Income," *Econometrica* **52**, 647-665.
- McDonald, James. (1995) "A Generalization of the Beta Distribution with Applications," *Journal of Econometrics* **66**, 133-152.
- Reklaitis, Victor. (2009) "Entertainment E-Com On Web's Retail Stage," *Investor's Business Daily* April 7, 2009.
- Roth, Alvin, E. and A. Ockenfels. (2002) "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review* **92**, 1093-1103.