

Information Asymmetries in Pay-Per-Bid Auctions

John W. Byers*
Computer Science Dept.
Boston University

Michael Mitzenmacher†
School of Eng. Appl. Sci.
Harvard University

Georgios Zervas*
Computer Science Dept.
Boston University

ABSTRACT

Recently, some mainstream e-commerce web sites have begun using “pay-per-bid” auctions to sell items, from video games to bars of gold. In these auctions, bidders incur a cost for placing each bid in addition to (or sometimes in lieu of) the winner’s final purchase cost. Thus even when a winner’s purchase cost is a small fraction of the item’s intrinsic value, the auctioneer can still profit handsomely from the bid fees. Our work provides novel analyses for these auctions, based on both modeling and datasets derived from auctions at Swoopo.com, the leading pay-per-bid auction site. While previous modeling work predicts profit-free equilibria, we analyze the impact of *information asymmetry* broadly, as well as Swoopo features such as bidpicks and the Swoop It Now option specifically. We find that even small asymmetries across players (cheaper bids, better estimates of other players’ intent, different valuations of items, committed players willing to play “chicken”) can increase the auction duration significantly and thus skew the auctioneer’s profit disproportionately. We discuss our findings in the context of a dataset of thousands of live auctions we observed on Swoopo, which enables us also to examine behavioral factors, such as the power of aggressive bidding. Ultimately, our findings show that even with fully rational players, if players overlook or are unaware any of these factors, the result is outsized profits for pay-per-bid auctioneers.

1. INTRODUCTION

One of the more interesting commercial web sites to appear recently from the standpoint of computational economics is Swoopo. Swoopo runs an auction website, using a nontraditional “pay-per-bid” auction format. Although we provide a more formal description later, the basic framework is easy to describe. As with standard eBay auctions, pay-per-bid auctions for items begin at a reserve price (generally 0), and have an associated countdown clock. When a player places a bid, the current auction price is incremented by a fixed amount, and some additional time is added to the

*E-mail: {byers, zg}@cs.bu.edu. Supported in part by Adverplex, Inc. and by NSF grant CNS-0520166.

†E-mail: michaelm@eecs.harvard.edu. Supported in part by NSF grants CCF-0915922, CNS-0721491, and CCF-0634923, and in part by research grants from Yahoo!, Google, and Cisco Systems.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Submitted to EC 2010, Cambridge, MA, 2010

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

clock. When the clock expires, the last bidder must purchase the item at the final auction price. The pay-per-bid twist is that each time a player increments the price and becomes the current leader of the auction, they must pay a bid fee. On Swoopo, the typical bid fee is 60 cents and the price increment ranges from 1 cent to 24 cents.

While there are other web sites using similar auctions, Swoopo has become the leader in this area, and recently has inspired multiple papers that attempt to analyze the characteristics of the Swoopo auction [2, 5, 7]. These models share the same basic framework, based on assuming players decide whether or not to bid in a risk-neutral fashion, which we explain in detail in Section 2. Some of these papers then go further, and attempt to justify their model by analyzing data from monitoring Swoopo auctions.

One of the most interesting things about the nearly identical analyses undertaken thus far is that the simple versions of the model predict negligible profits for Swoopo, in that the expected revenue essentially matches the value of the item sold. This fails to match the results from datasets studied in these papers, other anecdotal evidence [11, 12], as well as hard evidence we compiled in a dataset comprising over one hundred thousand auction outcomes that we collected, which show Swoopo making dramatic profits (see Figure 1).¹ Some suggestions in previous work have been made to account for this, including the relaxation of the assumption that players are risk-neutral [7], or the addition of a regret cost to model the impact of sunk costs [2].

In this paper, we take the previous analysis as a starting point, but we focus on whether there are intrinsic aspects of the pay-per-bid auction framework that can derive profit from even rational, risk-neutral players who correctly model sunk costs. Specifically, previous work has modeled the game as inherently symmetric, with all players adopting identical randomized strategies. However, there are natural asymmetries that can arise in the Swoopo auction, particularly *asymmetries in information*. A rational player’s strategy revolves around his assessment of the probability of winning the auction outright by bidding, based on the current bid, the number of bidders, the bid fee, and the value of the item. Let us focus on one of these parameters, the number of players n . Although previous models assume that n is known to all players in advance, in practice, there is no way to know exactly how many players are actively participating or monitoring the auction at any time. In Section 3, we show that even small asymmetries in beliefs in the number of active players can lead to dramatic changes in overall auction revenue, and these changes can grow sharply as the estimates vary from the true number of players.

As a related example, previous analyses assume that all

¹We estimate Swoopo’s net profits for an auction by summing up estimated bid fees plus the final purchase price and subtracting the stated retail value for the item.




Figure 1: Empirical estimate of profit margins for 114,628 Swoopo ascending-price auctions.

players both pay the same fee to place a bid in an auction and ascribe an identical value to an item. The latter is generally not the case. Less obviously, not all bidders on Swoopo are paying the same price per bid, for reasons we discuss in Section 5. In this case, players using less expensive bids have both a decided information advantage and a tactical advantage. Pushing this to the extreme, we have the case of shill bidders, who bid on behalf of the auctioneer, and can be modeled as bidders who incur no cost to bid (but also never claim an item). While we do not suggest shill bidders are present in online pay-per-bid auctions, our analysis in Section 6.2 nevertheless shows that they would have a striking impact on profitability.

Finally, our framework allows us to examine other interesting aspects of these auctions that are difficult to model analytically, but which can be studied via empirical observations. One question that we are particularly interested in is whether certain bidder behavior, such as aggressive bidding, is effective, as earlier work speculates [2]. In Section 7.2, we formulate a new definition of bidder aggression, and demonstrate that bidders range widely across the aggression spectrum. While aggressive bidders win more often, analysis of our dataset shows that the most aggressive bidders contribute the lion’s share of profits to Swoopo, and it is a surprisingly passive set of players who fare best.

We believe that modeling and analyzing these information asymmetries are interesting in their own right, although we also argue that they provide a more realistic framework and possible explanation for Swoopo profits than previous work. Indeed, our work reveals the previously hidden complexity of this auction process in the real-world setting.

We emphasize that while we provide data in an attempt to justify these additions to the model, in contrast to previous work, we eschew efforts to fit existing data to our model to parameterize and validate it. Given our understanding, we suggest that models at this stage can provide a high-level understanding, but it may be difficult to disentangle various effects through auction data alone without more detailed insight of user behavior. Moreover, current models appear as yet far from complete. We therefore suggest future alternatives and directions in the conclusion.

Finally, we note that, due to space limitations, we cannot fully describe all of our results in this paper. A much longer and more detailed version is available for download on the arXiv [3]. In particular, in many of our mathematical derivations here, we focus on the simpler case of fixed-price auctions, described in Section 2, for space reasons.

1.1 Related Work

Several recent working papers have studied pay-per-bid auctions [2, 5, 7]. While there are some differences among the papers, they all utilize the same basic framework, which is based on finding an equilibrium behavior for the players of the auction. We describe this framework in Section 2, and use it as a starting point. The key feature of this framework from our standpoint is that it treats the players as behaving symmetrically, with full information. Unsurprisingly, in such a setting the expected profit for Swoopo is theoretically zero.

Our key deviation from past work is to consider asymmetries inherent in such auctions, with a particular focus on information asymmetry. Information asymmetry broadly refers to situations where one party has better information than the others, and has become a key concept in economics, with thousands of papers on the topic. The pioneering work of Akerlof [1], Spence [10], and Stiglitz [9], for which the authors received a Nobel Prize in 2001, established the area. Typical examples of information asymmetry include insider trading, used-car sales, and insurance. Interestingly, the study of information asymmetry in auctions appears significantly less studied. We believe that our analysis of Swoopo auctions provides a simple, natural example of the potential effects of information asymmetry (as well as other asymmetries) in an auction setting, and as such may be valuable beyond the analysis itself.

1.2 Datasets

Where appropriate, we motivate our work or provide evidence for our results via data from Swoopo auctions. We have collected two datasets. One dataset is based on information published directly by Swoopo, which contains limited information about an auction. Information provided includes basic features such as the product description, the retail price, the final auction price, the bid fee, the price increment, and so on. This dataset covers over 121,419 auctions. We refer to this as the OUTCOMES dataset.

Our second dataset is based on traces of live auctions that we have ourselves recorded using our own recording infrastructure. Our traces include the same information from the Swoopo auctions as well as detailed bidding information for each auction, specifically the time and the player associated with each bid. This dataset spans 7,353 auctions and 2,541,332 bids. We refer to this as the TRACE dataset. Our methodology to collect bidding information entailed continuous monitoring of Swoopo auctions; however, in some cases we could not obtain all of the bids. In particular this happened when more than ten players were using BidButlers, automatic bidding agents provided by the Swoopo interface, to bid at a given level, as we collect at most ten bids with each probe of Swoopo. Overall, we captured every bid from 4,328 of the 7,353 auctions; only results from these complete auctions are included in our study. Further details regarding our dataset and collection methods, including how to download the data, can be found in the full version of our paper [3].

2. A SYMMETRIC PAY-PER-BID MODEL

We start with a basic model and analysis of Swoopo auctions from previous work, following the notation and framework of [7], although we note that essentially equivalent analyses have also appeared in other work [2, 5]. This serves to provide background and context for our work.

We consider an auction for an item with an objective value

of v to all players. There are n players throughout the auction. The initial price of the item is 0. In the *ascending-price* version of the auction, when a player places a bid, he pays an up-front cost of b dollars, and the price is incremented by s dollars. The auction has an associated countdown clock; time is added to the clock when a player bids to allow other players the opportunity to bid again. When an auction terminates, the last bidder pays the current price of the item and receives the item. In a variant called a *fixed-price* auction, the winner buys the item for a fixed price p ; bids still cost b dollars but there is no price increment. When $p = 0$, this is called a *100%-off* auction. In our analysis, we simplify players' strategies by removing the impact of timing (but we do study this empirically in Section 7.2). Instead of bidding at a given time, players choose to bid based on the current price, with ties broken at random. A player that chooses not to bid at some price may bid later on.

The basic formulation for analyzing this game is that a player who makes the q th bid is betting b than no future player will bid. Let μ_j be the probability that somebody makes the j th bid (given that $j - 1$ previous bids have been made). Then the expected payoff for the player that makes the q th bid is $(v - sq)(1 - \mu_{q+1})$; a player will only bid if this payoff is non-negative. Note that when $q > Q \equiv \lfloor \frac{v-b}{s} \rfloor$ it is clear that no rational player will bid, as the item price plus bid fee exceeds the value. For convenience in the analysis we will assume that $\frac{v-b}{s}$ is an integer, to avoid technical issues when this does not hold (see [2] for a discussion); this assumption ensures that a player that makes the Q th bid is indifferent to the outcome (the expected payoff is 0). In the fixed-price variant, the payoff is $(v - p)(1 - \mu_{q+1})$; as long as $v > p$, bidding may occur.

The equilibrium behavior is found by determining the probability that a player should bid so that the expected payoff is zero whenever $q \leq Q$, leaving the players indifferent as to the choice of whether to bid or not to bid. Hence the indifference condition is given by

$$b = (v - sq)(1 - \mu_{q+1}),$$

or

$$\mu_{q+1} = 1 - b/(v - sq)$$

in the ascending-price auction, and

$$\mu_q = 1 - b/(v - p)$$

at all steps in the fixed-price auction.

In what follows it is helpful to let β_q be the probability that each player chooses to make the q th bid with given that the $(q - 1)$ st bid has bid made and that the player is not the current leader. Note that by symmetry each player bids with the same probability. Hence, for $q > 1$, for ascending-price auctions we must have

$$1 - \mu_q = (1 - \beta_q)^{n-1}$$

$$\beta_q = 1 - \left(\frac{b}{v - s(q-1)} \right)^{1/(n-1)}.$$

Similarly, we have

$$\beta_q = 1 - \left(\frac{b}{v - p} \right)^{1/(n-1)}$$

for the fixed-price auction.

We point out that the first bid is, effectively, a special case, since at that point there is no leader. To maintain consistency, we want the indifference condition to hold for the first bid; that is, players still bid such that their expected profit is zero. This requires a simple change, since at the first bid there are n players who might bid instead of $n - 1$, giving for the ascending-price auction

$$\beta_1 = 1 - \left(\frac{b}{v} \right)^{\frac{1}{n}},$$

and similarly $\beta_1 = 1 - (b/(v - p))^{\frac{1}{n}}$ for the fixed-price variant.

The expected revenue for the auction can easily be calculated directly using the above quantities. However, we suggest a simple argument (that can be formalized in various ways, such as by defining an appropriate martingale) that demonstrates that Swoopo's expected revenue is v if there is at least one bid, and zero if no player bids. (A similar argument appears in [2].) First note that in auctions where there is at least one bid, an item of value v is transferred to some player at the end of the auction. Also, by the indifference condition, the expected gain to the player that places any bid is zero. (Think of a bid b as counterbalanced by the auctioneer putting an expected value b at risk.) Therefore, by linearity of expectations, the auctioneer recoups a sum of payments equal to v in expectation over the course of the auction, conditioned on there being at least one bid. The probability that no player bids is $(1 - \beta_1)^n$ by definition of β_1 , and thus the expected revenue is $v(1 - (1 - \beta_1)^n) = v - b$.

To be clear, in what follows, we will always consider revenue conditioned on the auction having had at least one bid, since otherwise, the auction is essentially a non-operation for the auctioneer. We call such auctions *successful*.

3. ASYMMETRIC PLAYER ESTIMATES

The analysis of Section 2 assumes that the number of players is fixed and known throughout. This assumption has been questioned in previous work; for example, in [7], they propose a variation where the expected number of players at each time step is known and the distribution is assumed to be Poisson, to model participants entering and leaving the auction over time. The end result is a small variation on the previous analysis. Here we take a different approach and remove the assumption that every player has the same estimate of n , the number of players in the game.

Before diving into the analysis, we provide some motivating data from our datasets. During an auction, Swoopo provides a list of the bidders that have been active over the last 15 minute period. Analysis of our TRACE dataset indicate that this significantly underestimates the total number of participants in the auction, so players who rely on this information to estimate the number of players may be misled. Our analysis shows that players underestimating n can dramatically inflate Swoopo's expected revenue.

Following Swoopo, we define an *active bidder* as someone who has bid in the last fifteen minutes. Using our TRACE dataset, we observed each auction at one minute intervals, and at each time instant we computed the number of active bidders as a percentage of the total number of players who ultimately participated in the auction. We found that auction participation typically builds to a crescendo at the end of the auction; on average, ten minutes from the end of the auction only 20% of all bidders have participated, and

five minutes from the end only 40% of all bidders have participated. Further, due to the nature of the auction, there is no fixed time at which the auction ends, so even bidders making predictions based on past observations are using a certain degree of guesswork.

We now consider the analysis of fixed-price auctions. To initially frame the analysis, we further assume that the true number of players is n , but *all* players perceive the number of players as $n - k$ for some k in the range $[1, n - 2]$. In this case, there is still symmetry among the players, but they choose to bid based on incorrect information. Following the previous analysis, to maintain the indifference condition that the expected revenue for a player that bids at each point should be equal to their bid fee, we have $(v - p)(1 - \nu_q) = b$, where now ν_q is the perceived probability that someone else will place the q th bid. As before, $\nu_q = 1 - \frac{b}{v - p}$. Again we let β_q be the probability that a player chooses to make the q th bid. For $q > 1$ we have $(1 - \nu_q) = (1 - \beta_q)^{n - k - 1}$, or $\beta_q = 1 - (1 - \nu_q)^{\frac{1}{n - k - 1}}$.

Crucially, ν_q is not equal to μ_q , the true probability that someone will make the q th bid. Since $(1 - \mu_q)$ equals the probability that nobody makes the q th bid, we have

$$\begin{aligned} 1 - \mu_q &= (1 - \beta_q)^{n - 1} \\ \mu_q &= 1 - \left((1 - \nu_q)^{\frac{1}{n - k - 1}} \right)^{n - 1} \\ \mu_q &= 1 - \left(\frac{b}{v - p} \right)^{\frac{n - 1}{n - k - 1}}. \end{aligned} \quad (1)$$

Remember that the above holds for $q > 1$, as for the first bid the bidding probabilities are slightly different, as explained in Section 2. Recall we assume a successful auction. Then μ_q is the same for all bids, so we simply call the value μ . The probability that the auction lasts another r bids, after the first, is given by $\mu^r(1 - \mu)$. If we let R be the revenue from a successful auction, we calculate Swoopo's expected revenue as:

$$E[R] = b + p + b \sum_{r=0}^{\infty} r \mu^r (1 - \mu). \quad (2)$$


In the simple case where $p = 0$ the expected revenue is:

$$E[R] = b \left(\frac{v}{b} \right)^{\frac{n - 1}{n - k - 1}}. \quad (3)$$


When $k = 0$, the expected revenue is v . But as k appears in the exponent of the v/b term, even small values of k can have a significant effect on the revenue. This impact as k varies is depicted for a representative auction for \$100 in cash with a bid fee of \$1 and 50 players in Figure 2(a). These will be our default parameters for fixed-price auctions throughout this work.

Conversely, one could consider what happens when players overestimate the population, that is to say $k < 0$. As expected, the revenue for Swoopo then shrinks, incurring an overall loss as demonstrated in the left half of Figure 2(a). Note the considerable asymmetry in the plot, however. Even though Swoopo has far more to gain by pure underestimation of the player population, a mix of overestimation and underestimation in equal measures still yields markedly increased revenues, as depicted in Figure 2(b).

Similar analyses can be made for different settings, such as ascending-price auctions and other mixtures of estimates.



(a) All players underestimate the population by k . Negative values of k stand for overestimates.



(b) Half the players underestimate the population by k and half overestimate it by an equal amount.

Figure 2: Expected revenue for Swoopo in a successful 100% off auction; $n = 50$, $v = 100$ and $b = 1$.

4. MODELING GENERAL ASYMMETRIES

We now consider variations of the auction where there are asymmetries in information. For this we need to extend the symmetric model and make a crucial distinction between the *true* values of the game's parameters – v , b and n – and the way players *perceive* them. (We motivate the misperception of each these parameters in the appropriate sections.) For simplicity, we will assume henceforth that there are two groups of players: A , of size k , and B , of size $n - k$. We can extend our approach to a larger number of groups naturally, but with increased complexity.

Players in group A perceive the value of the item as v^A , the bid fee as b^A , and the number of participants in the game as n^A . Define v^B , b^B and n^B similarly. Initially, we assume that each player is asymmetry-unaware, i.e. each player assumes all players have identical parameters and thus the groups are not aware of each other. We will be also interested in cases where one group is aware of the split and therefore has an advantage over the other group. That setting will utilize the same basic structure; we develop it in later sections.

The parameters determine both the perceived and the true probability of the q th bid being placed. So, for group A , let ν_q^A be the perceived probability that anyone – in either group – will place the q th bid. In other words, ν_q^A is an estimate of μ_q from the perspective of players in group A . Also, define μ_q^A as the true probability that one or more players in group A places the q th bid and similarly define ν_q^B and μ_q^B for group B . If μ_q is the true probability of the q th bid being placed then we have $1 - \mu_q = (1 - \mu_q^A)(1 - \mu_q^B)$.

Players in group A will bid according to their perceived

indifference condition, which for ascending-price auctions is now $(v^A - s(q-1))(1 - \nu_q^A) = b^A$, and similarly for group B . (Similar derivations hold for fixed-price auctions.) Using the fact that $1 - \nu_q^A = (1 - \beta_q^A)^{n-1}$ we can easily derive the individual bidding probability for group A players:

$$\beta_q^A = 1 - \left(\frac{b^A}{v^A - s(q-1)} \right)^{\frac{1}{n^A-1}}. \quad (4)$$

The derivation for group B players is identical. Using the individual bidding probabilities we can compute the probability of a bid being placed by anyone in group A as

$$1 - \mu_q^A = (1 - \beta_q^A)^k \quad (5)$$

$$\mu_q^A = 1 - \left(\frac{b^A}{v^A - s(q-1)} \right)^{\frac{k}{n^A-1}}. \quad (6)$$

Note that generally $\mu_q \neq \nu_q^A \neq \nu_q^B$.

4.1 A Markov Chain Approach

To compute various quantities of interest when we have asymmetric behaviors requires a bit of work, primarily because the probability of a bid at any given time depends in part on what group the current auction leader belongs to. In the models we have described, however, the auction itself is memoryless, in that, given the leader and the current number of bids, the history to reach the current state is unimportant to the future of the auction. Essentially all of our models have this form. Hence, we can place these auctions in the setting of Markov chains in order to efficiently calculate the distribution of auction outcomes.

Specifically, the general case for two groups of players can be captured by an absorbing, time-inhomogeneous Markov chain. (Recall that in a time-inhomogeneous Markov chain, the transition probabilities can depend on the current time as well as the state, but not on the history of the chain.) The chain contains four states: in state A a member of the first group is leading the auction, while in absorbing state W_A the auction has been won by a member of the first group. We define states B and W_B similarly. Note that we overload the notation A and B to refer both to the sets of players and a state of the Markov chain. Finally, note that there is no state corresponding to the initial setting prior to the first bid. Instead we choose the starting state probabilistically from A or B according to the appropriate probabilities for the first bid, recalling that we assume our auction is successful.

We use $p_{AB}(q)$ to denote the transition probability of going from state A to state B after the q th bid, and similarly we can define $p_{BA}(q)$, $p_{AA}(q)$, and so on. For example, $p_{AB}(1)$ is the transition probability from state A to state B when one bid has already been placed. When considering fixed-price auctions, the bidding probabilities, and hence the state transition probabilities, are invariant from bid to bid. In this special case the Markov chain becomes time-homogeneous, and given the distribution on the initial state we can derive analytical expressions for the probability of terminating in state A or B . A good description of this approach can be found in many standard texts; we provide a summary based on [4] in the full version of the paper [3].

For ascending-price auctions, which are time-inhomogeneous, we resort to numerical methods employing simple recurrence relations. This can also be useful to obtain more specific information in the case of fixed-price auctions (or as an alternative approach for calculating various quantities). For




Figure 3: Winners' total cost of bidpacks as a percentage of retail price (Trace dataset).

example, let $P_A(q)$ be the probability of being in state A after q bids; here $P_{W_A}(q)$ represents the probability that a player from A has won the auction at some point up to bid q , so that $P_{W_A}(q) + P_{W_B}(q)$ becomes 1 for an ascending-price auction when q is sufficiently large and converges to 1 for a fixed-price auction as q goes to infinity. Then we have

$$P_A(q+1) = P_A(q)p_{AA}(q) + P_B(q)p_{BA}(q), \quad (7)$$

and other similar recurrences, including

$$P_{W_A}(q+1) = P_A(q)p_{AW_A}(q) + P_{W_A}(q). \quad (8)$$

Given these various equations, it is easy to compute quantities such as the expected revenue. For example, in an ascending-price auction, assuming all players have a bid fee of b , every time A is in the lead, he has paid a bid of b for this, and the price has gone up by s . Letting R be the revenue, we easily find

$$E[R] = (b + s) \left(1 + \sum_{i=1}^Q P_A(i) + \sum_{i=1}^Q P_B(i) \right). \quad (9)$$

Notice that the simple nature of the Markov chain framework allows us to derive all the important quantities, such as the expected revenue for Swoopo, directly from the appropriate transition probabilities. Hence, in the rest of the paper, we focus on finding these probabilities, and leave further details to the reader.

5. ASYMMETRIES IN BID FEES

We now consider asymmetries that arise when players have different bid fees. As one motivation, among other items offered on auction at Swoopo are *bidpacks*, or sets of prepaid bids. Players who win bidpack auctions at a discount to face value and participate in later auctions can effectively enjoy lower bidding fees compared to other participants, generally without the other participants' knowledge.

To provide evidence that bidpacks can lead to varying bid fees, we estimate the total cost of bidpacks for winners of bidpack auctions in our TRACE dataset. Costs include the winners' bid costs and the prices they paid in winning auctions, as well as the bid costs those winners incurred when *losing* other bidpack auctions in our dataset. We then plot the average cost of bidpacks as a percentage of the nominal retail cost in Figure 5. This leads to an overall discount of over 1/2 of the retail cost. While this may still be an underestimate of bidpack costs (as we cannot take into account auctions we have not captured, and our results are biased

towards winners) it suggests that winners of bidpack auctions enjoy a substantial discount in bid fees when applying those bids to other auctions.

Discounted bids are also available through seasonal promotions that Swoopo conducts. Moreover, further variation in bid fees is due to the remarkable fact that Swoopo auctions take place with bidders bidding in different currencies. Further details are given in the full version [3]. Overall, our evidence suggests varying bid fees are realistic in practice, and we turn to quantifying their impact.

We consider the simpler case of fixed-price auctions with price p . We assume that the n bidders are divided in two groups A and B , of size k and $n - k$ respectively. We will assume that $k \geq 2$; the case where $k = 1$ can be handled similarly but the case structure of the analysis is slightly different. Group A incurs a bid fee of b^A while group B incurs a bid fee of b^B with $b^A < b^B$. In context, we may presume that group A is the set of bidders who are bidding at a discount whereas group B is the set of players who are charged regular bid fees. In what follows we also assume A players are aware of the two groups while B players perceive everyone as belonging to the same group as themselves. This creates an information asymmetry. We believe this choice of model is natural; we suspect many (less sophisticated) players may not recognize that others are obtaining cheaper bids. It also provides an example of how our Markov chain approach of Section 4.1 applies to such a setting.

Let μ_q^A be the collective probability that some player in group A makes the q th bid, and similarly define μ_q^B . Then the probability that no player makes the q th bid is:

$$(1 - \mu_q) = (1 - \mu_q^A)(1 - \mu_q^B) \quad (10)$$

where μ_q is defined to be the *true* collective probability that anyone, in either group, bids.

Next, consider the game from the point of view of B players. Remember that, according to them, everyone belongs to a single group incurring the same bid fee. Define ν_q^B as the *perceived* probability that anyone, in either group, makes the q th bid according to the information available to B players. From the indifference condition for B players we have:

$$(v - p)(1 - \nu_q^B) = b^B \quad (11)$$

$$\nu_q^B = 1 - \frac{b^B}{v - p}.$$

We derive the *true* probability β_q^B that a B player bids as:

$$(1 - \nu_q^B) = (1 - \beta_q^B)^{n-1}$$

$$\beta_q^B = 1 - \left(\frac{b^B}{v - p} \right)^{\frac{1}{n-1}}. \quad (12)$$


We can then write the probability of group B bidding as:

$$1 - \mu_q^B = \begin{cases} (1 - \beta_q^B)^{n-k} & \text{if group } A \text{ is leading,} \\ (1 - \beta_q^B)^{n-k-1} & \text{if group } B \text{ is leading,} \end{cases}$$


which after manipulation becomes:

$$\mu_q^B = \begin{cases} 1 - \left(\frac{b^B}{v-p} \right)^{\frac{n-k}{n-1}} & \text{if group } A \text{ is leading,} \\ 1 - \left(\frac{b^B}{v-p} \right)^{\frac{n-k-1}{n-1}} & \text{if group } B \text{ is leading.} \end{cases} \quad (13)$$

Remember that μ_q^B is the *true* collective probability with which group B players bid. Furthermore, notice that players



(a) Revenue as the price of cheap bids varies.



(b) Relative likelihood of a specific player in group A winning the auction when compared with a specific player in group B .

Figure 4: A fixed-priced auction with k players provisioned with cheap bids; $n = 50$, $v = 100$ and $b^B = 1$.

in group A are aware of this probability.

Assuming the leader before the q th bid was from group B and using the indifference condition for group A we can derive an expression for μ_q^A :

$$(v - p)(1 - \mu_q^A)(1 - \mu_q^B) = b^A \quad (14)$$

$$\mu_q^A = 1 - \frac{b^A}{b^B} \left(\frac{b^B}{v - p} \right)^{\frac{k-1}{n-1}}. \quad (15)$$

The derivation for group A leading is similar, leading to:

$$\mu_q^A = \begin{cases} 1 - \frac{b^A}{b^B} \left(\frac{b^B}{v-p} \right)^{\frac{k-1}{n-1}} & \text{if group } A \text{ is leading,} \\ 1 - \frac{b^A}{b^B} \left(\frac{b^B}{v-p} \right)^{\frac{k}{n-1}} & \text{if group } B \text{ is leading.} \end{cases} \quad (16)$$

Next, using Equation 10 we can derive an expression for μ_q , the *true* probability that a q th bid is placed:

$$\mu_q = 1 - \frac{b^A}{v - p}, \quad (17)$$

which holds irrespectively of who is the current leader. It seems counterintuitive that neither the number of B players nor their bid fee play any role in determining the probability μ_q . However, this is similar to the original setting where all players pay the same bid fee, and μ_q was independent of n .

We can also write an expression for β_q^A :

$$\beta_q^A = \begin{cases} 1 - \left(\frac{b^A}{b^B} \right)^{\frac{1}{k-1}} \left(\frac{b^B}{v-p} \right)^{\frac{1}{n-1}} & \text{if group } A \text{ is leading,} \\ 1 - \left(\frac{b^A}{b^B} \right)^{\frac{1}{k}} \left(\frac{b^B}{v-p} \right)^{\frac{1}{n-1}} & \text{if group } B \text{ is leading.} \end{cases}$$

With these bid probabilities in hand, we can apply the

framework developed in Section 4.1. Consider our usual fixed-price auction with $n = 50$, $b = 1$ and $p = 0$. Some bidders have access to a discounted bid fee b^A , while the rest pay the regular rate of \$1 per bid. Figure 4(a) displays Swoopo's expected revenue as the fee b^A charged to group A bidders for bidding varies.

The expected revenue per successful Swoopo auction actually increases, superlinearly, in the gap between bid fees. This is somewhat surprising, given that the amount of revenue from each bid from group A is *decreasing*. However, Group B bidders not only pay full price for their bids, but are also participating in an auction that tends to last substantially longer than they expect. Consequently Swoopo's revenues increase as well. Our analysis hinges on the assumption that group B bidders never realize that they have been dealt a losing hand; recall for fixed-price auctions the underlying bidding behavior is memoryless. During actual auctions, Swoopo does not reveal bid costs, making our model plausible. (Extending our model to a setting where players' beliefs about other players evolve as the auction proceeds, and then adapt their strategies, is future work.)

Also of interest is the advantage gained by a specific player having access to cheap bids. Using the same example as above, Figure 4(b) displays the relative likelihood of a specific A player winning the auction compared to a specific B player as a function of the discounted bid fee. There is a clear synergy here: provisioning of cheaper bids helps the players who receive them and the auctioneer alike.

We point out before continuing that, using the same approach, we can also analyze the setting where players have differing intrinsic valuations for the item being auctioned [3].

6. COLLUSION AND SHILL BIDDERS

Our previous analysis allows us to consider other standard situations with information asymmetry due to hidden information. Here we examine the setting where a subset of players collude to form a bidding coalition, and a setting with *shill bidders*, or bidders in the employ of the auctioneer.

6.1 Collusion

A natural approach for collusion is for members of a coalition to agree to not bid against each other, so that if one of them is currently leading the auction, the others bid with zero probability. We wish to quantify the advantage gained by this form of collusion in terms of the size of the coalition.

For our analysis, we assume that there is a group A of k players in a coalition, and a group B of $n - k$ other players not in the coalition. To these $n - k$ players, there appear to be n identical players in the auction. Again, the coalition players bid as usual, provided a coalition member is not the leader. Proceeds and expenses are shared equally between coalition members.

Non-coalition members bid according to their *perceived* indifference condition:

$$\nu_q^B = 1 - \frac{b}{v-p}. \quad (18)$$

This yields

$$\beta_q^B = 1 - \left(\frac{b}{v-p}\right)^{\frac{1}{n-1}}. \quad (19)$$

From this we can derive the *true* probability of a bid by

group B :

$$\mu_q^B = \begin{cases} 1 - \left(\frac{b}{v-p}\right)^{\frac{n-k}{n-1}} & \text{if group } A \text{ is leading,} \\ 1 - \left(\frac{b}{v-p}\right)^{\frac{n-k-1}{n-1}} & \text{if group } B \text{ is leading.} \end{cases} \quad (20)$$

We observe that players in group B , just as a consequence of overestimating the total population, bid less frequently than they should. This fact alone is enough for the coalition of players in group A to gain an edge in winning the auction.

Next, we look at the indifference condition for a player in group A when someone from group B is leading the auction:

$$b = (v-p)(1 - \mu_q^A)(1 - \mu_q^B) \quad (21)$$

$$\mu_q^A = 1 - \left(\frac{b}{v-p}\right)^{\frac{k}{n-1}}. \quad (22)$$

Recall as we stated earlier when group B is leading the auction group A players act independently. Hence

$$\beta_q^A = \begin{cases} 0 & \text{if group } A \text{ is leading,} \\ 1 - \left(\frac{b}{v-p}\right)^{\frac{1}{n-1}} & \text{if group } B \text{ is leading.} \end{cases} \quad (23)$$

Finally, using the fact that $1 - \mu_q = (1 - \mu_q^A)(1 - \mu_q^B)$, we can derive the following expression for the probability of a bid being placed by either group:

$$\mu_q = \begin{cases} 1 - \left(\frac{b}{v-p}\right)^{\frac{n-k}{n-1}} & \text{if group } A \text{ is leading,} \\ 1 - \frac{b}{v-p} & \text{if group } B \text{ is leading.} \end{cases} \quad (24)$$

The increased chances of group A winning the auction are apparent, as the auction is more likely to end when A leads.


Equations 19 and 23 are nearly sufficient to determine the probabilities for our Markov chain analysis. The only remaining issue regards our choice of tie-breaking rule. Notice that a highly optimized coalition could act as a single player controlling many identities, only selecting a single one to use at each opportunity to bid (albeit with higher probability μ_q^A instead of β_q^A). In this case, the coalition would be less likely to win in case of ties. We refer to this as a *single bidder* coalition, and the original, independent case as a *many bidder* coalition.

One would expect two consequences of collusion. First, a coalition of k bidders should have more than k times the probability of a non-colluding bidder to win. Second, the overestimation of the actual player population should negatively impact Swoopo's revenues. We confirm both of these consequences empirically.


Figure 5(a) displays the revenue Swoopo can expect in the presence of a coalition of size k for both tie-breaking rules. As can be seen, revenue declines significantly when large coalitions are present. Figure 5(b) displays the relative likelihood of the coalition winning the auction compared to any *particular* outsider. Even small groups of colluding players can gain a very large advantage in winning the auction, superlinear in the size of the coalition, offering a significant incentive to collude.

6.2 Shill Bidding

A further consideration is the effect of shill bidders, or bidders under the employ of the auction site who attempt to drive up revenue by bidding in order to prevent auctions from terminating early. This is not a theoretical problem;



(a) Revenue for Swoopo when collusion occurs.



(b) Relative likelihood of the colluding group winning the auction vs. any specific outsider.

Figure 5: Fixed-price auctions with a coalition of size k ; $n = 50$, $v = 100$ and $b = 1$.

pay-per-bid auction sites other than Swoopo have been accused of using shill bidding [6]. In the working paper [7], shill bidders were considered, but it was assumed that they would behave equivalently to other players in the auction. This assumption was necessary to maintain the symmetry of the analysis, and was justified by the argument that if shill bidders behave exactly as other players, they would be more difficult to detect. We argue that sites employing shill bidding may be willing to shoulder the increased detection risk associated with increased shill bidding as long as it is accompanied by increased profit.

There are several possible ways of introducing shill bidders. Here we focus on the following natural one: we define a (ρ, L) -shill as one that enters the auction with probability ρ and bids with probability one at each opportunity when they are not the leader until L bids have been made (in total, by all players), at which stage he drops out of the auction. Such an approach provides useful tradeoffs; increasing ρ or L increases the probability of detection, but offers the potential for increased profit.

To analyze shill bidding we employ our usual framework. We have a standard auction with n players with probability $1 - \rho$. With probability ρ the shill enters and the auction has $n + 1$ players. In this case, we separate the bidders into two groups: group A consists of the lone shill, and group B consists of the n legitimate players who are not informed of the shill’s presence. As before we can determine the transition probabilities and use our Markov chain analysis. Recall that shill bidders produce no revenue for the auctioneer, so the expected revenue is determined by the expected number of times a legitimate player is the leader. For convenience we adopt our usual tie-breaking rule, so the leader is picked




Figure 6: Expected profit for Swoopo with a (ρ, L) -shill; $n = 50$, $v = 100$, $b = 1$ and $s = 0.25$.

uniformly at random from the players who bid at each step.

Rather than plot the per-auction revenue with shill bidders, we instead plot the per-auction *profit*. We do this for two reasons. First, since a symmetric, full-information auction results in zero expected profit for the auctioneer in our model, all profit in our plots can be attributed to the presence of the shill. Second, in this setting, there is some chance that the shill will win the auction, in which case the auctioneer’s revenue is all profit, a fact not well captured by a revenue plot. Figure 6 displays the expected profit for Swoopo in the presence of a (ρ, L) -shill for an ascending-price auction. Shill participation in ascending-price auctions has diminishing returns with L , which is to be expected; even though the shill is forcibly extending the expected length of the auction, as the price of the item goes up, legitimate players become less willing to participate.

7. CHICKEN AND AGGRESSION

In this section we address a recently added feature to Swoopo’s interface, Swoop It Now, that appears to have not been analyzed previously. This feature has the potential to significantly change the dynamics of Swoopo auctions. Our suggestion is that this feature may lead to a subclass of players whose strategy makes some Swoopo auctions resemble the game of chicken [8], in contrast to the Markovian games we have modeled in previous sections. In games of chicken, it is generally understood that it can be useful for players to signal their intentions, explicitly or implicitly, to other players, in order to cause them to give up and allow the signaling player to win. A natural signaling approach in the timed auction context is to bid both frequently and quickly after another player bids. This bidding strategy has been noted previously, in the work of [2], where the author dubs this “bidding aggressively” and finds that aggressive bids have higher expected profit. Here we undertake an independent study, making several new contributions. Besides presenting how this behavior can be viewed as a signaling mechanism for a game of chicken embedded in Swoopo, we provide a novel and natural definition of aggression for pay-per-bid auctions. Then, using our trace data, we analyze auctions for signs of aggressiveness, and estimate how aggressiveness correlates with winning auctions and profitability for players. A surprising finding is that both too little and too much aggression appear to be losing strategies.

7.1 Swoop It Now and Chicken

Swoopo recently added the Swoop It Now option to auctions on its site, which gives each player the ability to pur-




Figure 7: Profit for Swoopo and a committed player as α varies; $n = 50$, $v = 100$, $b = 1$ and $s = 0.25$.

chase the item at a given price even if one loses the auction. (Deployment on the US site appears to have occurred around July 2009, before we began taking traces of auctions.) That is, in many auctions, Swoopo provides a nominal retail value for the auction item, call it r . At the end of the auction, a player who has incurred a total bid cost of δ can purchase the item at a price of $r - \delta$, effectively transferring unrecoverable sunk costs to a partial payment for the auction item. The retail value r given by Swoopo is generally significantly higher than the lowest online retail price [2, 7].

Unfortunately we do not currently know how often Swoop It Now is used; to our knowledge such information is neither given by Swoopo, nor derivable from any data Swoopo makes available. While the high nominal retail value is unattractive, after a large losing investment in an auction, this option may become attractive to certain players.

Let us consider the behavior this additional feature introduces and its consequences in two settings: the case where only one *committed* player is willing to exercise this option, and the case where multiple players potentially are. Our assumption here is that $r = \alpha v$, where v is a common value of the item for all players and $\alpha > 1$. Our key finding is that when multiple players consider taking advantage of this option as a backstop, the game becomes a variant of chicken.

Let us first suppose that a single player has the opportunity to buy at the price r , including the amount they spend on bids. This player may attempt to win the auction early by bidding at every possible step, believing that the expected gain dominates the maximum possible loss of $(\alpha - 1)v$. This player will therefore bid until either winning the auction or spending enough so that it is cheaper to buy the item using Swoop It Now than to win it at the current auction price; the other $n - 1$ players will play as usual. The player is essentially equivalent to a shill bidder, except they bid until spending a certain amount, rather than until a certain number of bids have been made, and they actually purchase the item if they win. The approach of Section 6.2 can be applied with minor variations. The main outcome, naturally, is that such a player increases Swoopo’s profit by prolonging the auction, assuming their presence does not change other players’ strategies. Figure 7 displays the profit earned by the single committed participant as well as Swoopo for an ascending-price auction. When a player uses the Swoop It Now feature, if they have bid δ so far, they will have to pay an additional side payment of $r - \delta$ to complete the purchase. Also, we decrement Swoopo’s profit by an additional v to account for the transfer of a second item to the auction winner. Profit for the committed player decreases in α

while the reverse holds for Swoopo. This model ignores the possibility that a player might signal their intention through aggressive bidding so that other players drop out of the auction, resulting in less profit for Swoopo.

In the case of two (or more) players who are interested in using the Swoop It Now feature, the resulting game instead resembles the game-theoretic version of chicken. For convenience we consider a fixed-price auction with a price of 0. Suppose that two players plan to continue to bid until either obtaining the item or spending r in bids and then using the Swoop It Now feature. If both exhaust their bids, they will both lose $(\alpha - 1)v$ in value. But if, instead, one of them backs off, allowing the other player to win, that player will lose only what they have bid so far, and the other player will purchase the item at a discount.

Obviously, we have simplified things considerably in this description; for example, there may be more than two players willing to play chicken in this setting. This is clearly a subject in need of further study. However, the Swoop It Now feature, by keeping individual losses bounded, does appear to embed the potential for games of chicken to erupt within Swoopo auctions. As aggressive bidding is a natural way to signal intent in this setting, (and may be a sound tactic in its own right), we turn to a study of aggression, making use of our TRACE dataset.

7.2 Aggression


In earlier work [2], Augenblick has suggested that aggressive bidding, including bidding immediately after another player has bid and bidding frequently in the same auction, leads to higher expected value for a player. His analysis is based on individual bids rather than bidders; that is, he considers for each bid how the time since the previous bid and the number of bids by the bidder for that item correlate to the expected profit, using regression techniques.

We adopt a different approach, by looking instead at how aggressive bidding affects the profitability of a *player* within an auction, and by providing a single aggression metric to measure the aggressiveness of a strategy. As a bidder may vary his strategy across auctions, we define aggressiveness in the context of a given auction. We believe that considering the effects of aggressiveness at the level of player profitability offers important insights as it views the merits of an aggressive strategy holistically.


We define an aggressive strategy as one which consists of placing many bids in rapid succession to preceding bids. Specifically, let the *response time* for a bid be the number of seconds since the preceding bid. Aggression should be inversely proportional to response time and proportional to the number of bids a bidder places within an auction. Hence we define the aggression of a bidder in a given auction as:

$$\text{Aggression} = \frac{\text{Number of bids}}{\text{Average response time (seconds/bid)}}. \quad (25)$$

To investigate whether aggressive bidding is a successful strategy we look at the traces of 3,026 complete (no missing bids) “NailBiter” auctions (Swoopo auctions which do not permit the use of automated bids by a “BidButler”) in our TRACE dataset. Figure 8(a) displays the empirical CCDF of aggression for bidders who ended up 1) “in the black”, by winning the auction profitably; 2) winning, but perhaps not profitably; and 3) “in the red”, by incurring bid fees in excess of the value of any item won. Note that the classes



(a) Empirical CCDF of aggression.



(b) Cumulative profit vs. aggression rank.

Figure 8: Aggression and profitability.

are not exclusive.

Our first observation is that aggression follows a highly skewed distribution: the majority of players display little aggression, while a small number of players are highly aggressive. Also, not surprisingly, those players winning the auction were bidding much more aggressively than others. More interestingly, we see that successful players, i.e., those who not only won the auction, but did so profitably, are more aggressive than average, but less aggressive than those who win auctions. Arguably, aggression is successful in moderation.

Figure 8(b) provides more insight into this latter point. For all bidder-auction pairs in our dataset, we compute the aggression and profitability of each outcome, and rank these outcomes by aggression (least aggressive first). We then plot the *cumulative* profit for all outcomes through a given rank with dark shading. For reference, we also plot the aggression of bidders at a given rank using light shading and the scale depicted on the right-hand side of the plot. We see that successful strategies are mostly concentrated at aggression ranks lower than average. More interestingly, a fact not evident in Figure 8(a), the highly aggressive players are responsible for most of Swoopo’s profits.

We now revisit the question of whether games of chicken are also taking place within Swoopo. To do that we turn to the TRACE dataset and look at the 3,026 “NailBiter” auctions for evidence of *duels*: auctions culminating in long sequences of back-and-forth bidding between two opponents. We find that 9% of all auctions culminated in a duel lasting at least 10 bids, 5% lasted at least 20 bids, and 1% lasted at least 50 bids. The longest duel we observed was 201 bids long and somewhat humorously took place between users `Cikcik` and `Thedduell`. We believe this provides further evidence that at least some auctions are becoming essentially

games of chicken, and reiterate our supposition that aggressive bidding is used as a signaling method in such settings.

8. CONCLUSIONS AND FUTURE WORK

Swoopo provides a fascinating case study in how new, non-trivial auction mechanisms perform in real-world situations. Here we have focused on the key issue of asymmetry, and in particular, how various manifestations of information asymmetry may be responsible in large part for the significant profits Swoopo appears to enjoy today. At the same time, we have also shown that profitability of these auctions is potentially fragile, especially in cases where signaling by committed players willing to play a game of chicken or collusion between players can end the auction early.

There are clearly many interesting directions to follow from here. One area we have started to examine is asymmetric models of pay-per-bid auctions with full information. Players could have differing bid fees or valuations of the item, but with these fees and valuations known in advance to all players. Interestingly, such models have not yet been considered in any depth in previous work.

Other broad topics for future work include more extensive study of user behavior on Swoopo, the impact of timing (that we and others have abstracted away), models where users can dynamically change their belief and strategies, and the impact of automatic bidding agents such as BidButlers. Finally, there remains the thorny problem of attempting to quantify directly the impact of specific auction characteristics on real-world profitability.

9. REFERENCES

- [1] G. A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- [2] N. Augenblick. Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis. Unpublished manuscript. Available at www.stanford.edu/~ned789, 2009.
- [3] J. W. Byers, M. Mitzenmacher, and G. Zervas. Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank. Technical report, arXiv.org, January 2010. <http://arxiv.org/abs/1001.0592>.
- [4] C. M. Grinstead and J. L. Snell. *Introduction to Probability*. American Mathematical Society, 1997.
- [5] T. Hinnosaar. Penny Auctions. Unpublished manuscript at <http://toomas.hinnosaar.net/>, 2009.
- [6] <http://www.pennyauctionwatch.com/tag/shills/>.
- [7] B. C. Platt, J. Price, and H. Tappen. Pay-to-Bid Auctions. Unpublished manuscript. Available at <http://econ.byu.edu/Faculty/Platt>, 2009.
- [8] A. Rapoport and A. Chamah. The game of chicken. *American Behavioral Scientist*, 10, 1966.
- [9] M. Rothschild and J. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976.
- [10] M. Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973.
- [11] B. Stone. Sites ask users to spend to save. *New York Times*, August 17, 2009.
- [12] R. H. Thaler. Paying a price for the thrill of the hunt. *New York Times*, November 15, 2009.