“Entertainment Shopping”

An Analysis of Profit and Strategy in a New Auction Format

A thesis presented

by

Jeff Nanney

To

Applied Mathematics

in partial fulfillment of the honors requirements

for the degree of

Bachelor of Arts

Harvard College

Cambridge, Massachusetts

April 1, 2010

Abstract

Swoopo, an "entertainment shopping" website, implements *pay-to-bid* auctions to what appears to be great effect. Our estimates suggest they profit around $200 per auction, or 70% of the auctioned item's value, averaged across a dataset of 120,000 auctions. Swoopo forms the best-known case study in non-traditional auction design, offering a rich dataset and self-suggesting research questions that are largely untouched in the literature. Our approach is two-fold, combining a thorough theoretical treatment with simple, original and striking empirical results. We propose a general, asymmetric, full-information model and describe its equilibria, solving explicitly in those cases that prove relevant to our empirical analysis. We further propose and solve two general models of informational asymmetry, and borrowing motivation from [5], we propose models to explain the levels of revenue and their distribution over Swoopo auctions. We also offer the first treatment of lottery properties as potential explanations for observed auction behaviors; specifically, we use our models of informational asymmetry to model scenarios with risk-loving agents. Using historical auction data and bid-level records, we conclude that Swoopo's profits are outsized and highly consistent. Further, we compare the equilibrium predictions of an auction's termination probability to the data, concluding that the expected return to bidding decreases substantially as auctions progress. We provide robust analysis partially attributing this effect to our theoretical predictions concerning the arrival of agents to the auction over time. We present the first analysis of Swoopo's timer, establishing a straightforward yet large effect relating the time remaining to the profitability of a bid. We conclude that many characteristics of the Swoopo revenue distribution can be explained through our models, most suggestively in the sensitivity of revenue to the agent-arrival process. Formalizing and solving these models of agent-arrival represent a well-motivated avenue for further research.

Contents

1 Introduction

Consider an exercise in which you, as an auctioneer, have the opportunity to auction off a single dollar. By demanding that the highest bidder pay the amount of his bid, you are very unlikely to benefit from this exercise, as it would require a bidder to commit to losing money with no potential upside. Instead, imagine demanding the *two* highest bidders pay you the amount they bid. And, then, consider the auction in a sequential bidding format, so that any player may become the leader at any point by offering an amount greater than the current high bid. Now it is quite easy to see, whether by intuition or by implementation, that you have a potential upside. Indeed, anecdotally, many clever individuals have parlayed this simple format into a modest profit, and some parlor-game entertainment.

In recent years, Swoopo has taken on the role of dollar-auctioneer to similar effect, but on a global scale. Swoopo runs what many term *pay-to-bid auctions* (or, colloquially, *penny auctions*), so named because each bid placement incurs a fixed cost to the bidder. Given this structure in which non-winners can lose a substantial amount of money, many have drawn parallels [13] to Martin Shubik's classic dollar auction, itself an illustrative example of the "War of Attrition." In particular, Swoopo charges relatively small fees per bid, and increments the price of the item by a substantially smaller amount. This discrepancy appears to be crucial to explaining penny auction revenues. This structure permits many users to lose, while still allowing one agent to win, and often win big.

Swoopo primarily auctions consumer electronics (TVs, cameras, computers, game consoles), video games, and watches, though it also has auctioned cash, gold, platinum, and in the height of self-reference, packs of bids to be used on Swoopo. To the unsophisticated observer, Swoopo's auctions appear to be a bargain-hunter's paradise. As its ads communicate, Swoopo attracts clients because individuals focus on the level of savings (frequently on the order of 70 to 90%) on an item's value. To the keen observer, Swoopo is just as alluring, if not more so. It auctions suggest a highly and consistently profitable business that is both completely imitable and wildly popular.

Hence, we are motivated by two big questions: how do we *expect* bidders to behave in this format, and how do bidders *actually* behave in this format? The former question lies at the heart of the theoretical treatment of Swoopo, while the latter at the heart of its empirics. The high profile and profitability almost certainly have inspired many curious observers to discover how much and in what manner Swoopo is cashing in on these auctions. Likewise, Swoopo's online records of ended auctions has helped fuel these academic curiosities.

What sets Swoopo apart from most auction design perspectives is that it is an *all-pay auction*, insofar as every participant must pay each bid's fee. Unusual for this class of auctions, however, is the fact that Swoopo payoffs are unconstrained. A bidder who does not win the auction can lose more in bid fees than the total cost incurred by the winner. Thus, pay-to-bid auctions have almost entirely avoided academic treatment, which we do not consider surprising in that Swoopo is the first known example of this auction format on any meaningful scale. This changed in the latter half of 2009, when a number of unpublished manuscripts appeared.

With so much relatively low-hanging fruit, much of this early analysis of Swoopo bears substantial similarity to one another, despite the likelihood that this early analysis was developed independently. The same observation applies to our own work, though we try to interface better with the existing literature than previous efforts. In a few instances, we describe results that already exist in the literature, and if relevant, we offer our independent formulation. In general, however, we focus on describing the more novel aspects of our research, while putting them in the context of ideas in the recent literature.

1.1 Contributions

Our work is dually motivated, and so we conceive of our results as being a communication between theoretical work and empirical analysis. We have chosen a presentation that focuses first and foremost on discussing theoretical results, but we consider both aspects of the paper to be parts of the same whole. Because empirical questions form the most primitive motive, we attempt to maintain a constant dialogue in our theoretical work, so as to motivate results. We also take liberties in describing stories and intuitions that guide modeling choices, but emphasize that these purely illustrative, and do not form justifications by themselves.

The first task we undertake is describing a baseline model, which constitutes the springboard for both theoretical variations and empirical tests. We acknowledge that an analogous model underpins the work in [9], [10], [3], and [5]. We also stress the importance of what we term *natural equilibria*, which have the structure of the auction surviving each non-trivial time-step (or bid-level) with nonzero probability.

Next, in what is perhaps our broadest result, we propose a completely general asymmetric model, in which each agent faces a different vector of auction-related parameters. We offer substantial motivation, but the key is that in practice agents often face different conditions. We give a complete description of equilibria for two large cases of the problem, and we also cover the conditions in which *trivial equilibria* develop. Finally, we present and demonstrate a method of solving for these equilibria analytically, and put conditions on when these equilibria exist.

We further propose and solve a general asymmetric model in which two *types* of agents comprise the type space, but neither type is aware of the other. Although we note [5] does treat a specific case of this model, we discuss ours as a key generalization. We are also glad to see that our model was sufficiently well-motivated to have a subcase thoroughly examined. Although the emphasis on asymmetry is distinctly a property of [5] and not of our paper, we do note that this general model of ours preceded any asymmetric or imperfect information result in the literature. We also observe that in conceiving this model, we were motivated to consider the aforementioned general asymmetric model.

Motivated by an example of asymmetric information [5], we formulate an altered asymmetric model where exactly one group is aware of the presence of the other. We solve for and describe the equilibria, and discuss the impact and meaning of the informational asymmetry, while providing intuition for the structure of the mixed-strategy equilibrium. This generalization is especially useful, as we provide conditions for when it makes sense to swap the direction of the asymmetry, and when the model breaks down entirely.

By exposing conditions under which the informational asymmetry model is degenerate, we encapsulate our favorite result from [5]. Specifically, we demonstrate that their misestimation model, in which agents incorrectly estimate the population, satisfies the structure of one of our corner cases above. We are especially pleased with this result because it provides a broader theoretical framework in which the somewhat arbitrary "misestimation" makes sense. We proceed to propose a few models that match our intuition about population dynamics in an auction. These serve to motivate the stories of population dynamics that we posit.

Our final theoretical contribution is the first real treatment of the Swoopo problem as something other than an auction. We explore the lottery-like properties of the auctions, and attempt to connect Swoopo observations to lottery literature. Then, we model these lottery intuitions by incorporating them into the models in the most parsimonious way possible. In particular, we provide a discussion of how to model the problem with a subpopulation of risk-seeking bidders, and the effect this has on equilibrium. We use our generalized results for asymmetries to illustrate possible interpretations of risk-lovingness.

We begin our empirical section by thoroughly detailing what has become increasingly well-known: Swoopo's auction format appears to earn substantial, consistent profits, robust over time and auction type. Though unable to resolve it, we are the first to identify a significant trend in the revenue data for Swoopo from April 2009 to October 2009 (approximately). We propose an explanation, and evaluate it in the context of changes to Swoopo. We are particularly surprised that this trend is not in the literature.

We also make a strong argument that skew-lovingness is playing some role in the revenue story, a latent belief that is not addressed in the literature. We consider this significant in the context of questions about user preferences, as well as the appropriateness of a lottery paradigm for Swoopo. We see completely consistent trends of the lowest price-incremented auctions having the highest overall revenues. We also control for item value, and demonstrate that revenue as a percentage has the same pattern of being (essentially) monotonically decreasing in the bid-increment.

The most significant empirical contribution is a robust demonstration that the profitability of a bid rapidly decreases over time. We also consider this finding extremely important in that it meaningfully challenges the lottery intuition: time-steps where the skew in payoff is lowest (at the right tail) appear to be the least profitable.

This decreasing return to bids through the auction is apparent in a variety of interpretations in a variety of subsets of the data, and we find it particularly noteworthy in that [10] predicted the opposite: that due to a risk-lovingness argument, the expected return on a bid would approach 1 as the length of an auction increased. We credit [3] for independently discovering this key trait of Swoopo. [3] postulates the effect is caused by the interplay between naivety and sunk cost, a model we briefly discuss just to measure plausibility. Though we find the intuition palatable, and believe sunk costs probably do play a role, we find no evidence of this claim in the bid-level data.

We do, however, demonstrate a strong, positive effect between the number of unique bidders in the early stage of an auction and the overall profitability, even when we control for the auction lasting at least as long as some arbitrary level. This result uses bid-level data, and we find this effect to be robust to the parameter values. The connection is indirect, but this analysis *strongly* suggests an explanation for the aforementioned decreasing returns to bids: auctions which have a higher ratio of unique bidders almost certainly have a larger population. Related closely to our theoretical work on agent entry, and the result in [5] showing the impact of misestimates on revenue, we believe this unique-bidder effect partially captures an agent entry process, during which agents do not account for the increased population.

We use bid-level data in a novel way to pick out evidence of a particularized bidding population, whom we call *early bidders*. This idea of profiling bidders seems very open, and we believe our general method of looking for concentrations of users in bid-level data provides a realistic method of discovering the categories bidders may fall into.

We also reject suggestions in the literature that the timer aspect of Swoopo is uninteresting.

Using a meticulous formatting of our dataset, but otherwise entirely unsophisticated techniques, we show that profitability varies greatly with the amount of time left on the timer at the time of placement. This confirms the intuition that bidding at the last minute is profitable, and bidding 2 to 3 seconds after the timer has just been reset appears to be optimal.

1.2 Relevant Literature

Although Swoopo has been in existence in some form since 2004, its recent rise to prominence has induced academic interest only recently. When this research was originally undertaken, there was no academic literature on penny auctions. Unsurprisingly, as Swoopo has become an increasingly prominent auction website, a number of attempts to explain the behavior of auction participants have emerged. A few of these efforts have resulted in unpublished manuscripts that engage the same spirit as this paper, namely to describe the equilibrium properties of a penny auction and compare them to the empirical results evident in Swoopo data.

The penny auction format shares structural similarity with Shubik's dollar auction [11]. In the original formulation, Shubik's all-pay auction format demonstrated the potential for an unusual auction structure to produce, in practice, outsized profits for the auctioneer. Practically speaking, the significance of Shubik's all-pay paradigm has been seen to lie in its application to lobbying games, in which rent-seeking firms must pay upfront costs with the hopes of securing a valuable government contract or policy. In the context of penny auctions, the relevance of Shubik's observations seems obvious. Particularly, Shubik established a precedent for the behavior seen in the apparently unprofitable auction participants of Swoopo and other penny auctions.

Hinnosaar's game-theoretic analysis of Swoopo was among the earliest working papers [9] on Swoopo. The model introduces n agents, each of whom has symmetric and perfect information about the value of the auctioned item as well as the beliefs of his competitors. Of course, this model produces the dissatisfying result that the auctioneer receives zero or small negative profits in equilibrium. Hinnosaar acknowledges the inadequacy of this conclusion in light of the evidence of Swoopo profits, and offers a few cursory glances at potential caveats to the model that might yield auctioneer profits.

Platt, Price, and Tappen provide another early economic interpretation [10] of the Swoopo problem. They also treat the problem as a symmetric game, evaluating the mixed-strategy equilibrium in which every player bids at any given level with some probability. This model also predicts zero or small negative profits. However, the authors claim this prediction matches the data adequately for the majority of auction types (i.e., auctions on items other than video games). This empirical claim seems inappropriate for a number of reasons discussed herein. The authors do, however, offer a valuable modeling alternative to explain perceived profits in video game auctions. Specifically, they propose risk-loving auction participants, a possibility explored in more detail in this paper. Given the heuristic similarity between the payoffs advertised in Swoopo and those of lotteries, the existence of risk-loving agents in Swoopo seems entirely plausible.

The Augenblick job market paper [3] is the most thorough version of analysis similar to those described above. Augenblick describes an analogous model that predicts no profits, a result qualified by introducing the notion of sunk costs on the part of bidders. The model ascribes to bidders a regret cost of participating in the penny auction, and the key feature is that the bidder discounts (naively) his future regret cost. These understated sunk costs permit bidders to bid above the original equilibrium level and gives nonzero profits to the auctioneer. Augenblick also offers a more detailed empirical analysis than the previous papers, particularly by examining the firm's strategy in supplying auctions to the marketplace.

The most creative analysis to date appears in Byers, Mitzenmacher, and Zervas's working paper [5]. Their work builds off the standard model common to the above papers, admitting the potential for imperfect information on the part of the bidders. Similarly, their work relaxes the symmetry assumption, and describes a few analytical cases of specific instances of asymmetry. Their paper also surveys the cases of collusive behavior amongst participants, shill bidders employed by the auctioneer, and the role of aggressive bidding. Many of these variations on the general Swoopo model are creative and somewhat motivated, but the paper stops short of connecting much of the modeling possibilities to the data provided by Swoopo's auctions. One of the most interesting variations discussed is the potential for bidders to underestimate the number of competitors in the auction. This contingency can have a dramatic influence on the length of the auction and revenue.

1.3 Swoopo Structure

Swoopo has been listing auctions since 2004, originally as the company TeleBid, where individuals participated in auctions by placing bids over the phone. However, Swoopo has primarily been an online auction mechanism. In order to place a bid in a Swoopo auction, a bidder must pre-purchase a BidPack, which contains a fixed number of bids that the user may then place in individual auctions. Through this purchase, each bid has an implicit cost to the bidder. Upon placing the bid, however, the price of the item being auctioned only increments by a fixed amount, specific to each auction.

By placing a bid, the bidder must pay the price of the item if and only if the auction subsequently ends with his bid as the final bid. Crucial to the results of most Swoopo auctions is the fact that this fixed incremental amount is substantially smaller than the actual cost to the bidder of placing his bid. Also important, Swoopo auctions feature a timer, visible to all bidders, which resets after every bid. The stated purpose claims that this clock seeks to prevent sniping, so that every bidder has an adequate chance to respond to the most recent bid. In practice, it seems clear that the updates to the timer serve the express purpose of extending the length of the auction to increase revenue.

Up through July 2009, the cost of each bid on Swoopo was 75 cents, after which point the cost transitioned to 60 cents. Swoopo has always offered a variety of auction structures, primarily by varying the amount by which the final price is incremented. Auctions have featured increments of 24, 15, 12, 6, 2, 1 and 0 cents. The 0-cent increment took the form of "free auctions" or "fixed price auctions," which allow the winning bidder to pay nothing, or a predetermined fixed amount, regardless of the final incremented price of the item. Over time, Swoopo appears to have experimented substantially with each of these auction categories. In 2009, Swoopo stopped offering fixed price and free auctions. Moving into 2010, the variety of increments decreased sharply, and now Swoopo offers exclusively penny auctions (i.e., auctions in which the increment is exactly one cent) on their auctions as of this writing.

The last notable feature of Swoopo auctions is the BidButler, an extremely simplistic bidding agent at the disposal of every Swoopo bidder. The BidButler essentially serves the purpose of submitting automated bids within a user-specified price range when the last 10 seconds of the auction have been reached. Perhaps due to complaints related to this feature, or perhaps for other reasons, Swoopo offers "NailBiter" auctions, which do not allow users to use the BidButler, so that all bids are placed manually. In similar spirit, Swoopo offers "Beginner" auctions, in which bidders must have not yet won an item in any Swoopo auction.

Also in July 2009, Swoopo introduced a somewhat radical feature to the auction format, termed "Swoop it Now", which permits users who lose an auction to use the cost of the bids placed in that particular auction towards buying the item at Swoopo's stated retail price. This feature creates pathological cases for our analysis, primarily because Swoopo does not provide any data on which users, if any, exercise their "Swoop it Now" option on a given auction. This clearly hinders any method of precisely measuring Swoopo revenue, as well as any method of precisely measuring the cost to any particular bidder. We recognize that this can have a confounding effect on our analysis, but we observe that the micro-data does not suggest any sort of paradigm shift around the introduction of "Swoop it Now."

1.4 Equilibrium Concept

In general in this paper, we consider Subgame Perfect Nash Equilibria. Particularly in our analysis related to the *corner cases* of auctions, we perform backward induction, either to derive an indifference condition in the mixed-strategy equilibrium, or determine a pure strategy if possible. As a convention, we primarily consider equilibria in this paper as a function of the termination probability p_k, which represents the probability the auction will end on the k^{th} bid. The point here is that to specify fully the equilibrium we must always define the individual mixed strategies that yield this termination probability. For ease of notation we often omit this step, primarily because our intuition is better tied to the termination probability than to the individual mixed-strategies. This ease of interpretation is directly attributable to the fact that a given termination probability can give reason to arbitrarily many mixed strategies unless further conditions are imposed. Where necessary or revelatory, we do describe the structure of the equilibrium with respect to individual mixed-strategies, and we always make the symmetric case our example if one is given.

The other key equilibrium idea in this paper is the distinction between *natural equilibria* and *trivial equilibria*. The former refer to equilibria in which there is a nonzero probability of reaching any given time step (not including time steps for which no one has an incentive to bid). The latter refers to equilibria that backward induct into the degenerate cases of auctions terminating either with 0 bids placed, or with 1 bid placed, after which the auction ends with certainty. The reason for the distinction will be very obvious in our usage: the natural equilibria are those we use for predictions, as we think they more accurately represent how agents might actually play.

Finally, unless otherwise specified, we consider only equilibria that satisfy our *anonymity principle*. This merely requires that players cannot condition their strategies on the identity of the current leader. In a few instances we discuss what happens to the set of equilibria if we relax this anonymity condition. In general, though, we are more interested in the mixed-strategies themselves. We prefer to have a fuller specification of the relevant equilibria than an underspecified set of the complete equilibria.

2 A Simple Auction Model

We begin with what might be considered the first-best of the Swoopo problem: a simplistic model treating the bidders as symmetric. Admittedly, the most striking characteristic of the Swoopo auction format is the outsized firm profits, an observation readily available to the inquisitive visitor to Swoopo's website. However, any game theoretic solution with full information amongst rational, risk-neutral agents will not be able to explain firm profits. Because firm profits in equilibrium require equilibrium losses for at least one player, who has the opportunity not to participate, we cannot expect to encounter an equilibrium with profits for the auctioneer.

Regardless, in order to explain agent behavior, we need a baseline model, preferably one emphasizing simplicity, albeit perhaps at the expense of making precise predictions. An analogous model was independently demonstrated in [10]. Because the model essentially suggests itself, it is also chosen as the baseline of [3] and [5].

2.1 The Model

There are n bidders participating in a fully symmetric auction with full information. Each bidder values the item being auctioned at v, and in order to win the item, each non-leading agent may place bids that can be purchased for b. The current leader, if there is one, is not permitted to bid. Each auction can be described by an incremental (i.e., price-incrementing) parameter d. For example, the price of the item after the k^{th} round, paid only in the state of the world where the auction has terminated, by the winner is kd.

For simplicity, we treat the structure of payment for bids as an irrelevant detail of accounting. In particular, the placement of a bid amounts to a cost of b at the time of bidding. The auction begins with an initial price 0, and the auction takes the form of sequential bidding. In each period, all participating agents may elect to pay this cost b for the right to bid, which increments the price by d. The termination condition of the auction requires that a period expire without an additional bid. Assuming there is an incumbent bidder, he has "won" the item, pays the *final price* kd, where k is the round of termination, and receives the item valued at v. If there is no leading bidder, which is to say the auction terminated in round 0, the auctioneer retains the item and collects no revenue. Note that the assumption that the winning bidder pay the *final price* of kd is incomplete; the aforementioned cases of "100% off" and "fixed price" auctions clearly differ in this regard, and will be treated as special cases with a fixed final price, and an increment parameter $d = 0$. So that the auctioned item has incentive to any players, we assume that $v > b + d$. This assumption is not restrictive: were it not met, the equilibrium is trivial, and involves no agent bidding.

The last component of describing the auction is the tie-breaking procedure. In practice, this aspect of the auctions is quite unclear, as it relies on the consistency and thoroughness of the auctioneer. However, for this model, we assume that if two or more players opt to bid in a given round, the tie-breaking is done uniformly. That is, for m bidders, all of whom have placed a bid in a given round, each has probability of $\frac{1}{m}$ of being the leader in the subsequent round. Obviously, only the chosen bidder, or leading bidder, has to pay his bid cost b. More to the point, after this tie-breaking procedure, it is treated as if the non-leading bidders had not placed a bid.

For now, we consider the standard case of $d > 0$. Denote by $\bar{k}$ the upper bound at which a rational player will place a bid at all, which is just $\bar{k} = \frac{v-b}{d}$. Indeed, for $k > \bar{k}$, we easily observe that the final price of the item to the winner is $kd + b > v$ which is strictly worse than not bidding. Since bidding rounds are necessarily integral, we consider the value $\hat{k} = \lfloor \bar{k} \rfloor$. If $\hat{k} = \bar{k}$ is integer, then players are indifferent to bidding and not bidding for $k = \bar{k}$, since they can conclude that their bid will win (no rational opponent can bid in $k > \bar{k}$) and guarantee zero profit.

If $\hat{k} < \bar{k}$, this complicates our analysis. To illustrate this complication, we consider the backward induction this causes. Since players *pass* (i.e., elect not to bid) with certainty in $\hat{k} + 1$, all players will bid with probability 1 in round $\hat{k}$, guaranteeing an expected profit (for each of n players) of $\frac{1}{n}[v - \hat{k}d - b] > \frac{1}{n}[v - \bar{k}d - b] = 0$, by construction of $\bar{k}$. Thus, with players bidding with certainty in round $\hat{k}$, we conclude that no player will rationally bid in $\hat{k} - 1$. Clearly, this logic descends for lower rounds, with every player bidding in rounds of the form $\hat{k} - 2z$ and every player passing in rounds of the form $\hat{k} - 2z + 1$ for integer z. Clearly, then, the auction terminates in period 0 or 1, depending only on the parity of $\hat{k}$. For even $\hat{k}$, the auction terminates with no one bidding, as every player passes, and the auctioneer keeps the item. For odd $\hat{k}$, every player will bid with probability 1 in round 1, after which the auction will terminate, with the randomly chosen agent profiting $v - b - d$.

The above description summarizes what we refer to as *trivial equilibria*. Because $\hat{k} < \bar{k}$ describes auctions with this backward induction, they will terminate no later than the end of the first round. With this observation in place, we are motivated to consider the case $\hat{k} = \bar{k}$, equivalent to assuming $\frac{v-b}{d}$ is integral. Obviously this assumption is restrictive. However, the equilibrium results are much more interesting in this context. Furthermore, [3] argues for the plausibility of this assumption, by demonstrating that the equivalent equilibrium of the game in which $\hat{k} = \bar{k}$ form an ϵ-equilibrium in the $\hat{k} < \bar{k}$ case, and that the two are numerically indistinguishable. The intuition behind the robustness of this latter set of more interesting equilibria is quite clear: the stage of profitable deviation appears in stage $\hat{k}$, which is reached with low probability, each player's bid is accepted with low probability, and the magnitude of the mistake is very small (one bid). For reasonable parameter values, [3] predicts the cost of these non-equilibrium strategies to be approximately 5 orders of magnitude smaller than the value of the auctioned item.

2.2 A Natural Equilibrium

Given the specification to this point, the model permits a very wide class of equilibria. Below, I will attempt to address the structure of the general equilibria. However, for now, I will limit my scope to those equilibria in which there is non-zero probability of the auction reaching each stage $k \leq \bar{k}$.

Denote by p_k the probability that there is no k^{th} bid placed, conditioned on there being $k-1$ bids placed. In other words, p_k is the probability that the auction terminates with price $d(k-1)$. Therefore, p_k envelops the behavior of all n players insofar as it represents the collective probability of none of them bidding at the k^{th} stage.

Now, with $\hat{k} = \bar{k}$, we consider equilibrium behavior. Recall, for any $k > \bar{k}$, we have $p_k = 1$, as no agents will bid and the auction terminates with certainty. For $k \leq \bar{k}$, since there is nonzero probability of the auction reaching that stage, at least one player must be mixing over his strategies: bid or not bid. Each of these agents mixing in the equilibrium must be made indifferent over his choices. This defines for $1 < k \leq \bar{k}$ the indifference condition:

$$(v - kd)p_{k+1} - b = 0. \tag{1}$$

We interpret this equation with p_{k+1} as the probability that this bid causes the auction to *terminate*, making the agent who placed it the winner. Likewise, we note that the indifference argument assumes a bidder in the lead to be made indifferent, but this requirement is not met for $k = 1$. Without an indifference condition, the value of p_1, the probability of anyone placing the first bid, can take on any value in the range $(0, 1]$. In general, $p_1 = 0$ also satisfies the properties of equilibrium, but we are here restricting ourselves to equilibria in which there is a nonzero probability of each stage $k \leq \bar{k}$ being reached. Thus, given this assumption, for any $\hat{p} \in (0, 1]$ we have constructed the equilibria of this auction:

$$p_k = \begin{cases} \hat{p} & k = 1 \\ \frac{b}{v-(k-1)d} & 1 < k \leq \bar{k} \\ 1 & k > \bar{k} \end{cases} \tag{2}$$

Defining the termination probability of the auction is really sufficient for a distributional representation of auction lengths. However, this set of equilibria lends itself fairly easily to describing the individual agents' strategy profiles. In short, recall that we have n agents, and for each agent i define the set of probabilities $q_{k,i}$, which represent the probability agent i will elect to *pass* in the k^{th} round. This notation is somewhat inconsistent with previous attempts, but the choice to describe the no bidding probability rather than the bidding probability should become evident.

Recall the realistic assumption that only a non-leading bidder may place a bid, which is the case for Swoopo. It also seems unlikely that players would outbid themselves, though we do note that [3] provides evidence that highly experienced agents exhibit a profitable learning effect with respect to aggressive bidding. Regardless, we ignore the reputational aspects of bidding, so that if there is a current leader, the other $n-1$ players comprise those whose aggregate probability of bidding must correspond to the probability of auction termination. Specifically, in the k^{th} round, we already identified the termination probability, p_k. And the probability of no one bidding, equivalent to the termination probability, is computed directly $p_k = \prod_{i \neq j} q_{k,i}$, where j is the leading bidder. For completeness, we observe that more equilibria can be permitted if we assume that each participating agent can observe the identity of the leader, thereby conditioning his bidding probability on that identity. For now, we maintain anonymity and abstract away the cumbersome process of conditioning strategies based on the current leader.

We linearize this equation as $\ln(p_k) = \sum_{i \neq j} \ln(q_{k,i})$, and note that all $q_{k,i}$ satisfying this equation form a mixed-strategy equilibrium. We note that this equation only applies for $1 < k \leq \bar{k}$. However, we also note that this equation must hold for each agent j being the leader of the auction. Thus, consider the substitution $s_i = \ln(q_{k,i})$ for a given k where $1 < k \leq \bar{k}$, and further denote $S = \sum_i s_i$. Thus, since p_k is a constant regardless of who is the leading bidder (the leader is assumed anonymous), we have that $\sum_{i \neq j} \ln(q_{k,i}) = \sum_{i \neq j'} \ln(q_{k,i})$, which can be rewritten as $S - s_j = S - s_{j'}$ for all j, j'. And of course, this implies that $s_j = q_{k,j} = q_{k,j'} = s_{j'}$, implying perfect symmetry in the mixing probabilities for $1 < k \leq \bar{k}$. So, by symmetry, we have that $q_{k,i} = (p_k)^{\frac{1}{n-1}}$.

For $k > \bar{k}$, $p_k = 1$ forces $q_{k,i} = 1$ for all i. For $k = 1$, we employ the same logic, but note that there is no leading bidder, so we have $\ln(p_1) = \sum_i \ln(q_{1,i})$. So, this only gives us one condition, whereas before we had n, allowing asymmetric bidding probabilities in the first round. Indeed, for all $q_{1,i}$ satisfying $\ln(p_1) = \sum_i \ln(q_{1,i})$, we have an equilibrium.

Thus, in the equilibrium with fixed $\hat{p} \in (0, 1]$, player strategies will take the form:

$$q_{k,i} = \begin{cases} q_{1,i} & k = 1 \text{ s.t. } \ln(\hat{p}) = \sum_j \ln(q_{1,i}) \\ \left(\frac{b}{v-(k-1)d}\right)^{\frac{1}{n-1}} & 1 < k \leq \bar{k} \\ 1 & k > \bar{k} \end{cases} \tag{3}$$

This equilibrium is largely symmetric, with the only asymmetries permitted in the first round bidding behavior, as a result of there being no leader to be made indifferent. Indeed, we note the fully symmetric equilibrium with $\hat{p} \in (0,1]$ is of the form:

$$q_{k,i} = \begin{cases} (\hat{p})^{\frac{1}{n}} & k = 1 \\ \left(\frac{b}{v-(k-1)d}\right)^{\frac{1}{n-1}} & 1 < k \leq \bar{k} \\ 1 & k > \bar{k} \end{cases} \tag{4}$$

Returning to the idea of simplicity, we note that the generality introduced by allowing the parameter $\hat{p}$ to remain unspecified is not especially insightful. Therefore, we will usually refer to the special case $\hat{p} = \frac{b}{v}$, since for $k = 1$, $v - (k-1)d = v$.

2.3 Fixed-Price Variant

Although the extension is trivial, we briefly describe the fixed-price analogue of the aforementioned equilibria. These are motivated by Swoopo's "100% off" and "fixed price" auctions, whose final prices are predetermined and do not increment. Thus, we have $d = 0$. Most of our analysis still holds, with the crucial caveat that there is no round after which bidding will terminate with certainty. Since the final price is fixed at some level $w < v$, every round after the first has identical structure. Namely, there is no $\bar{k}$ for fixed-price auctions.

Indeed, with this in place, we extend the natural equilibrium in the obvious way with $\hat{p} \in (0,1]$:

$$q_{k,i} = \begin{cases} (\hat{p})^{\frac{1}{n}} & k = 1 \\ \left(\frac{b}{v-w}\right)^{\frac{1}{n-1}} & 1 < k \end{cases} \tag{5}$$

We again are primarily interested in the simplest structure, where $\hat{p} = \frac{b}{v-w}$ so that each stage, including the first, has the same termination probability. In fact, with this constant probability in mind, we denote $\bar{p} = \frac{b}{v-w}$. The probability of the auction terminating at a given stage k, no longer conditional on it reaching that stage, is just

$$f(k;n) = p_{k+1} \prod_{i=1}^{k} (1 - p_i) \tag{6}$$

which follows immediately from observing that an auction terminating in the k^{th} round must simply survive the first $k-1$ rounds, and then terminate in the next round. The probability mass function in this our fixed-price case is $f(k;n) = \bar{p}(1-\bar{p})^k$, which is just the geometric distribution. So, the model predicts that auction lengths will be geometrically distributed for fixed-price auctions, with parameter $\bar{p} = \frac{b}{v-w}$. Using this we can easily derive the moments of the geometric distribution, which provide us easily with the moments of the revenue distribution. Specifically, we can compute the variance and skew of the revenue distribution, which has an obvious motivation into work on other utility functions.

2.4 Revenue Properties of the Baseline Model

With the observation that fixed-price auctions are geometrically distributed, we next turn to the probability mass function of the price-incrementing auction ending in the k^{th} round. The distribution of auction outcomes changes slightly, as the termination probability is increasing in the round number, whereas in the fixed-price problem, the termination probability remained constant at each stage of the game. So substituting the price-incrementing termination probabilities, we have probability mass function:

$$f(k;n) = p_{k+1}\prod_{i=1}^{k}(1-p_i) = \left(\frac{b}{v-kd}\right)\prod_{i=1}^{k}\left(1-\frac{b}{v-s(i-1)}\right) \tag{7}$$

This immediately gives us expected revenue

$$\mathbb{E}[Rev] = \sum_{k=1}^{\bar{k}}\left[(b+d)kp_{k+1}\prod_{i=1}^{k}(1-p_i)\right] = \sum_{k=1}^{\bar{k}}\left[(b+d)k\left(\frac{b}{v-kd}\right)\prod_{i=1}^{k}\left(1-\frac{b}{v-s(i-1)}\right)\right] \tag{8}$$

This expectation can be computed in a straightforward, yet tedious manner.

A more enlightening approach is to consider a story, akin to the earlier observation that the ability of each agent to decline participation bounds the firm's revenue above by v. Assume that the first bid is indeed placed. Subsequently, every bidder is made indifferent from our construction of the equilibrium. Therefore, the expected payoff from bidding for every player is exactly 0. Since the sum of the expected payoffs to the bidders and the auctioneer must in total be 0, we conclude that the expected profit of the auctioneer, given the first bid is placed, is also 0. Thus, the expected revenue, given an initial bid, is just v. However, there is $\hat{p}$ probability that no bidder places an initial bid, in which case the auctioneer's revenue is instead 0. Of course, in the Swoopo framework, the firm keeps the item in this case, so it is not a particularly meaningful case. Still, for completeness, this gives expected revenue of $(1-\hat{p})v$. In the familiar case of $p_1 = \hat{p} = \frac{b}{v}$, $p_1(0) + (1-p_1)(v) = 0 + (1-\frac{b}{v})v = v - b$. We note that this revenue solution applies equally well in both the price-incrementing and fixed-price auctions. For the fixed price case, expected revenue is of the form

$$\mathbb{E}[Rev] = \sum_{k=1}^{\infty}\left[(b+d)kp_{k+1}\prod_{i=1}^{k}(1-p_i)\right] = \sum_{k=1}^{\infty}\left[(b+d)k\left(\frac{b}{v-w}\right)\prod_{i=1}^{k}\left(1-\frac{b}{v-w}\right)\right] \tag{9}$$

Clearly, given that the firm retains the item when no bids are placed, we conclude in practice that firm profits are indeed 0. As discussed at the introduction of this section, the most dissatisfying aspect of this model is its failure to describe the profits realized by Swoopo as the auctioneer. Regardless, we have a theoretical framework from which we can make precise predictions about the auction empirics, with the probability mass function of auction termination even more important than the revenue observations.

2.5 Complete Description of Equilibria

Up to this point, we have emphasized *natural equilibria*, conveniently defined here to include only those equilibria that predict each feasible (i.e., $k \leq \bar{k}$) stage of the game has nonzero probability of being reached. Based on work in [9] and [3], we now fully describe the class of equilibria. The key observation is that we can simply combine the two different analyses we have employed thus far in the following manner.

Consider an arbitrary stage $\tilde{k} \leq \hat{k} = \bar{k}$, since we have retained the integrality assumption, after which players bid according to the indifference condition. Therefore, by the arguments in our natural equilibrium analysis, we have that $p_k = \frac{b}{v-(k-1)d}$ for $k \in (\tilde{k}, k]$. Similarly, in round $\tilde{k}$, which is now analogous to stage 1 of the natural equilibrium model, agents have expected payoff 0 of bidding (and of not bidding) because the indifference conditions in $k > \tilde{k}$.

So we return to the general case of $\hat{p}$ as the stage 1 termination probability, which now refers to the stage $\tilde{k}$ termination probability. Finally, to analyze the phase of the game prior to $\tilde{k}$, we revisit the analysis of *trivial equilibria*, where players backward induct and the game does not last beyond the end of the first stage.

Now, if $\hat{p}(v-(\tilde{k}-1)d) < b$, no agent will want to bid in round $\tilde{k}-1$, because the leading bidder in that round has expected payoff: $\hat{p}(v-(\tilde{k}-1)d) - b < 0$. Therefore, in round $\tilde{k}-2$, every agent will bid with probability 1, since the auction terminates with certainty in round $\tilde{k}-1$. Thus, as predicted, we use our backward induction from the trivial equilibria before to argue that the auction either terminates at the beginning or end of the first stage. Similar analysis holds for $\hat{p}(v-(\tilde{k}-1)d) > b$. Finally, note that the case of $\hat{p}(v-\tilde{k}d) = b$ implies $\hat{p} = \frac{b}{v-(\tilde{k}-1)d}$, which is equivalent to setting $\tilde{k}$ one lower, so that the natural equilibrium's indifference condition applies.

Thus, the termination probabilities of the auction are fully specified. Controlling for the parity of $\tilde{k}$, which as before affects whether the trivial equilibrium ends at the beginning or end of the first stage, the general equilibria are nothing more than a combination of the two separate cases we have already addressed. By simply choosing the stage at which the game switches form (from the trivial equilibrium to the natural equilibrium) $\tilde{k}$, the game is completely described.

3 A General Asymmetric Model

Thus far, our attention has been focused primarily on symmetric interpretations of the auctions. Although we have discussed the construction of all equilibria in Section 2.5, we have emphasized the symmetric case. More to the point, we have exclusively considered the case where every agent faces the same valuation, bid cost, and bid increment in each auction (i.e, the parameters v, b, and d). Although our empirical analysis is rooted in the predictions of the symmetric model, we are interested in describing a less restrictive structure for the auction equilibrium. Moreover, given the inadequacy of the revenue predictions, we are motivated to propose an alternate analysis of the game.

The only treatment of asymmetric penny auctions appears in [5], in which asymmetry is presented as a plausible explanation for the apparent profits in Swoopo auctions. Specifically, their work focuses on informational asymmetries, in which the agent population is split into two subpopulations, which have asymmetric knowledge of the parameters facing the other subpopulation. Although this focus on informational asymmetries has its own motivations, and indeed provides interesting results we examine later, for now we attempt to solve a model that allows full asymmetry in a perfect information environment.

3.1 Motivation

As is the case in any theoretical-empirical correspondence, many aspects of real-world Swoopo auctions severely complicate a straightforward analysis. Given its unusual nature and apparently wild profitability, Swoopo has captured a large degree of scrutiny [8]. Most published complaints amount to observers marveling at the auction format as a (questionably) legal form of gambling, a criticism we address later. Others object to the obvious potential for shill bidding, an aspect of great interest to [5], [3], and [10], which we ignore in our analysis as it offers very little in the way of predictions.

The bigger complaints about Swoopo ought to be rooted in the asymmetry of the parameters and the failure of Swoopo to describe precisely the microstructure of auction implementations. We offer a brief overview of these subtleties of Swoopo to motivate the asymmetric model.

Estimating v is especially difficult, and is a common problem facing auction analysis. For each auction, Swoopo offers a retail price for the item, ostensibly serving to anchor the agents' valuation at this level. The most obvious issue is that the listed price is generally far above the prevailing market price. [3] shows that for a substantial majority of auctioned items, a simple Amazon search implies Swoopo's retail price overstates price by 35%. This of course, is in addition to the standard problem of different agents in an auction having varying valuations of the item.

Similar problems arise around bid costs b, which [5] argues can differ widely for users. They argue that FreeBids auctions (auctions where the item at stake is simply a pack of Swoopo bids) have the effect of transmitting bids at a cost lower than the selling rate b. While it is true the winners generally receive the bids at a steep discount, the revenue of these auctions far exceeds the bid cost. Moreover, this analysis requires a non-standard behavioral model in which agents treat the value or cost of bids as different from the market price of bids, which in our view seems implausible. However, [5] does provide evidence of a one-day discount (a holiday promotion), where Swoopo offered a price reduction 15%. In principle agents would forecast their bid needs and purchase in this environment, but given uncertainty and budget constraints, we can easily imagine this rebate suggests that the cost of bids can indeed vary greatly across agents.

However, we argue that the smoking gun of asymmetry on Swoopo is its international structure. Like any ambitious business, particularly an internet business, Swoopo has sought increased market share by offering auctions in a variety of countries. At the time of this writing, Swoopo has websites in the US, Canada, the UK, Spain, Germany, Austria, and South Korea. The surprising (and perhaps perverse) fact is that Swoopo offers most of their auctions at the same times in all of these markets. Their motivation here is obvious: the different times of day mitigate lulls in activity, the aggregation of bidders is crucial to supporting more auctions, and it is administratively easier to provide (mostly) the same auctions across the websites. Although rules and regulations in each country differ slightly, the primary issue involves the value and cost parameters in each market. In short, Swoopo bids in the UK cost .50 pounds, in EU countries, .50 euros, in the US .60 dollars, and so on. Furthermore, the same auction viewed at www.swoopo.com, which might read current price 10 will have a current price of 10 pounds in the UK, which are substantially different. Obviously, the value of items v is also variant across countries, though this is primarily just an extension of the variance in individual valuations. Finally, just like b, the increment parameter d varies across the 7 different websites.

3.2 The Model

With the likelihood of asymmetries in auction parameters hopefully well-motivated, we proceed to the model. The structure of the auction is exactly as in Section 2.1, with the same tie-breaking procedure for multiple bids in a single round. Each player now has an individual valuation of the item, which we denote v_i for $i = 1, 2, \ldots, n$. Similarly, we consider the cost of each bid to vary across players, so that the bid cost is b_i for each agent, as well as the price increment d_i.

Analogous to the original model, we define $p_{k,i}$ to be the probability that no one places the k^{th} bid given that bidder i is the current leader. Similarly, $q_{k,i}$ denotes the probability that bidder i passes (does not bid) in the k^{th} round given that he is not leading after $k-1$ rounds.

We begin by setting our attention to the more complicated case where $d_i > 0$, and will treat the fixed price case of $d_i = 0$ for all i as a consequence of these results. We make the simplifying assumption that either: $d_i > 0$ for all i, or $d_i = 0$ for all i. This seems quite obvious, as d_i represents an incremental parameter. Although its magnitude differs across countries, there are no known instances of an auction being conducted as fixed-price in one country and price-incrementing in another.

Now, we proceed to define the upper bound for each player's bidding range. That is, we define $\bar{k}_i = \frac{v_i - b_i}{d_i}$, which represents an upper bound for the rounds k that a player can profitably bid. For ease of notation, we apply an arbitrary reordering of the players, so that we have a monotonic ordering of the $\bar{k}_i$. Thus, we reorder player identities and parameters so that $\bar{k}_1 \geq \bar{k}_2 \geq \ldots \geq \bar{k}_n$. Also, motivated by our previous methodology, we define $\hat{k}_i = \lfloor \bar{k}_i \rfloor$ for all i.

Case 1. *Where* $\hat{k}_1 < \bar{k}_1$.

> This case is analogous to the symmetric case. When round $\hat{k}_1$ is reached, player 1 will bid with probability 1 since no opponent j can bid, since $\bar{k}_j \leq \bar{k}_1 < \hat{k}_1$ for all j. Thus, with player 1 always bidding in $\hat{k}_1$, we have no agents willing to bid in $\hat{k}_1 - 1$, and backward induction continues into a *trivial equilibrium*. Specifically, we have for odd $\hat{k}_1$ an auction terminating at the end of round 1, and for even $\hat{k}_1$ an auction terminating at the beginning of round 1 (i.e., with no bidding occurring and the auctioneer keeping the item).

Case 2. *Where* $\hat{k}_1 = \bar{k}_1$. *We consider three subcases.*

Subcase 1. *Where* $\bar{k}_2 = \bar{k}_1$

This is the simplest case, and it corresponds nicely with the equilibria treated in the symmetric model. We have the indifference condition at $\bar{k}_1$, which we can extend downward. So, we consider this an analogue of the complete description of equilibria in Section 2.5, which I will detail below.

Subcase 2. *Where* $\bar{k}_2 \in [\bar{k}_1 - 1, \bar{k}_1)$

This case also permits the analogues of the complete equilibria. However, this set of equilibria correspond only to the case where $p_{\bar{k}_1,2} = \frac{b_2}{v_2 - (\bar{k}_1 - 1)d}$. If this is not the case, we cannot have indifference and the game returns to the trivial equilibria ending at the beginning or end of the first period.

Subcase 3. *Where* $\bar{k}_2 < \bar{k}_1 - 1$

This case is also simple, and a bit striking. At round $\bar{k}_1 - 1$, player 1 will bid with certainty, and no other player may bid since $\bar{k}_2$ is the largest among them. Thus, no opponent will bid in round $\bar{k}_1 - 2$. In this variant of the trivial equilibrium, the backward induction implies that player 1 will bid with probability 1, while all other agents must refrain from bidding. The reason is precisely because player 1 has an assured victory by bidding in round $\bar{k}_1 - 1$.

Similar to how we analyzed the symmetric case above, we realize that in describing these corner cases, we suppress full descriptions of equilibrium. In particular, those equilibria that are classified as trivial are relatively uninteresting, and we summarize them in the sense that they terminate with certainty either with no one bidding, or with a single agent winning at the end of period 1. We note that Subcase 3 is a slightly more interesting version of the trivial equilibrium. Here, the intuition is that the parameters facing player 1 provide a sufficiently higher valuation that no other agents can compete. The result is that every equilibrium has player 1 bidding with certainty, and consequently no other agents bidding. Albeit extreme, this motivates the potentially strong impact of different valuations.

Now, we must describe the equilibria which occur in Subcases 1 and 2, which match the standard case where the indifference condition is met. We proceed with analysis similar to that in Section 2.2. We first consider an arbitrary period of the game $\tilde{k} \leq \bar{k}_1$. This $\tilde{k}$ represents the transition between the alternating stages of bidding and not bidding, versus the second phase of the game, during which players are indifferent at each stage where they bid with positive probability.

We observe that for $k > \bar{k}_1$, $p_{k,i} = 1$ for all i. However, we note that we cannot say the same for $k > \bar{k}_i$; the reason being that other agents with higher valuations may still be able to bid. So we cannot exactly pin down $p_{k,i}$ in general for $\bar{k}_i < k \leq \bar{k}_1$. Fortunately, this does not affect our analysis for the *termination probabilities*. Specifically, these cases are irrelevant, as i will not bid in this range anyway.

Furthermore, we note that for k on the interval $(\tilde{k}, \bar{k}_i]$, the indifference condition holds. Finally, we can describe this equilibrium by choosing, for each agent i, a probability $\hat{p}_i$ which applies at stage $\tilde{k}$. We recall that the indifference in this stage, and the fact that the indifference condition does not hold in $\tilde{k}-1$ leaves this probability unspecified. However, we recall the parity argument from before: namely, we have alternating periods of bidding and not bidding, which differ based on whether $\hat{p}_i$ is greater than or less than $\frac{b_i}{v_i-(\tilde{k}-1)d_i}$ (the equality case again corresponds to decrementing $\tilde{k}$, thereby moving the indifference case downward).

With this in mind, we define the equilibria as follows, using the structure from [3]. Choose $\hat{p}_i \in [0, 1]$ arbitrarily, as well as an arbitrary $1 \leq \tilde{k} \leq \bar{k}_1$, and define

$$\alpha_i = \begin{cases} 0 & \hat{p}_i < \frac{b_i}{v_i-(\tilde{k}-1)d_i} \\ 1 & \hat{p}_i > \frac{b_i}{v_i-(\tilde{k}-1)d_i} \end{cases} \tag{10}$$

Then, the set of equilibria must satisfy the following termination conditions:

$$p_{k,i} = \begin{cases} 0 & k < \tilde{k} \text{ s.t. } (\tilde{k} - k) + \alpha_i \equiv 0 \mod 2 \\ 1 & k < \tilde{k} \text{ s.t. } (\tilde{k} - k) + \alpha_i \equiv 1 \mod 2 \\ \hat{p}_i & k = \tilde{k} \\ \frac{b_i}{v_i-(k-1)d_i} & \tilde{k} < k \leq \bar{k}_i \end{cases} \tag{11}$$

We offer a few observations. First, we reiterate that the only incompleteness is in describing $p_{k,i}$ for $k > \bar{k}_i$. This is an irrelevant case, and if we solve for equilibrium strategies, we avoid this problem by restricting our analysis to the set of agents remaining in the game (i.e., those players for whom $\bar{k}_j \geq k$). Second, these termination probabilities are indeed confusing, but they are viewed most simply as the straightforward generalization of the complete equilibrium in the symmetric case. Third, in practice, we would solve the equilibrium via backward induction: beginning at stage $\bar{k}_1$ we identify the indifference condition moving backward, and incorporate the fact that the set of agents in the auction changes as we backward induct, always including only those agents i for whom $\bar{k}_i \geq k$.

3.3 Toward a Solution

In general, we recall the comments in Section 3, where we described the extremely large set of equilibria permitted by the given termination conditions. Still, for completeness, we will attempt to describe the process of deriving equilibrium strategies based on the termination probabilities. As before, we focus on the natural equilibria of Section 2.2, corresponding to the indifference condition phase of the game (i.e., when $\tilde{k} = 1$). If $1 \leq k \leq \bar{k}_1$ there is a nonempty subset of *remaining players*. Without loss of generality, I will describe the case in which all n players remain, which implies $1 \leq k \leq \bar{k}_n$ so that all players have positive payoff from winning the auction.

We employ the same indifference solution technique, observing that $b_i = p_{k,i}(v_i - (k-1)d_i)$, where d_i denotes the bid increment, indexed to each agent. We further observe that the auction terminating at this stage requires all bidders $j \neq i$ electing not to bid at the k^{th} stage, an event with probability $\prod_{j \neq i} q_{j,k,i}$. We recall the complication of anonymity versus full information: when the identity of the leader is known, players may condition their strategies on this information; anonymity restricts the set of equilibria.

Describing the equilibria when players can condition their strategies is to describe a very large set. However, we can do so easily by realizing that for each leader i, a player has probability of passing (not bidding) in stage k: $q_{j,k,i}$. Thus, we require in equilibrium only that

$$p_{k,i} = \frac{b_i}{v_i - (k-1)d_i} = \prod_{j \neq i} q_{j,k,i}, \tag{12}$$

which holds for a given stage k (which we are backward inducting through) and all players i.

As it is more restrictive, we can describe more fully what the anonymous equilibrium looks like. In this case, we require that $\ln(p_{k,i}) = \sum_{j \neq i} q_{k,j}$, where $q_{k,j}$ is defined as usual: the probability of player j not bidding in round k, given that he is not the leader. Because this condition must hold for all potential leaders i, we have a set of n equations in n unknowns, the $q_{k,j}$. We can solve this system by linearizing. Specifically, for each round k, we have a system of n linear equations that define the mixing probabilities of each agent. Then, by introducing logarithms, for each agent i we have indifference condition:

$$\sum_{j \neq i} \ln(q_{k,j}) = \ln\left(\frac{b_i}{v_i - (k-1)d_i}\right) \tag{13}$$

To solve, we introduce

$$t_{k,j} = \ln(1 - q_{k,j}) \text{ and } r_{k,i} = \ln\left(\frac{b_i}{v_i - (k-1)d_i}\right) \tag{14}$$

So, the set of n indifference conditions for our fixed k is just $\sum_{j \neq i} t_{k,j} = r_{k,i}$ for all $i \in \{1, 2, \ldots, n\}$. We denote the vector of all $t_{k,j}$ by $\overrightarrow{t_k}$, and the vector of all $r_{k,i}$ by $\overrightarrow{r_k}$.

We adopt the standard notation of I_n as the identity matrix of dimension n, and J_n as the matrix of dimension n in which every entry is a constant 1. Therefore, in matrix form, we have that $(J_n - I_n)\overrightarrow{t_k} = \overrightarrow{r_k}$. We observe that the inverse of $J_n - I_n$ exists for all $n > 1$, and takes the form

$$(J_n - I_n)^{-1} = \frac{1}{n-1} J_n - I_n \tag{15}$$

We confirm this because

$$(J_n - I_n)(\frac{1}{n-1} J_n - I_n) = \frac{1}{n-1} J_n^2 - \frac{n}{n-1} J_n I_n + I_n \tag{16}$$

(because $J_n I_n = I_n J_n$ and $I_n I_n = I_n$ by definition of identity). We verify the inverse relationship by concluding $\frac{1}{n-1}(nJ_n) - \frac{n}{n-1} J_n + I_n = I_n$, since $J_n^2 = nJ_n$.

Therefore, we have that $\overrightarrow{t_k} = (\frac{1}{n-1} J_n - I_n)\overrightarrow{r_k}$. Identifying each term in isolation, we conclude that

$$t_{k,i} = \frac{1}{n-1} \sum_{j=1}^{n} r_{k,j} - r_{k,i} \tag{17}$$

for each $i \in \{1, 2, \ldots, n\}$.

For notational convenience, denote $S_k = \sum_{j=1}^{n} r_{k,j}$ so that $t_{k,i} = \frac{S_k}{n-1} - r_{k,i}$. Substituting our logarithmic transformations, we have that

$$q_{k,i} = e^{\frac{S_k}{n-1} - \ln(\frac{b_i}{v_i - (k-1)d_i})} = \left(\frac{v_i - (k-1)d_i}{b_i}\right) e^{\frac{S_k}{n-1}} \tag{18}$$

Thus, in order for this equilibrium to be well-defined, we need that $\frac{S_k}{n-1} - r_{k,i} \leq 0$ for all $i \in \{1, 2, \ldots, n\}$.

3.4 Equilibrium Description

We note that the full set of equilibria, much like in the symmetric model, contains a large set of equilibria we find uninteresting. Above, we offer the cases in which *trivial equilibria* develop, and even discover a new variant, in which the strongest-valued player is assured first round victory. This differs from the *trivial equilibria* seen to this point, in that previously all players chose to bid or not bid with certainty in the first round, and in the latter case each received the item with uniform probability.

In our subcases, we identify the conditions under which we can specify equilibria of the form in Section 2.5: namely when $\hat{k}_1 = \bar{k}_1$ and $\bar{k_2} \in [\bar{k}_1 - 1, \bar{k}_1]$. Furthermore, when $\bar{k}_2 = \bar{k}_1$ this analysis covers all equilibria. When $\tilde{k} = 1$, these equilibria are again *natural* in the sense that each stage up to $\bar{k}_1$ has nonzero probability. Moreover, for general $\tilde{k}$, we are still primarily interested in the latter phase of the game, during which players are employing more interesting mixed-strategies.

We recall in our solution in Section 3.3 we assumed that $k \leq \bar{k}_n$ so that all players remain. Moreover, this immediately gives us the necessary condition that $r_{k,i} < 0$. In the case that $k > \bar{k}_n$, we observe that the game has changed: namely, we exclude all agents i for whom $k > \bar{k}_i$, and proceed with the auction. The remaining agents still bid according to the same termination probabilities.

The only condition not already accounted for is $S_k \leq (n-1)r_{k,i}$ for each agent i remaining. We only need $S_k \leq (n-1)r_k$ where $r_k = \min_i(r_{k,i})$. The intuition for this restriction is simple: if one of the bidders has a termination probability (of indifference) that is too low, it will be inconsistent with the rest of the termination probabilities. Another way of seeing this is to imagine how the probabilities are actually built: the difference between two players' indifference conditions is limited in the sense that this change only reflects swapping one player for the other. In the extreme case, we realize that if a player is willing to bid with extremely small termination probability, and no others are, he cannot be made indifferent by the mixed strategies the other players will be willing to play. Therefore, the *strength* of the bidder with minimal $r_{k,i}$ can break the equilibrium, forcing us into a *trivial equilibrium*.

3.5 The Fixed-Price Variant

In this context, extending these results to the fixed-price auction proves straightforward. While there remain the trivial equilibria, we no longer have to worry about the corner cases involving the last round of bidding, different for each agent, $\bar{k}_i$. This simplifies matters immensely. In order to characterize the natural equilibria, we need only consider that $p_{k,i} = \frac{b_i}{v_i - w_i} = \prod_{j \neq i} q_{k,j}$. Thus, this corresponds precisely to the linearizing solution technique employed above. Furthermore, we avoid the complication of descending via backward induction because the probability of passing is constant at each stage of the auction k.

Thus, we can fully describe the general asymmetric equilibria. For any $\hat{p}_i \in [0,1]$, the set of equilibria must satisfy the following termination conditions:

$$p_{k,i} = \begin{cases} \hat{p}_i & k = 1 \\ \frac{b_i}{v_i - w_i} & k > 1 \end{cases} \tag{19}$$

We note that what follows is valid for general $\hat{p}_i$, but for consistency across all periods, we choose $\hat{p}_i = \frac{b_i}{v_i - w_i}$.

Our linearizing technique can be used to solve for the actual mixed strategies

$$q_{k,i} = e^{\frac{S}{n-1} - \ln(\frac{b_i}{v_i - w_i})} = \left(\frac{v_i - w_i}{b_i}\right) e^{\frac{S}{n-1}} \tag{20}$$

where $S = \sum_{j=1}^{n} \ln\left(\frac{b_i}{v_i - w_i}\right)$ no longer depends on k. Of course, we still have the restriction that $S \leq (n-1)\ln\left(\frac{b_i}{v_i - w_i}\right)$ for all i, to prevent a sufficiently *strong* bidder from breaking the equilibrium.

4 Imperfect Information

As discussed in the introduction to Section 3, [5] introduces the notion of imperfect information. Their coverage focuses on *asymmetry*, which they use in a variety of contexts. In the previous section, we presented and solved a completely general model of asymmetry with perfect information. By contrast, [5] treats primarily informational asymmetries in the context of two subpopulations, and again leverages multiple meanings of asymmetry. In one context, they consider subpopulations facing different valuations v^A and v^B, where informational asymmetry means each subpopulation assumes the entire population shares their valuation. Using this contextualized meaning of asymmetry, we had previously solved for equilibrium in a more general case, a result presented below.

In an alternative context, [5] considers each subpopulation facing a different bid cost, where the high-cost group remains unaware of the asymmetry, but the low-cost group has perfect information (crucially, they are aware of the existence of the high-cost bidders). After summarizing the methodology they employ, we derive a more general result in a straightforward manner. In doing so, we expose a key feature of the asymmetric analysis.

Using this observation, we demonstrate that one of their most interesting results is actually a degenerate case of our generalization. We credit their emphasis on a key feature of this case, namely the convexity of the revenue function with respect to misestimates of the population size. The purpose of summarizing this degenerate case is to motivate our stylized interpretation of auction population dynamics, which interfaces well with their insight.

Briefly, we note that all subsequent models, both in this section and otherwise, are presented in the fixed-price variant. Unless otherwise specified, our analysis does indeed extend (in the obvious manner) to price-incrementing auctions. In the interest of emphasizing the key predictions and high-level effects of our variations, we restrict ourselves to solving for the fixed-price case. Moreover, unless otherwise specified, we abstract the generality permitted in choosing the first-stage termination probability $\hat{p}$ that defines equilibrium. This has no effect on the analysis, but allows us to use more compact notation in referring to termination or mixed-strategy probabilities independent of k.

4.1 Two Groups Unaware of One Another

We return to our first Swoopo auction model, but now we split the population of n agents into two subpopulations, which we denote A and B. There are k bidders in group A, who face parameters v_A, b_A and w_A. The remaining $n-k$ in B face v_B, b_B and w_B. Every agent assumes that the auction specification he is following applies to everyone identically. Therefore, every bidder assumes that the remaining $n-1$ bidders share his type.

In this presentation, we retain the assumption that agents are risk-neutral. However, for the sake of completeness (and a subsequent result), we note that no aspect of this analysis depends on any particular aspect of risk neutrality; indeed, each subpopulation may have its own utility function.

We introduce a new concept of *assumed termination probability*, which we denote by π_A and π_B. As the name suggests, each represents the termination probability that, in the natural (symmetric) equilibrium, make agents of that type indifferent. They are assumed, and not *actual* termination probabilities because these agents are incorrect in believing their competitors to be of the same type.

So, each A assumes the termination probability will satisfy the indifference condition

$$\pi_A(v_A - w_A - b_A) + (1 - \pi_A)(-b_A) = 0 \tag{21}$$

Because of their assumption, this assumed termination probability exactly corresponds to the actual termination probability when all agents are of type A. Regardless, we have $\pi_A = \frac{b_A}{v_A - w_A}$. With this in mind, each A agent derives their mixed-strategy profile, electing to pass with probability

$$q_A = (\pi_A)^{\frac{1}{n-1}} = \left(\frac{b_A}{v_A - w_A}\right)^{\frac{1}{n-1}} \tag{22}$$

Note that we ignore the corner case of stage $k = 1$, as it only differs in terms of their being n agents as potential bidders rather than $n - 1$. We henceforth overlook this case, by making the simple assumption that the first bid has already been placed. Conditioning on the first bid being placed only eliminates those cases in which no one bids and the auctioneer retains the item, which do not offer us any analytical insight.

We repeat the analysis, and although parameters differ, the structure of the game is identical for agents in A and B. So, we get mixed-strategy profile

$$q_B = (\pi_B)^{\frac{1}{n-1}} = \left(\frac{b_B}{v_B - w_B}\right)^{\frac{1}{n-1}} \tag{23}$$

Combining the full specification , this gives us the *actual* termination probability. However, this probability depends on the type of the current leader, which matters only here, as this is the only context in which we use the size of each subpopulation. This defines

$$p = \begin{cases} \left(\frac{b_A}{v_A - w_A}\right)^{\frac{k-1}{n-1}} \left(\frac{b_B}{v_B - w_B}\right)^{\frac{n-k}{n-1}} & \text{when leader is in } A \\ \left(\frac{b_A}{v_A - w_A}\right)^{\frac{k}{n-1}} \left(\frac{b_B}{v_B - w_B}\right)^{\frac{n-k-1}{n-1}} & \text{when leader is in } B \end{cases} \tag{24}$$

This fully specifies the natural equilibrium by describing the mixed strategies and the resulting termination probabilities. For more detailed computations, most notably revenue, we recognize the complication of having termination probabilities dependent on the type of the leading bidder. This issue is trivially resolved by realizing this system has the Markovian property of memoryless-ness. We suppress the derivation, but direct the curious reader to [5] which offers a straightforward derivation of the Swoopo Markov chain.

4.2 Two Groups: One Aware of the Other

We define our two groups as above, but with one key change to their set of information. Agents in A are aware of the size and existence of subpopulation B, while agents in B remain oblivious to the existence of A. As usual, we treat the leading bidders as anonymous to avoid agents conditioning their strategy against the identity (and type) of the current leader.

The analysis for agents in B has not changed at all from before, as their condition has not changed. They have the same assumed termination probability, which defines their probability of passing in the mixed-strategy equilibrium

$$q_B = (\pi_B)^{\frac{1}{n-1}} = \left(\frac{b_B}{v_B - w_B}\right)^{\frac{1}{n-1}} \tag{25}$$

And the actual probability of no agent in B bidding we denote $p_{B,A}$ when an A agent is leading, and $p_{B,B}$ when a B agent leads. These are simply

$$p_B = \begin{cases} p_{B,A} = q_B^{n-k} = \left(\frac{b_B}{v_B - w_B}\right)^{\frac{n-k}{n-1}} & \text{when leader is in } A \\ p_{B,B} = q_B^{n-k-1} = \left(\frac{b_B}{v_B - w_B}\right)^{\frac{n-k-1}{n-1}} & \text{when leader is in } B \end{cases} \tag{26}$$

The analysis for agents in A is a bit more subtle, to reflect their more detailed knowledge of the auction. Specifically, they must be made indifferent by the actual termination probability. Therefore, we have that $\pi_A = p = \frac{b_A}{v_A - w_A}$. We recognize that the probability of termination is exactly determined by the separate probabilities of no one in A bidding, and of no one in B bidding.

Conditioning on the case of the leading bidder being in A, we have $p = p_{A,A} p_{B,A}$ from which we can solve for using $p_{B,A}$ from above

$$p_{A,A} = \frac{1}{p_{B,A}} \left(\frac{b_A}{v_A - w_A}\right) = \left(\frac{b_A}{v_A - w_A}\right) \left(\frac{v_B - w_B}{b_B}\right)^{\frac{n-k}{n-1}} \tag{27}$$

We must be careful here, as there is nothing to keep $p_{A,A}$ from exceeding one, which is clearly not possible. This is a contingency that depends on the exact parameter values, which we detail below.

In the same manner we solve for $p_{A,B}$, which gives us finally

$$p_A = \begin{cases} p_{A,A} = \left(\frac{b_A}{v_A - w_A}\right) \left(\frac{v_B - w_B}{b_B}\right)^{\frac{n-k}{n-1}} & \text{if leading agent is in } A \\ p_{A,B} = \left(\frac{b_A}{v_A - w_A}\right) \left(\frac{v_B - w_B}{b_B}\right)^{\frac{n-k-1}{n-1}} & \text{if leading agent is in } B \end{cases} \tag{28}$$

We also require that $p_{A,B} \leq 1$ so that it is a well-defined probability.

To address this concern, we first note that $p_{A,A} = p_{A,B} \left(\frac{v_B - w_B}{b_B}\right) > p_{A,B}$, since they differ only in the number of A bidders who are left to bid. Therefore, it is sufficient to require that

$$p_{A,A} = \left(\frac{b_A}{v_A - w_A}\right) \left(\frac{v_B - w_B}{b_B}\right)^{\frac{n-k}{n-1}} \leq 1 \tag{29}$$

If this is not the case, and $p_{A,A} > 1$, the equilibrium breaks down as described. Specifically, as we saw in the general asymmetric model of Section 3.2, no A agent would be willing to bid in this environment. So, the auction degenerates into the agents of type B bidding in isolation. So, in this degenerate case, the impact of their misinformation is only that they assume there are n agents, when actually the $n - k$ agents are the full bidding population.

In the standard case, we have $p_{A,A} \leq 1$ and we finish off specifying our equilibrium. We recall that the termination probability must be equal to the probability of no agent in A bidding, multiplied by the probability of no agent in B bidding. We know these probabilities conditional on the leading bidder's type. However, this information is superfluous. Given that $p_{A,A} \leq 1$, we know that both cases satisfy the same $p = p_A p_B$ which is demanded by the indifference condition for A. Indeed, we have $p = \frac{b_A}{v_A - w_A}$.

This is result is remarkably clean, but this should come as no surprise. By requiring that $p_{A,A} \leq 1$, we essentially require that the corresponding parametric asymmetry (the difference in the v, b, and d parameters) be either in favor of A, or only modestly in favor of B. If this is the case, then A continues to be interested in bidding. And the fact that there are other A agents (which explains why we need $k \geq 2$) means they all hold one another to the same indifference condition as before. Alternatively, we can interpret this in the context of the termination probability for group A binding first (when $p_{A,A} \leq 1$), and allowing the A agents to employ the mixed strategy that retains their indifference.

When $p_{A,A} > 1$, the parameters do not permit any such mixing, and agents in A necessarily drop out, which leaves us with the degenerate case below. In the interest of finalizing our analysis of this model (covering this degenerate case), we proceed to summarize the [5] result that resolves this case.

4.3 Mistakes in Estimates of n

One of the most interesting variations examined in [5] concerns the case of auction participants forming incorrect assumptions about the number of opponents against whom they are bidding. The actual model is quite simplistic, but that is part of what makes it appealing. The obvious intuition is indeed correct, in that underestimating the population of bidders increases revenue (as agents overbid), while overestimating the population decreases revenue. The meaningful prediction, however, is the convexity of revenue in the magnitude of these mistakes.

Because we are completing our analysis from Section 4.2, we preserve that notation and introduce it into the technique of [5]. Specifically, there are $n - k$ agents in B, each of whom assumes there are n bidders in the population. These agents face the same parameters as before, with the set of agents in A no longer relevant, we return to the simpler notation of parameters v, b and w.

As we did above, we abstract away the corner case of the first bid. That is, we assume that the initial bid is placed, throwing out the cases where the auction terminates with no bids. Henceforth, we will make this assumption when calculating parameters related to auctions particularly revenue. The intuition here is natural: auctions terminating without bids are not particularly interesting (and indeed do not appear in our dataset), and we can describe distributional properties more simply if we condition on the auction have a winner.

Thus, agents will bid to the point where they are made indifferent (according to their mistaken beliefs), so that

$$\pi = \frac{b}{v - w} \tag{30}$$

where w is again the fixed price of this particular auction. Of course players use these probabilities to define their strategy profiles, so that we have $\pi = q^m$, or $q = (\pi)^{\frac{1}{m}}$.

However, the actual termination probability must take into account all n players that are actually participating. Therefore, for any stage $k > 1$, there is a leader who must survive against the other $n - 1$ agents, implying

$$p = \left(q^{\frac{1}{m}}\right)^{n-1} = \left(\frac{b}{v - w}\right)^{\frac{n-1}{m}} \tag{31}$$

So, with our natural equilibrium in place, we observe revenue. Fortunately the geometric distribution observation from Section 2.3 holds, so that

$$\mathbb{E}[Rev] = w + b + \sum_{k=0}^{\infty} bkp(1-p)^k = w + b\left(\frac{v - w}{b}\right)^{\frac{n-1}{m}} \tag{32}$$

using the expectation of the geometric distribution.

An interesting consequence of this revenue curve is that it is convex in m, the number of players estimated by the bidders. This convexity implies that population estimates with mean n, which are correct in the sense of the average case, produce additional expected revenue due to the variance in estimation. Although there is no particularly strong reason to think that players would be correct on average even though they are incorrect in particular cases, we still consider this convexity as a candidate explanation of auctioneer profits.

4.4 Auction Entry

A key feature of the mistakes made by players in the above misestimation is that they do not learn or update their beliefs in the auctions. In general, this violates the assumptions of a rational agent model; in practice, however, it seems plausible that agents might find it difficult to adjust their population estimates given their limited knowledge. Swoopo does display a count of an auction's *active users*, defined as the number of individuals placing a bid in the last 15 minutes. However, it is easy to see cases in which this vastly understates (or overstates) the actual population of bidders.

Thus far we have abstracted way the timer that Swoopo uses to govern its auctions, primarily because it can severely complicate the analysis. Most observers of a Swoopo auction will intuitively assume that the timer does indeed play a role. In the simplest (and therefore likeliest) account, we see that auctions with substantial time remaining have little to no bidding, whereas auctions with just a few seconds left are the most hotly contested. Similarly, it is probably insufficient to consider the time remaining as the sole determinant of the number of active bidders. Indeed, an auction whose timer has been reset to 10 seconds a hundred times almost certainly has a population of bidders much larger than an auction that has reached 10 seconds for the first time. There are a variety of explanations, the likeliest of which being that agents have limited attention, and are likely to join and remain in a single auction more often than jumping around across active auctions. Bid-level data strongly supports this claim.

4.5 A Stylized Model of Agent Entry

Although we are cognizant of the leap in assuming that agents consistently misestimate the size of their opposition, for this analysis we consider a stylized model in which agents have a constant assumption about the population size as n. In the early stages of the game, when the auction has not been actively bid on for long, we assume that the actual population is $m < n$. By contrast, in the later stages of the game, when the auction has seen active bidding, we assume the population is M, exceeds n, perhaps by a substantial amount. We choose to focus on the case $M = 2n - m$ so that we are dealing with a singe parameter m, and the mean of the population extremes is n.

Further stylizing our model, we consider a deterministic entry process. In the first round of bidding, there are m players. After each round, 1 player joins with certainty, until we have reached the maximal population size, which we have set at $M = 2n - m$. Assuming the auction reaches this point, we assume that these agents remain active in the auction until it has completed.

Because the agents' estimate of the population is constant n, their assumed termination probability is also a constant in the fixed-price environment. As above, we have $\pi = \frac{b}{v-w}$, which also defines the mixed strategy for each player, where he passes with probability $q = \left(\frac{b}{v-w}\right)^{\frac{1}{n-1}}$.

By contrast, the *actual* termination probability must be a function of the number of players participating. Therefore, for each using the q which are the employed mixed strategies, we have the following termination probabilities:

$$p_j = \begin{cases} \left(\frac{b}{v-w}\right)^{\frac{j-1}{n-1}} & j \in [m, 2n-m] \\ \left(\frac{b}{v-w}\right)^{\frac{2n-m-1}{n-1}} & j > 2n-m \end{cases} \tag{33}$$

We realize that once the $2n-2m$ case is reached, the population stabilizes at $2n-m$, and we have the standard underestimate case from Section 4.3. To handle this, we compute the probability of period $2n-m$ being reached, which is just $P = \prod_{j=m}^{2n-m}(1-p_j)$, and also denote by $\bar{p}$ the probability in the final (constant population) stage of the game: $\bar{p} = \left(\frac{b}{v-w}\right)^{\frac{2n-m-1}{n-1}}$. With this in mind, we can express revenue in this stylized auction (again, conditioning on one bid having already been placed):

$$\mathbb{E}[Rev] = w + b + \sum_{k=0}^{2n-2m} bkp_{k+m}\left[\prod_{j=m}^{m+k}(1-p_j)\right] + P\left(w + b + \sum_{k=0}^{\infty} bk\bar{p}(1-\bar{p})^k\right) \tag{34}$$

The last term in this revenue sum is of course just the revenue result from the previous part. Substituting, we have that

$$\mathbb{E}[Rev] = w + b + \sum_{k=0}^{2n-2m} bkp_{k+m}\left[\prod_{j=m}^{m+k}(1-p_j)\right] + P\left(\frac{b}{v-w}\right)^{\frac{2n-m-1}{n-1}} \tag{35}$$

Analytically, this revenue seems messy to interpret. Still, we can describe this revenue in terms of a few effects. The first is that the probability of reaching the $k = 2n-2m$ stage is actually lower than the probability of reaching the $k = 2n-2m$ stage when agents correctly identify there are n bidders. This follows directly from the fact that $P \leq p_n(1-p_n)^{2n-2m}$. More broadly, we can view this as a consequence of the AM-GM inequality: $pp' \leq \left(\frac{p+p'}{2}\right)^2$ with equality if and only if $p = p'$.

Second, we note that the actual revenue collected in the first $2n-2m$ stages is smaller in this entry model than it is in the standard perfect information model. This observation is related to the previous one, but the intuition follows directly from the fact that the inflated termination probabilities in the early stages overcome the effect of the depressed termination probabilities in the later stages, because the probability of reaching those later stages has gone down.

Third, however, we note that the dominant effect is that revenue, conditional on reaching the $2n-2m$ stage, has grown nonlinearly (growing rapidly). Indeed, in the original model, the revenue conditional on reaching this level would be $w + b(2n-2m) + (v-w) = b(2n-2m) + v$. However, with the population at $2n - m$ for the latter stage of the auction, the conditional revenue is now

$$\mathbb{E}[Rev|2n - 2m] = w + b(2n - 2m) + b\left(\frac{v - w}{b}\right)^{\frac{2n-m-1}{n-1}} \tag{36}$$

Outside of these revenue observations, the most important takeaway from this model is the distribution of auction terminations, as well as the allocation of profits. In the early stages, the auction terminates with higher probability than the standard model predicts, because for $j < n$, $p_j > p_n$. Therefore, since players are indifferent with respect to the overestimated n, the early bidders are actually receiving modest expected profits on bids placed prior to the $k = n - m$ stage, when the population reaches the estimate level of n. In the later stages, the termination probability (conditional on reaching those stages) falls off quickly. Furthermore, the expected profits on each of these bids is negative and decreasing in the round k of the game, until stage $k = 2n - 2m$, after which the constant population implies a constant level of expected loss on each bid.

4.6 Alternative Formulations

In the sense that the above model is stylized, we here offer a few obvious extensions that preserve the same key observations. The main extension to be considered is that the entry need not be deterministic, and need be bounded precisely in the range $[m, 2n - m]$. We can easily imagine having an agent probabilistically enter the auction at each time step with some probability β, and potentially having an agent probabilistically leave the auction at each time step with some probability $\gamma < \beta$.

Moreover, the assumption that the number of players in the auction be centered at n is entirely arbitrary, and was chosen primarily to isolate the effect of varying the population, and not confuse this with the effect of underestimation alone as described in [5]. A generalization of the model we present above might feature agents who have very slightly delayed information about the actual population of bidders. Specifically, we envision a simple symmetric model, where the population at a given stage k is just n_k. However, each player believes the population is n_{k-1}, where $k > 1$. In this environment players have relatively accurate knowledge of the bidding population, but it lags the actual population by one period. The motivation here is somewhat obvious, as players in practice might be slow to update their population estimates (especially when revising upward), since you only have evidence of a new opponent after his bid has been accepted.

Regardless, the key observation remains the convexity of revenue. If we assume the population drift is completely random and mean zero, the random variation will produce increased expected revenue. This follows in much the same way that the stylized model predicts.

In the extreme, we could consider agents with sticky estimations of the bidding population. Specifically, each player upon entering the auction has an estimation of the population, which we assume to be perfectly accurate. However, that player fails to update this estimation, and so is anchored to his initial (correct) value. The predictions of this model, again, do not differ qualitatively from those already offered, but they provide yet another account for how informational imperfections might develop.

5 Lottery Specifications

Although most Swoopo research is primarily motivated by the striking data on its revenues, we are not content to explain this feature in isolation. We have, to this point, modeled Swoopo quite literally, treating it as an auction featuring profit-maximizing agents as bidders. The first general observation was that agents have the option not to bid, so with perfect information, revenue never exceeds the item value. In the above section, we borrowed the key insight about population mis-estimation from [5], and connected them to some stylized models of the entry of agents into the auction.

One of the original hypotheses motivating this research was that a conventional game-theoretic or auction-based analysis would miss a key aspect of the Swoopo paradigm. Specifically, we conjectured that Swoopo behavior might be better explained by research on lotteries.

From a high-level, the strength of this intuition is that it matches extremely well the stories agents appear to tell about themselves, as well as the marketing campaign of Swoopo. A simple glance at Swoopo advertisements shows that Swoopo emphasizes the substantial savings earned by the winner of a recent auction. Rather than pitching Swoopo as a convenient online retailer using a non-traditional auction mechanism, the advertisers stress the large upside to bidders who are lucky enough to win.

The problem with this hypothesis, however, is that it attempts to answer an isolated problem (Swoopo) with a famous, much-researched paradox on lotteries. This complication, however, does not prevent us from considering some straightforward methods of connecting existing observations about lotteries to the Swoopo framework. Most importantly, the fact that lottery behavior requires adjustments to our assumptions motivates some alterations to our model.

5.1 Lottery Research

Conventional explanations for why individuals participate in (negative expected value) lotteries take a variety of forms. One common justification is that such lotteries are themselves a source of entertainment or utility. That is, to model the lottery as a utility maximization based solely on payoffs is to ignore an important aspect of the lottery itself: that individuals derive utility from participation itself. Many behavioral economists have offered explanations based on some measure of bounded rationality. For one, humans are very poor at conceiving of extremely small or large orders of magnitude. The win probabilities, at least of jackpot payoffs, might be so small that individuals cannot accurately compute their expected utility. Alternatively, the magnitude of payoffs might be treated asymmetrically, with ticket purchases based on entertainment budgeting (rather than cost), while the prize money gets treated as pure winnings. Some controversial theories even attribute the popularity of lotteries to credit-constrained agents attempting to attain a threshold level of wealth that enables entrepreneurial or other income-generating activity.

While the stories behind each of the above explanations vary in their plausibility, they hopefully address the complexity of the lottery problem. In the Swoopo context, treating the lottery as a game is particularly attractive, as many commentators (and indeed, Swoopo itself) views the website as *entertainment shopping*, a perverse euphemism for a casino-warehouse amalgam. The behavioral explanations of lotteries extend quite literally to the Swoopo environment, insofar as the assumed risk-neutral, utility maximizing agent in our previous models has been implicitly endowed with the ability to reason about small probabilities and moderately large payoffs.

The most obvious challenge is transforming these explanations into testable hypotheses. For now, we remain interested in explicating our model in the context of a few variations. In the subsequent empirical section, we will provide some connection between the data and these lottery-motivated changes.

5.2 Entertainment Shopping

Again, a variety of observers [8] have proposed entertainment as the hidden explanation accounting for the gap between predictions and reality in auction revenue. Essentially, we seek to tweak the model to provide this previously-unaccounted for utility from participation. In a trivial but realistic approach, we consider a dollar-value Δb representing the surplus to the bidder received simply for placing the bid.

The analysis does not change in any significant way. For simplicity we recall the fixed-price, symmetric model. The indifference probability (in the natural equilibrium) has $p = \frac{b-\Delta b}{v-w}$, since Δb can be literally interpreted as a rebate on each bid. So, in the distribution of outcomes, there is no substantial change; only the parameter of auction termination p changes, but the distribution remains geometric.

The impact of course is that revenue predictions rise in lockstep. Conditioning on the fact that at least one bid has already been placed, we have:

$$\mathbb{E}[Rev] = w + b\left(\frac{v-w}{b-\Delta b}\right) \tag{37}$$

We recognize that the simplicity of the change explains the relatively uninteresting interpretation: the model lowers the effective bid cost, so that agents pay the same price to play, yet employ mixed strategies that make termination less probable.

Despite its simplicity, we observe that this model does lend itself to some later empirical approaches. Namely, if we assume the entertainment premium remains constant over time, a change in the bid cost will have a nonlinear impact on revenue. Specifically, the revenue will decrease nonlinearly in b.

Note that this analysis can be extended in the obvious way to individualize the entertainment premia. Specifically, we can imagine the general, asymmetric model of Section 3.2 with b_i replaced by $b_i - \Delta b_i$. Although a parameter will have been added, the analysis remains intact.

5.3 Risk Attitudes

In principle, we would like to be able to propose and test all plausible explanations for the behavior exhibited by Swoopo bidders. However, in the compromise between generality and specificity, we consider the case of risk-loving agents, which can be viewed as a manifestation of many of the effects described related to lotteries.

Our first model of risk-lovingness used a power utility function. Based on [10] we reconsider our analysis for a constant absolute risk aversion (CARA) utility function with risk aversion parameter λ. Specifically, CARA takes the form $u(w) = -e^{-\lambda w}$, motivated by the fact that bidder wealth is unknown. The parameter λ refers to the risk aversion (or risk lovingness) of that agent. There is a tradeoff at work in the choice of utility function. While the power-law utility requires arbitrary assumptions about wealth, it does have the advantage of separating preferences for risk and preferences for skew, discussed in [2]. Regardless, for this analysis we follow the suggestion of [10] and use CARA utility. The model itself, however, follows our original analysis with only the utility function changing.

We again consider the standard fixed-price auction, with valuation v, bid cost b, and fixed price w. The only asymmetry in this model concerns there being two potential risk profiles in the bidding population. We consider k agents with CARA utility functions with parameter λ and the remaining $n-k$ agents with risk-neutral, dollar-maximizing utility. We denote the CARA subpopulation A, and the risk-neutral population B.

In terms of knowledge, we assume that each agent assumes all players have the same utility function that he does. Therefore, this model is structurally identical to the work in Section 4.1. Although we solved that model for risk-neutral bidders, we recall the observation that none of the analysis depended on the form of agent utility functions.

Each agent in A mixed according to an indifference condition that satisfies

$$\pi_A u_A(W+v-w-b)+(1-\pi_A)u_A(W-b)=u_A(W) \tag{38}$$

where π_A is the predicted termination probability by the A agent, u_A is the CARA utility function with parameter λ, and W is an arbitrary wealth level which will drop out.

We suppress the derivation in the interest of space, however we conclude that

$$\pi_A=\frac{1-e^{\lambda b}}{e^{\lambda(w+b-v)}-e^{\lambda b}} \tag{39}$$

From this, we can define the mixed-strategy profile of A agents, who will pass with probability, in stages $k>1$:

$$q_A=(\pi_A)^{\frac{1}{n-1}}=\left(\frac{1-e^{\lambda b}}{e^{\lambda(w+b-v)}-e^{\lambda b}}\right)^{\frac{1}{n-1}} \tag{40}$$

Recall that in the first stage there are n potential bidders, whereas in subsequent stages there are $n-1$, which would adjust the mixed strategy accordingly. We continue to ignore these cases for simplicity in presentation.

For agents in B, the analysis is exactly as it was in the original model. Specifically, $\pi_B=\frac{b}{v-w}$. This gives us the mixed strategy for stages $k>1$

$$q_B=(\pi_B)^{\frac{1}{n-1}}=\left(\frac{b}{v-w}\right)^{\frac{1}{n-1}} \tag{41}$$

Finally, with these mixed strategies determined, we can easily find the *actual* termination probabilities. We have separate termination probabilities that result depending on whether the leading agent is of type A or type B. This is the only place in which the size of each group's population actually appears. The termination probability with an A incumbent is $p_A = q_A^{k-1} q_B^{n-k}$, and with B, $p_B = q_A^k q_B^{n-k-1}$. We summarize this as

$$p = \begin{cases} \left(\frac{1-e^{\lambda b}}{e^{\lambda(w+b-v)} - e^{\lambda b}}\right)^{\frac{k-1}{n-1}} \left(\frac{b}{v-w}\right)^{\frac{n-k}{n-1}} & \text{for incumbent from } A \\ \left(\frac{1-e^{\lambda b}}{e^{\lambda(w+b-v)} - e^{\lambda b}}\right)^{\frac{k}{n-1}} \left(\frac{b}{v-w}\right)^{\frac{n-k-1}{n-1}} & \text{for incumbent from } B \end{cases} \tag{42}$$

Unsurprisingly, interpreting these probabilities is largely contingent on the value of the parameter λ. We are here primarily motivated by the risk-loving case $\lambda < 0$. In this case, we have $\pi_A < p_B < p_A < \pi_B$. The intuition is very clear: A agents predict the lowest termination probability; if a B agent is the leading bidder, termination is slightly higher since there are more A agents to bid; and the largest termination probability is what B agents predict.

5.4 Informational Asymmetry

We consider an alternative story behind the risk profiles motivating the bidding population. The space of bidders continues to be split into two types, CARA utility maximizers denoted A, and risk-neutral maximizers denoted B. In the above, we assumed every agent naively assigned his type to the entire population. Now, we imagine that the risk-neutral bidders continue to behave naively, assuming that their opponents are also all risk-neutral, so that Swoopo auctions are merely a non-traditional method of selling goods at fair value. The CARA utility maximizers, however, have perfect information. They correctly identify the auction participants as either risk-neutral or like themselves.

This story clearly matches the asymmetry discussed in general in Section 4.2. Again, those results were specifically derived for risk-neutral utility, but the analysis holds through if agents of each type have a distinct utility function. Only the parameter names have changed for B, the asymmetry-unaware subpopulation, so we still have that

$$q_B = (\pi_B)^{\frac{1}{n-1}} = \left(\frac{b}{v-w}\right)^{\frac{1}{n-1}} \tag{43}$$

$$p_B = \begin{cases} p_{B,A} = (q_B)^{n-k} = \left(\frac{b}{v-w}\right)^{\frac{n-k}{n-1}} & \text{when leader is in } A \\ p_{B,B} = (q_B)^{n-k-1} = \left(\frac{b}{v-w}\right)^{\frac{n-k-1}{n-1}} & \text{when leader is in } B \end{cases} \tag{44}$$

The analysis for A is structurally the same, but we now have an indifference condition with the

CARA utilities (instead of the expected payoff utilities), which gives assumed termination probability:

$$\pi_A = \frac{1 - e^{\lambda b}}{e^{\lambda(w+b-v)} - e^{\lambda b}} \tag{45}$$

We recall that the the identity $p = p_A p_B$ must always hold. Since A has perfect information, their assumed termination probability (from the indifference condition) is accurate, such that $p = \pi_A = p_A p_B$. These probabilities apply in both cases, when an agent from A is leading, or when an agent from B is leading. We conclude that

$$p_A = \begin{cases} p_{A,A} = \left(\frac{1-e^{\lambda b}}{e^{\lambda(w+b-v)}-e^{\lambda b}}\right)\left(\frac{v-w}{b}\right)^{\frac{n-k}{n-1}} & \text{if leading agent is in } A \\ p_{A,B} = \left(\frac{1-e^{\lambda b}}{e^{\lambda(w+b-v)}-e^{\lambda b}}\right)\left(\frac{v-w}{b}\right)^{\frac{n-k-1}{n-1}} & \text{if leading agent is in } B \end{cases} \tag{46}$$

The only case we need to check is that these are indeed valid probabilities. So, we must exclude the cases $p_{A,A} \geq p_{A,B} > 1$. We are interested in solving $p_{A,A} \leq 1$ for λ, which would provide us an exact condition on λ in terms of v, b, d, as well as k and n. This equation cannot be solved in general using algebraic techniques. So, we settle for a numerical approach. Motivated by the existence of risk-lovingness, we assume that $\lambda < 0$.

Indeed, for $\lambda < 0$, we have that $\pi_A < \pi_B < p_B$. Therefore, we conclude that

$$p_A = \frac{\pi_A}{p_B} < 1 \tag{47}$$

So while $\lambda < 0$ is not necessary, it is a sufficient condition for our work in Section 4.2 to apply. In this case, the termination probability is uniquely determined by the A agents, as

$$p = \frac{1 - e^{\lambda b}}{e^{\lambda(w+b-v)} - e^{\lambda b}} \tag{48}$$

In terms of the probability of termination, this case is indistinguishable from the case of n CARA utility maximizers. The difference, of course, lies in the mixed-strategy equilibrium itself. The intuition behind this auction format is that B agents continue to behave as they always have in this equilibrium type. However, the auctions last longer because there are risk-loving A agents, whose mixed-strategy induces a lower termination probability than B agents anticipate. We note that the termination probability satisfies

$$p = \frac{1 - e^{\lambda b}}{e^{\lambda(w+b-v)} - e^{\lambda b}} < \frac{b}{v - w} \tag{49}$$

implying that revenues for Swoopo exceed the risk-neutral case. Therefore, the previous intuition still holds. The risk-loving agents are left indifferent, as they bid until their indifference is met. The risk-neutral agents are playing a game without full information, and so are bidding more frequently than they would optimally. The result is a utility (or, expected value) loss for B agents, which directly translates to auctioneer profits.

5.5 Risk-Lovingness

As a general principle, we are skeptical of explanations that rely on arbitrary manipulations of agent utility functions. We consider the prior for risk-neutrality to be very strong, and therefore demand substantial evidence in favor of risk-loving behavior before it is accepted. We are particularly averse to the idea of over-fitting by constructing utility functions that exactly replicate observations, without providing appropriate justification.

Regardless, there is substantial literature to support the idea of agents behaving in a risk-loving manner, particularly in relation to lotteries and other forms of gambling. In the context of Swoopo, some of the best empirical justification of risk-seeking behavior is analyzed in [4] and [7]. Both papers investigate the role of *myopia* in inducing risk-seeking behavior. Particularly of interest are the results of [7], which demonstrate that individuals purchase a significantly higher rate of lottery tickets when making purchase decisions one at a time, as opposed to determining how many tickets to buy at a single time step. The sequential nature of the purchase decisions was isolated and shown to cause agents to engage in higher levels of risk-seeking. The authors also justify this in the context of the *peanuts effect*, which closely relates to the idea of *skew lovingness*, in that agents are particularly attracted to gambles in which the (likely) downside is small in magnitude, and the (unlikely) upside is large in magnitude.

This *myopic risk-seeking* matches closely our intuition for Swoopo participants. During an auction, each agent makes purchase decisions with the bids he owns, deciding at each time step whether to place a bid or not. His decisions during the auction are unlikely to take the form of contemplating the full allocation of his bid endowment. Instead, it seems plausible that bidding takes the sequential structure of small decisions made, perhaps with myopia.

Of course, agents must purchase their bids ahead of time, in fixed-size packages, which seems to contradict the idea of myopic decisions. However, this effect itself has alternative explanations. For one, Swoopo may be trying to minimize the number of transactions, to avoid agents from purchasing a single bid every time one is needed. In the time-crucial context of an auction, Swoopo wants its bidders well-provisioned with bids. Alternatively, research in [12] shows that agents make significantly riskier investment choices when they decrease the frequency with which they make investment decisions. That is, [12] motivates us to consider the purchase of bids as an investment decision for each agent, whereas the use of a bid is where (mentally) the cost is incurred.

6 Empirical Work

The theoretical work discussed to this point is primarily motivated by the empirical work described in this section. Many elements of our models are indeed motivated by informal observations as well as our intuition, which is hopefully communicated throughout. However, we have endeavored primarily to use the models as a starting point, offering a variety of high-level predictions about how we might expect auctions to progress in practice. This section then should provide more formal motivation, more rigorous than aforementioned intuitions. In places, we look to this data-driven work to confirm or contradict certain predictions, and hopefully to suggest certain explanations for the behavior of auction participants.

As before, many of our empirical results have been subsequently investigated in the recent round of working manuscripts, especially in [3]. Where appropriate, we emphasize the cases where others have employed similar methods. We consider this primarily to be a useful feature of our research context, as we can use these other analyses to challenge our conclusions, or perhaps consider them as more robust.

6.1 Methods

As much as possible, we try to present our data-based conclusions in the clearest manner, prioritizing simple, intuitive methods that offer clear, meaningful conclusions. We do, for completeness, parameterize certain models under different conditions, but our interest is primarily in estimating the general likelihood of that particular explanation by testing it on independent subsets of our data. We give preference to graphical and tabular presentation of summary data over complicated econometrics. This is justified by the high level of uncertainty surrounding penny auctions, as well as the myriad potential explanations that remain, facts that push us toward a broader scope of high-level observations.

The first important technique we employ throughout this section is a method of relating termination probabilities across different-valued auctions. This method is motivated by the intuition that an auction for an item of value v at stage k should have half the termination probability of an auction for an item of value $2v$ at stage k (assuming the same bid costs and increments b and d). We recall the general form of our termination probability,

$$q(k, v) = \frac{b}{v - (k-1)d} \tag{50}$$

which is parameterized by the item value v and the stage k. Another way of interpreting this across-auction method is the following observation: an auction with item value $2v$ has an indifference condition giving some termination probability $\bar{q}$ at a stage k. Then, an auction with item value v will have this same termination probability when

$$\frac{b}{\bar{q}} = 2v - (k-1)d = 2\left[v - \left(\frac{k-1}{2}\right)d\right] \tag{51}$$

The straightforward conclusion is indeed correct. In testing the natural equilibrium, where this indifference condition holds, we note that $q(k, v)v$ should be constant across auctions. Therefore, we choose 1 as our arbitrary unit of item value, and have termination probabilities

$$\bar{q}(k, v) = vq(k, v) = \frac{b}{1 - \left(\frac{k-1}{v}\right)d} \tag{52}$$

so that the stages of the auction are computed in units relative to value. In the same manner, we have, for fixed-price auctions,

$$\bar{q}(v) = \bar{q}(k, v) = vq(k, v) = \frac{b}{1 - \left(\frac{w}{v}\right)} \tag{53}$$

since the stage k is arbitrary. We note that [3] uses a nearly identical method.

Furthermore, we are challenged by the fact that bids at different junctures of an auction cannot be treated as equals. Specifically, in price-incrementing auctions, the indifference condition changes at each stage because the payoff in the win-state decreases. Conditioned on our testing the natural equilibrium, we would expect the profitability of a bid to remain constant at each stage. In testing this, and other questions, we offer the following notion of *expected return on a bid*

$$R(k; b, v, d) = \frac{\tilde{q}(k, v)(v - (k-1)d)}{b} \tag{54}$$

where $\tilde{q}(k, v)$ is an empirical estimate of the probability of termination for this auction at stage k.

In practice, we can use the same time-adjustment procedure described for termination probabilities on these bid-returns, allowing us to compare auctions of different-valued items:

$$R(k; b, v, d) = \left(\frac{\tilde{q}(k, v)v}{b}\right)\left(1 - \left(\frac{k-1}{v}\right)d\right) = \frac{\bar{q}(k, v)}{b}\left(1 - \left(\frac{k-1}{v}\right)d\right) \tag{55}$$

In the natural equilibrium, our hypothesis is implicitly that these bid-returns are uniformly 1 across the auction-space and stage-space, as a bid-return of 1 implies the empirical termination probabilities satisfy the indifference condition. Likewise, $R > 1$ suggests positive profitability of a bid, scaled by item value, and $R < 1$ suggests negative profitability of a bid.

6.2 Datasets

A key feature of Swoopo's design that has attracted researchers is the easily accessible data on completed auctions. In February of 2009, we conducted our first crawl of Swoopo's historical data, recording for every terminated auction every piece of available data. This data includes the item being auctioned, the price level at auction end (paid by the winner in price-incrementing auctions), the fixed price in fixed-price auctions, Swoopo's listed retail value of the item, the cost of a bid, the bid-increment (possibly 0), the time of auction end, and the type of auction (i.e., NailBiter, Beginner, Fixed-Price). We proceeded to repeat this data collection in May 2009, September 2009 and January 2010. We observe consistency in the Swoopo data, insofar as auctions crawled at different collection points have the same published records.

We also attempted to collect real-time bid data, as historical auctions offered little in the way of describing how bid histories developed (only the usernames corresponding to the last 10 bids are recorded). However, we encountered difficulty in this attempt, as it proved infeasible to monitor a large percentage of concurrent auctions. Fortunately, in January 2010, the authors of [5] made publicly available their dataset, which included the same historical data we had collected, but also with bid-level data for over 7000 auctions in October through December 2009. We have chosen to present our empirical work using their dataset, primarily for ease of formatting and consistency across the bids dataset and the auctions dataset.

More details about their dataset are available in [5], but a feature of particular note is that only 4328 of these auctions are *full* in the sense that every bid is captured. In order to be conservative, they elected to throw out the auctions with incomplete bid histories. While we agree with this conservatism, and practice the same judiciousness with the data, we also emphasize the pernicious selection bias introduced into the dataset. Namely, the auctions with full datasets appear to have been selected for by lower activity of BidButlers (which frequently induce chains of simultaneous bids that are impossible to capture fully). This alone is concerning, but the relationship between these BidButlers and revenue is much more troubling. By excluding auctions with incomplete bid histories, we are restricting ourselves to a very particularized subset of auctions which enables us to draw very inaccurate conclusions. Wherever possible, we attempt to avoid relying on bid-level data because of this selection bias. Where it is used, we discuss the potential ramifications of this bias.

6.3 Auction Values

One of the most immediate observations in a Swoopo auction concerns the retail price that Swoopo lists. Ostensibly serving to inform the bidding agent, the naive analysis assumes agents are perfectly anchored to this valuation. In practice, some work has been done in [3] and [10] to show that these valuations are frequently better estimated by a more thorough search. Both demonstrate that for items available on Amazon.com, the revenue of each auction is more highly correlated with the Amazon.com price than the Swoopo suggestion. We note that the price discrepancies are on the order of 30% for those items it affects. This is indeed an interesting result, as it lends a particular degree of rationality to Swoopo bidders that requires knowledge or search outside of the auction format itself.

In our revenue calculations, we treat the value of an item as equal to that which Swoopo publishes. Since these valuations are likely to be overstated, this method actually understates Swoopo revenues.

6.4 Revenues

As the feature that motivated this research in the first place, revenues are the single most striking aspect of the Swoopo empirics. As we have discussed at length in our theoretical sections, any risk-neutral (or risk-averse) agents participating in a perfect information game cannot provide expected profits for the auctioneer. This fact follows directly from the *optionality* available to every bidder: apprised of the appropriate information, each player may elect not to play.

Therefore, the fact that Swoopo consistently generates profits, across auction types over time, deeply challenges our theoretical models. In our dataset, Swoopo averages a profit of over $210 per auction, which corresponds to 70% profit margins on average. In Figure 1 we show that Swoopo has been highly profitable over this span of over a year. This presentation justifies the motivating intuition behind studying Swoopo: this unusual method of auctioning items, coupled with the appropriate set of bidders, has produced substantial auctioneer profits. This contradicts any predictions made by our equilibrium analysis in full information, provided we do not introduce risk-seeking behaviors, or other non-standard assumptions related to rationality.

One peculiar observation from Figure 1 concerns the dramatic decline in both measures of profitability from the heights of March and April 2009 to a sharp fall in May and June, continuing to a valley in October 2009, before a sharp assent. We observe that the explanation is not in the number of auctions being offered: specifically, total revenues decreased at the same or faster rate over this period, as the number of auctions being offered also declined.

One potential explanation concerns some sort of structural change in Swoopo's auction mechanism or characteristics. Although the summer of 2009 featured a wide variety of changes to Swoopo's format, we note that the revenue declines begin too early to be an effect of Swoopo's alterations. Indeed, these alterations are much more likely to be a strategic response to lost revenue. That is, responding to the sharp drop in profits, Swoopo retooled its strategy: lowering the bid cost, implementing a wide variety of bid-increments, and adding promotional features. Regardless, these changes cannot account for the early evidence of the decline.

The only other plausible explanation we can offer is a fundamental decline in demand for Swoopo's services. It is possible this is a lagging response to the depressed macroeconomic climate of late 2008. This explanation is purely speculative, but the effect observed is too large to be ignored.

We also present Figure 2, which breaks down the same discussion above into a variety of auction types. Most of the profit levels match our intuition or are discussed in the table as well, but this graph does indirectly illustrate the changing Swoopo environment in the summer of 2009. The standard auctions of 15 cent-increments disappeared, as did fixed-price auctions, in favor of a wide variety of price-increments and a 60 cent bid fee. We recall that, at the time of this writing, Swoopo offers exclusively penny auctions.

Although the completeness of the data in Figure 2 encourages our claims about Swoopo profits, to provide further evidence of this robustness we present Table 1. This table confirms that even conditioning on a variety of auction subsets, the qualitative results relating to profitability are highly consistent. Only one auction type, the 24 cent auction, which is unsurprisingly a short-lived offering, has negative profits. We observe that profitability for price-incrementing auctions is (nearly) monotonically decreasing in the increment. The lone exception is the 15 cent variety, which is fundamentally difficult to compare, as 15-cent auctions represent only the first half of the time period. This observation matches our original *lottery intuition*, namely that the lower-incremented auctions' lottery characteristics are earning them high revenues. More formally, which we investigate

Type	Number of Auctions	Profit ($)/Auction	Profit (%)/Auction
Beginner	19549	26.6499	10.0015
Fixed-Price	2203	662.271	70.0318
100%-Off	4556	243.761	119.198
NailBiter	32765	59.2896	50.6573
NailBiter & Beginner	6450	33.7421	27.2803
1 cent	13505	895.19	185.313
2 cent	3421	631.247	118.362
6 cent	9010	224.834	87.6631
12 cent	25233	56.9628	29.2765
15 cent	59243	101.691	59.6417
24 cent	4218	-9.57342	-15.6253
Total	121419	216.442	70.8618

Table 1: This table summarizes the profits for Swoopo in our auctions dataset. Profitability (as a dollar amount and a percentage of item value) are reported for various subsets of auction types.

later, there is a strong trend in this data suggesting that auctions with more skewed payoffs are more profitable. Free, fixed, penny and 2 cent auctions are the highest earners, and they necessarily have highly skewed payoff distributions. This provides a good justification for our intuitive expectation that the lottery model, by admitting *skew lovingness* in agents, could explain heightened revenues.

Another noticeable trend in the data is the low profitability of NailBiter and Beginner auctions. This is hardly surprising. In [5], aggressive, rapid bidding (often associated with a BidButler) is demonstrated to be responsible for a substantial portion of Swoopo's profits. Further, [3] demonstrates that there is a strong relationship between experience (as measured by bids used) and aggressive behavior. They also demonstrate that the profitability from aggressiveness has an associated learning effect. Regardless, the big picture confirms the same effect discussed in the selection bias of the dataset from [5]. Namely, we reason that Beginner and NailBiter auctions strongly select for lower levels of aggressive or rapid bidding patterns, which in turn correlate strongly with the lower revenue levels.

6.5 Termination Probabilities

Having established that Swoopo is earning consistent, outsized profits, we turn to an analysis of the distribution of outcomes. In principle, we are interested in capturing the long-run distribution of auction lengths, factoring in constant levels of bid fees, price-increments or fixed price levels, and item valuations. In practice, this is not possible, so we rely on approximate methods.

The parameter hardest to control in the dataset is of course item valuation. We still have the problem addressed in Section 6.3, as well as the general auction concern about individual agent

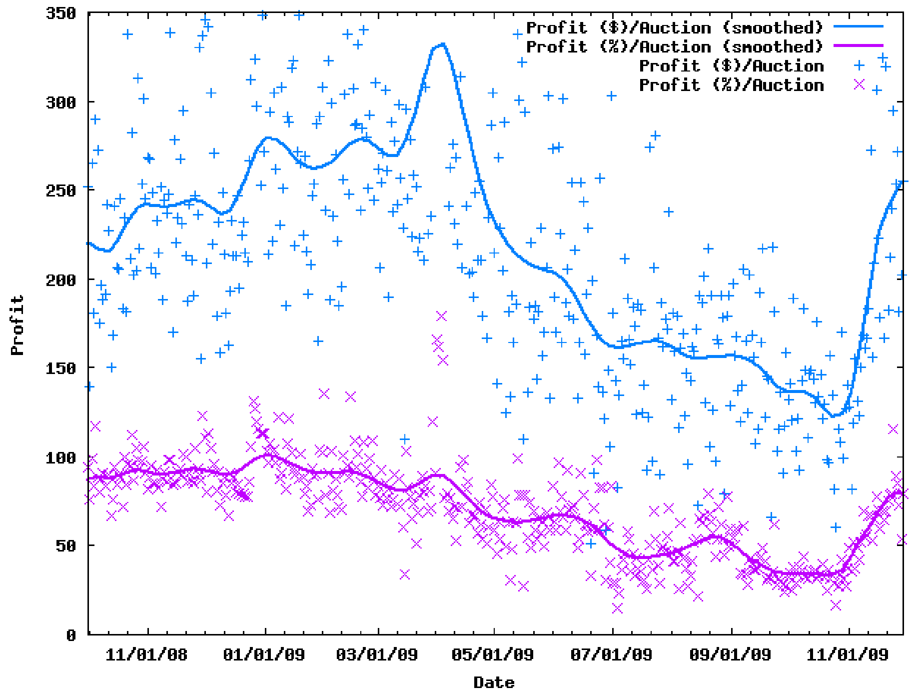

Figure 1: This scatterplot demonstrates the daily profitability of auctions over time. Note that the y-axis refers both to dollar amounts and percentage levels for simplicity of presentation. We have also smoothed the data to show large effects over time.

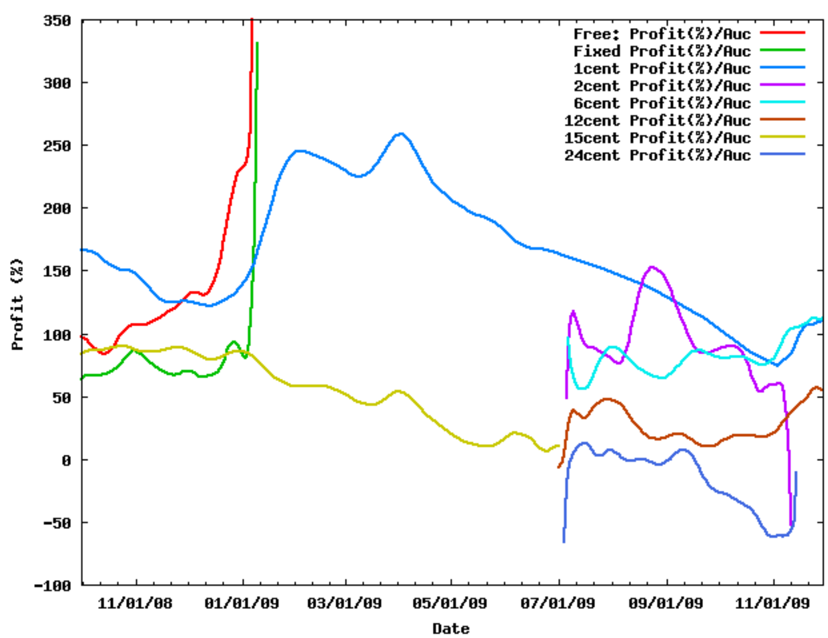

Figure 2: Plot of (smoothed) levels of daily profitability (measured in percentage) for a variety of auction subsets.

valuations. However, these are insignificant effects, compared to the general problem of auctions covering items of vastly different values. We recall our notion of normalizing the *termination probability* by an item's value, discussed in Section 6.1. Using this framework, if we fix the bid fee and the increment for each bid, we are able to leverage the adjusted times (normalized by value) and compare directly across auctions.

In Figure 3, we illustrate the normalized termination probability for a substantial subset of the historical auctions dataset. These three graphs apply to pre-July 2009 auctions, to avoid complications created by changes in the bid fee. Each graph represents a different type of auction, fixed-price, penny auctions, and the original 15 cent auction. We also include graphs of the theoretical termination probabilities. These lines are the predictions made in the *natural equilibrium*, and represent the aggregate probability of every agent passing, which we previously denoted p_k. We have normalized these theoretical graphs so that direct comparison can be made to the empirical data.

The most immediate observation is that the probability of termination, without exception, falls below the predicted level. Of course, this result is not at all surprising according to our revenue findings and our original intuition; it is only surprising given the prior that the natural equilibrium held. We also observe that the deviation is largest in the penny and fixed price auctions. This also should come as no surprise given our revenue analysis by panel in Table 1.

The most powerful effect captured in these graphs, however, is the time-progression of termination probability. At the earliest stages, the deviations are relatively small, and in the case of the 15-cent auctions are negligible. However, as the value-adjusted measure of time increases, the predicted level of termination exceeds the realized termination by an increasing amount. This is particularly striking in the penny auction graph, where the probability of termination steadily decreases, while the predicted level of course rises. This observation is consistent across all auction types, before and after the change in the cost of a bid (and other structural changes in the summer of 2009). To illustrate the robustness of this observation, we present Figure 8 in Appendix A that have the same interpretation for various auction formats.

As we observed in Section 6.1, we can just as easily interpret these termination probabilities in terms of the expected return on a bid at a given time step. Indeed, the definition of R, the return on a bid, confirms the obvious: at the same (adjusted) time-step, the ratio of expected termination probability to theoretical termination probability represents the expected return to a bidder at that time step. We confirm this by charting R against the value-adjusted time. We omit these graphs in the interest of space, primarily because they offer no additional insights. Relating termination probabilities to bid returns, however, is fundamental: the expected return to a placed bid decreases through an auction. This result is particularly striking in that it contradicts a previous finding in [10]. The authors claim that bid profitability increases over time, which they support by looking at auctions for Wii video game consoles. We strongly believe their predictions are based on an inappropriately small dataset, and argue for the robustness of our findings.

This observation is at the center of our empirical analysis. We consider it particularly valuable in that it defies some conventional intuitions that many observers consider obvious. Specifically, as strongly illustrated in the 15-cent figure, the probabilities do not coalesce well with stories about skew-seeking bidders. In auctions with larger price-increments, the effect of the rising price-level has a substantial effect on the attractiveness of the winning state. This effect is obviously captured in the convexity of the theoretical graphs for the penny and 15 cent auctions. The fact that bidders are decreasingly profitable as the auction progresses suggests greater competitiveness or over-bidding when the payouts are less skewed. Although this does not rule out risk-seeking behaviors as a potential explanation for some effects in the data, it makes clear that skew-loving behaviors alone are not driving Swoopo's profits.

This observation also challenges the unsophisticated intuition that an auction participant should bid late in auctions. Although this belief does not have the sound behavior ground that the skew-intuition does, it is surprisingly common, and the authors of [10] lend it some credence. These graphs are of course trimmed, but it seems quite clear that waiting alone is not a profitable strategy for the Swoopo bidder (though waiting forever might be optimal).

6.6 Bid Returns

With our uncertainty now largely focused on the reason for decreasing returns to bidding over time, we seek potential theoretical explanations for this effect. Though we are still using the *natural equilibrium* to make predictions, we inevitably consider relaxing the assumptions about agent behavior. Many potential explanations can be imagined, but we consider two particularly strong candidates.

6.6.1 Sunk Costs

In [3], Augenblick presents a discussion of the *naive sunk cost fallacy* as a potential explanation for agent behavior on Swoopo. In the context of bid-returns decreasing over time, the notion of sunk costs seems entirely plausible. To demonstrate this effect, we offer a brief overview of Augenblick's work. We note that his work on sunk costs is essentially an implementation of the sunk cost model in [6], which is similarly implemented in [1].

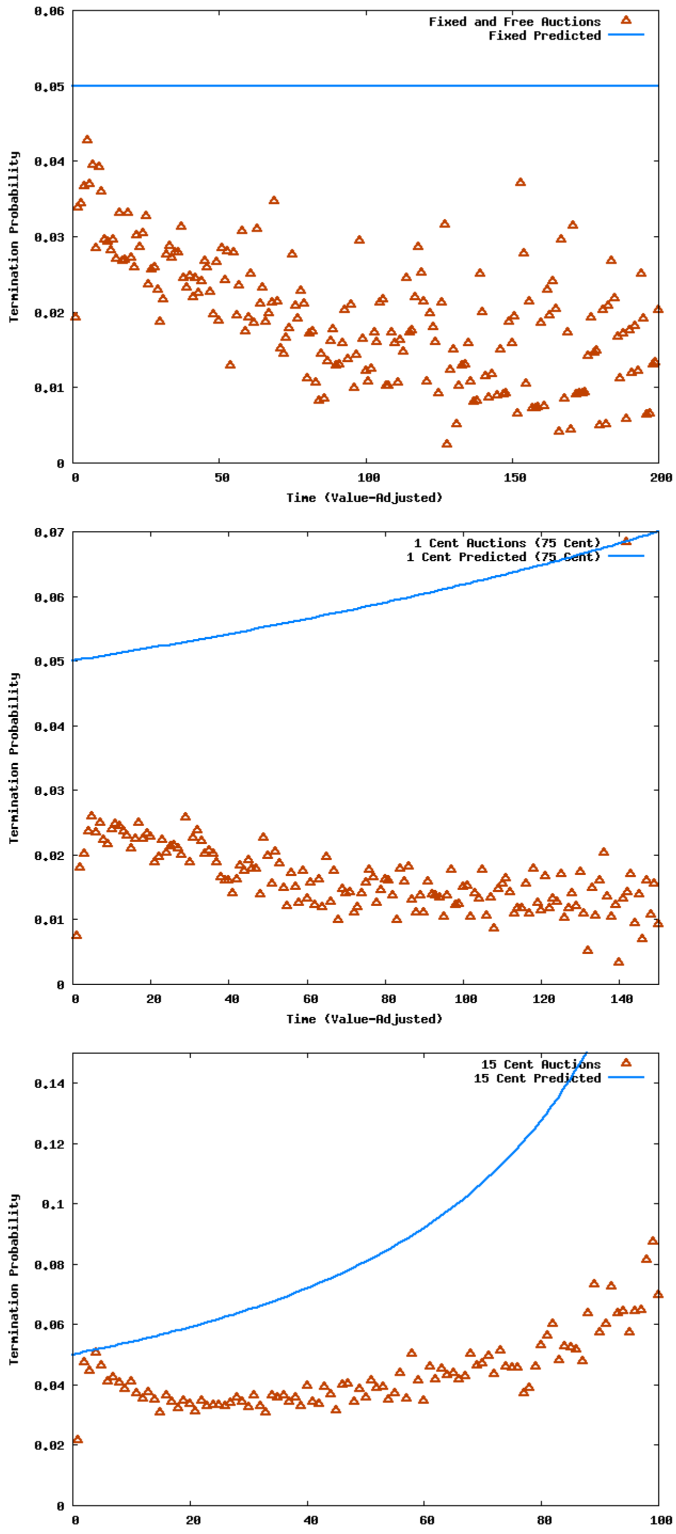

Figure 3: These figures demonstrate a comparison of equilibrium predictions and Swoopo data concerning the termination probability at a given time step. The graphs are normalized according to the procedure described in Section 6.1.

The driving element of the model is a regret parameter ρ representing a (utility) cost, on top of the dollar cost, to losing the auction after placing bids. Then, having placed m bids, the player experiences $-mb(1+\rho)$ if he never bids again. Clearly this cost alone would suggest a decreased incentive to bid. However, the model hinges on the additional parameter η, representing the naivety of the agent in his inability to forecast his own regret.

Therefore, he *believes* his utility from the scenario where he bids again and subsequently loses will be $-b(m+1)(1+(1-\eta)\rho)$, understating the downside. The last step to the model is defining utility in the state where this subsequent bid is a winner. Using this utility we can define indifference. The model's last twist requires that the naive bidder, after winning, must still pay his regret cost $b(m+1)\rho$ if and only if auction revenues exceed the item value v. When the item is worth more than the auction's revenue, this regret cost is foregone.

Solving this model is very straightforward analytically. Using the solution in [3], we can determine our *termination probability* as predicted by the naive sunk cost equilibrium:

$$p_k = \begin{cases} \max\left(0, \frac{b+b\rho-b\rho\eta(\frac{k}{2}+1)}{v-kd+(1-\eta)c\rho(\frac{k}{2}+1)}\right) & k \leq \frac{2(v-b)}{b+2d} \\ \max\left(0, \frac{b+b\rho-b\rho\eta(\frac{k}{2}+1)}{v-kd}\right) & k > \frac{2(v-b)}{b+2d} \end{cases} \tag{56}$$

We avoid analyzing this equilibrium further as it is not our result, but we emphasize the intuition. The termination probability has two different regimes, depending on the level of k relative to $\frac{2(v-b)}{b+2d}$. However, in both regimes, the termination probability p_k is decreasing for all k where $\rho, \eta > 0$. That is, assuming the bidders have this regret cost, and are naive in their projections, then the probability of the auction terminating falls over time. In short, the intuition of the sunk cost effect suffices: the bids of an agent that prove unsuccessful draw him in (with his regret cost), and he is able to keep making these troublesome bids because of his naivety.

We view this as a potentially acceptable explanation for the empirical result of declining termination probabilities, primarily because it is simple, intuitive, and offers clear predictions. However, we are reluctant to place any confidence in this explanation thus far. There are no bid-level empirics in [3] to support (or refute) this connection.

6.6.2 Auction Entry

Following the theoretical work we presented in Section 4.2, we recall the powerful impact that imperfect information and asymmetries can cause. In particular, we recall the special case considered by [5] in which players misestimate the population of bidders participating in the auction.

Motivated by the stories discussed in Section 4.6, we consider the claim that termination probabilities decline over time because our entry model in Section 4.5 provides accurate predictions. Again, this model suggests that agents, deterministically or probabilistically, enter the model with some fixed assumptions about the population size of their fellow bidders. As agents enter over time, these estimates do not adjust (either at all or adequately) to reflect the new population. Also, we demonstrated that the convexity of revenue with respect to these mistakes allowed for large changes in revenue.

To test these predictions, we used our bid-level data and constructed subsets in which at least m bids were placed (where m took on the values 30, 50, 75 and 100). Then, for each of these defined subsets, we considered the number of unique bidders that appear in the first k bids (where $k \leq m$ took on values ranging from 30 to 75). We then plot, for each ordered pair (k, m), revenue for each observed auction as a function of the number of unique bidders recorded in the first k bids. Again, we note that we have controlled for the portion of the auctions that has already passed by throwing out auctions with fewer than m bids.

Our results demonstrate that there is a strong positive relationship between the number of unique bidders in the beginning of an auction and total revenue. The effect ranges from 2 to 15 expected dollars (in profit) per unique bidder. This range is substantial, but the qualitative prediction is robust across a large set of parameters. We offer a representative case ($k = m = 50$) in Figure 4, which demonstrates the effect of unique bidders in the early stages of an auction on final profit.

That graph actually includes bid-histories that were *incomplete*, in the sense that one or bids were missing. We henceforth refer to these as *incomplete bid histories*. As discussed in Section 6.2, the choice between the complete and incomplete bid histories presents a tradeoff between two forms of selection bias. To address this, we were careful only to include bid histories for which the first k bids were observed. Otherwise, we would run the risk of understating the number of unique bidders. We believe this adequately neutralizes the effect of missing bids, as the player identities are irrelevant for bid numbers exceeding k, and we could verify the revenue in the historical auctions dataset.

Also, to be thorough, we have included Figure 5, which repeats the analysis for the complete bid histories. The effect still exists, though is smaller in magnitude. We conclude the completeness of the bid histories does not qualitatively alter our interpretation.

Furthermore, the selection effects (that make particular bid histories incomplete) are ambiguous with respect to the quantity of unique bidders. While having many unique bidders increases the likelihood of simultaneous bidding (which might cause skipped observations), this effect seems small compared to the effect of selecting for BidButlers. That is, incomplete histories are far more likely to have been selected as such because BidButlers induced a chain of simultaneous bids. Although we have not controlled directly for this BidButler effect, our data indicates BidButlers have a strong positive effect on revenue, and a negative effect on the number of unique bidders. Therefore, if the data-supported intuitions about the selection effects are correct, we would expect the effect to be even stronger than we have estimated.

6.7 Early Bidders

One observation left out of our analysis of Figures 1 and 2 relates to the left tail of each graph, in the neighborhood of the adjusted time value of 0. These early time steps exhibit behavior visibly different from the rest of the distribution. Though the effect only applies to the first handful of data points, its consistency in all 3 auction types, as well as the data in the appendix draws our concern. We take special care of the strange behavior at the very first time step, where the probability of termination is much lower than the surrounding observations.

We are immediately drawn to consider explanations related to potential selection effects, since auctions with no bids are not observable. However, there is no clear story to motivate why this would impact the decision facing the second (or later) stage bidder. Moreover, although the effect quickly dissipates, the data shows that this increase in termination probability exists for adjusted time steps on the interval 0 to 5.

Motivated by this observation, we consider an alternative analysis. First, we confirm the existence of this observation in a variety of contexts. Among them, as our general technique suggests, we test for robustness across all major categories of auction (each bid increment, beginner auctions, NailBiter auctions, etc). We continue to observe that there is a definite interval effect, which we summarize as follows: consistently, across all subsets, the first bid has substantially lower profitability than the subsequent bids. The rest of the interval effect involves the rising profitability and probability of winning, which consistently occurs on the range of the next 3 bids.

With this qualitative observation in hand, we proceed to posit the existence of a subpopulation of bidders, who are, for whatever reason, particularly attracted to bidding at the beginning of auctions. We can easily imagine certain players self-identifying in this manner, closely watching upcoming auctions and preparing to be one of the earliest bidders.

We briefly test for the existence of such a population. By using the bid histories, we consider all histories for which the first 15 bids are recorded. Then, for each of the first 15 bid steps, we determine

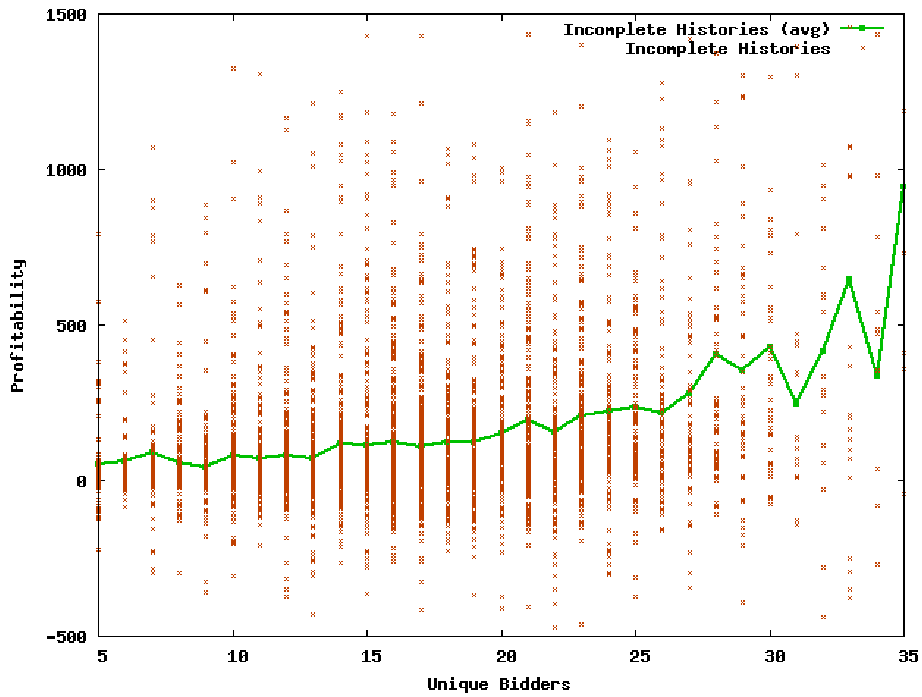

Figure 4: Scatterplot showing each observation of auctions with 50 or more bids, comparing number of unique bidders to realized revenue. The average case is also plotted. This scatterplot includes the incomplete bid histories, which are complete in the sense that all the first 50 bids are recorded.

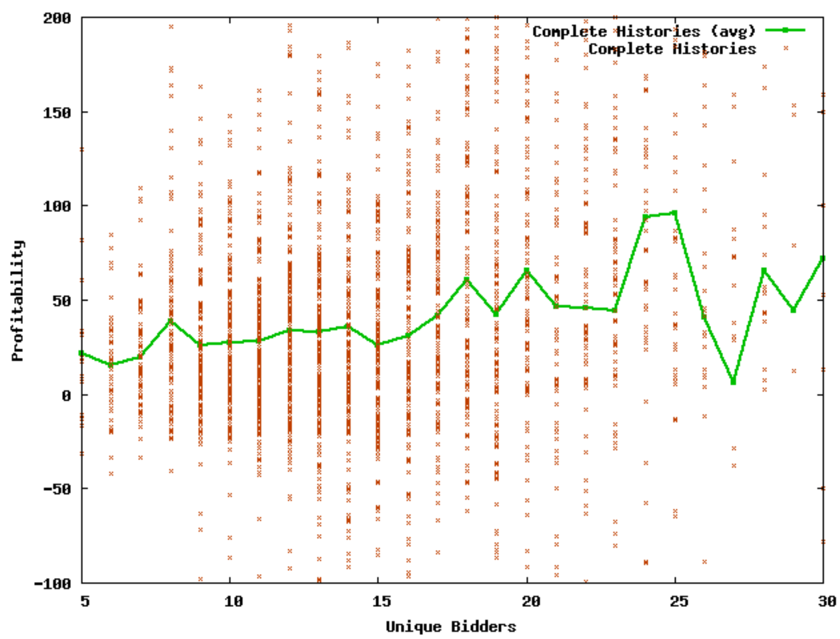

Figure 5: As above, except this graph includes only observations for which *every* bid were observed, including those after 50.

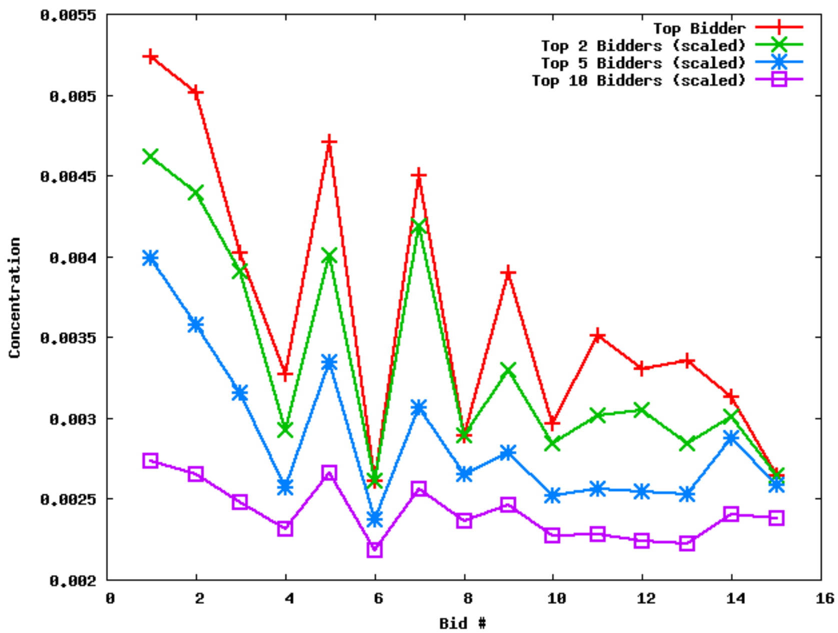

Figure 6: This figure demonstrates the relationship between the k^{th} bid and the level of concentration amongst the i most frequent bidders at that level. The concentration effect is robustly decreasing across the first 15 bids, most visibly in the first 6 to 8.

the distribution of users who bid in any auction at that time step. For example, for the first bid, we determine how many times each user in the dataset bid first in an auction. Repeating this for all 15 time-steps, we then consider the top i bidders for each time-step, where i takes on all integer values from 1 to 20. In short, this procedure gives us a clustering measure of how concentrated the bidder population is for each time-step.

We find that there is a very obvious effect related to the clustering: the concentration of bids by the top i bidders (robust for different values of i) declines rapidly over the first 15 bids. Although this does not answer questions about why agents are forming this early bidding subpopulation, our hypothesis of its existence is strongly confirmed. In Figure 6, we demonstrate a subset of our analysis, where i takes on the values 1, 2, 5 and 10. The sawtooth patterns at alternating time-steps are the result of the data being on too small a scale to perform any smoothing, for fear of losing the demonstrated effect. The key observation is the rapid descent of how concentrated the bidder population is at each bid time-step.

So, we have profiled a particular type of bidder, which we call *early bidders*. We hesitate to declare definitely what causes these early clusters, but we do note that it could be a rational response to the revenue observations from earlier. These early bidders could be a manifestation of small parts of the population *learning* that profitability declines over time. We leave it to further work to investigate if there is any correlation between the early bidding population members and experience (number of bids placed), or other evidence of a learned effect.

Finally, we recall the initial motivation for discovering these early bidders was the extremely strange probability of termination for the first bid. We of course believe this is related broadly to the dynamics of this *early bidder* population, but we have not demonstrated how. One story, which seems somewhat appealing is that users mentally segregate auctions which have no bids placed from those which do. Under this unsophisticated approach, auctions that have a single bid already placed could attract bids with higher probability than those with no bids placed. We cannot test this hypothesis adequately because there is no data on auctions terminating without a winner.

6.8 Timer

To this point, we have abstracted away the aspect of the Swoopo timer. In practice, it is an extremely complicated feature of the auction, though in principle its purpose is obvious. The timer is necessary as a termination condition, but in Swoopo's case the resetting feature is absolutely necessary to prevent a completely random lottery at the last possible second to enter a bid. The rules for how exactly Swoopo adds time to the clock are not spelled out, though we appear to have captured the major heuristics from data and live auction monitoring. The basic structure is that each bid adds seconds to the timer, and in the vast majority of cases, this involves resetting the timer to 20, 15 or 10 seconds, depending on how many bids have already been placed. The 20 second case generally precedes the 15 second case, which always precedes the 10 second case.

Because conditioning strategies on the timer is unattractive from a game-theoretic and modeling perspective, there is no analysis of this crucial aspect of the auction format. Indeed, the generally very thorough treatment of [3] suggests that the timer is an uninteresting feature of Swoopo. Although we agree the timer is not a worthwhile addition to our formal penny auction model, we feel strongly that the timer plays a crucial role in what we might call a *heuristic model* of how agents strategize.

One key observation about the timer is that players exhibit the same style of sniping strategy we see on eBay. Nearly twice as many bids are placed at 1 or 2 seconds as at 3 or 4 seconds. Second, since the timer resets after a bid, and players have varying computational and mental abilities to respond, there are many more bids placed at 20, 15 and 10 seconds than at any other bid level. The column of observations in Table 7 illustrates this effect, which is well captured in that there are 55,000 bids at 20 seconds, but only 19,000 at 19 seconds, and only 10,000 at 17 seconds.

According to these two heuristic facts, we would expect very low profits for bids placed in the last 1, 2 or 3 seconds (contingent on where players set their snipe heuristic), as well as very low profits for the 10, 15 and 20 second levels. The intuition is a bit subtle, but very important: bids that are marked at 10, 15 or 20 seconds are (in almost every case) intended to be bids at a lower timer-level. That is the player whose bid is recorded at 10 seconds, almost certainly attempted to bid at $t < 10$, but another user's bid was recorded in between his decision and his action. The result is that his bid is treated after the other bid, and so is recorded as a 10 second bid. Then, the selection effect for 10, 15 and 20 second labeled bids is very poor: this particular auction, at this particular juncture has at least two bidders (and probably more) who are actively bidding at the snipe threshold. This will inevitably produce very poor results for the bids that are recorded at these levels. We note that these poor results are only compounded when BidButlers are active. Since BidButlers can trigger one another, their bids are recorded simultaneously, and all but one of them (and usually all of them) are losing bids.

To test these heuristics and their impact on winning probability, and more importantly profitability, we constructed a sequence of timer data from the bid histories. The dataset includes a record of what the time read after that particular bid was placed, as well as the number of seconds Swoopo added to the timer as a result of that bid. We had to be careful with the dataset because of irregularities associated with *groups of bidders* (which occur when multiple bids were still recordable independently, but Swoopo publishes them in one stream of data). To be thorough, we use the timestamp associated with each bid to provide independent verification of the timer claims of the dataset. Wherever these values differ by more than one second (the limit of precision given the granularity of our data), we throw out that observation. Fortunately, we repeated this analysis without cleaning the dataset and produced nearly identical results. We conclude that the selection effect of throwing out discrepant data does not appear to be pernicious.

Then, we use these data to compute the time left when each bid was placed, and indexing on these, we consider the profitability of bids placed with any of 1 to 20 seconds left. Figure 7 summarizes our results. As predicted, the most profitable bids occur in the lulls that follow the 10, 15 and 20 second flurries (again, caused by simultaneous bidding). We are not surprised at these effects, but we are impressed with how strong and visible the effect is in the data. For robustness the data is presented graphically for the complete bid histories, as well as the incomplete bid histories separately. We exclude the complete bid histories from the incomplete set in this context to make them independent.

Time remaining	Per-Bid Profit	Observations (number)
1	-0.144109	21126
2	-0.118111	21305
3	0.0727426	15434
4	0.248126	12347
5	0.2693	11298
6	0.456764	11912
7	0.437005	14231
8	0.162748	18802
9	0.175941	22146
10	-0.0871238	34281
11	0.132115	8772
12	0.0702723	11753
13	0.190701	17042
14	-0.0466994	22481
15	-0.40606	54390
16	0.167896	7045
17	0.30579	10179
18	0.20606	14964
19	-0.0144929	18910
20	-0.436752	54928

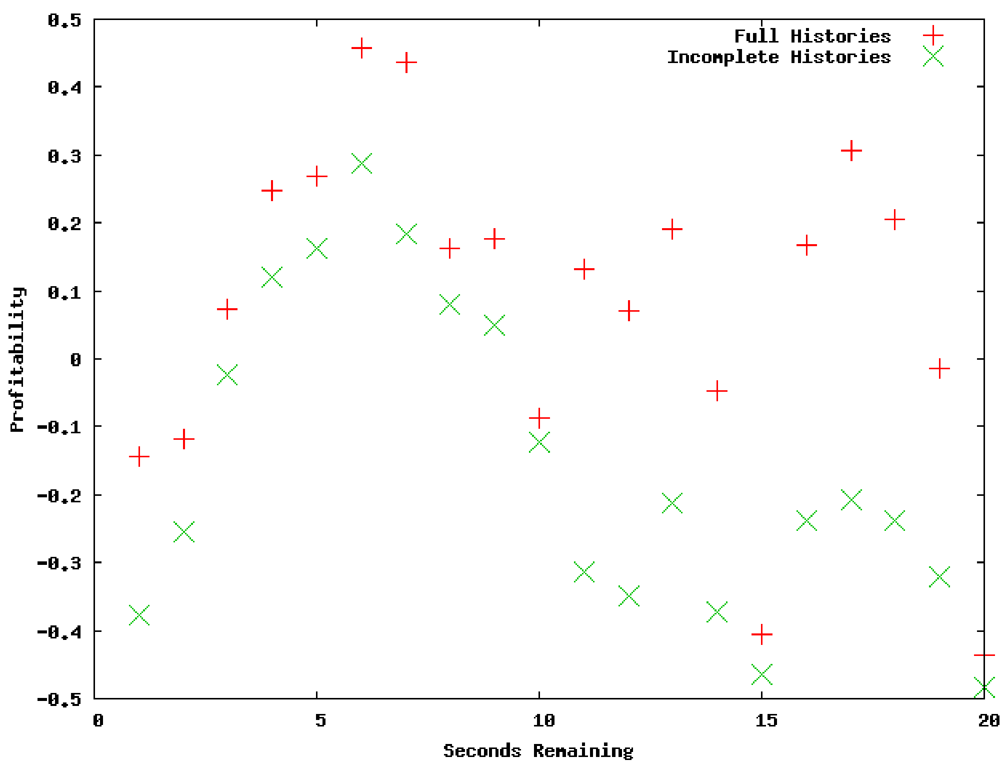

Figure 7: Table and graph demonstrating the timer effect. The graph includes series for the complete bid histories, as well as the complement (i.e., those bid histories which are not complete). The table corresponds to the data in the complete bid histories graph.

6.9 Risk Attitudes

We refrain from providing detailed analysis of fitting theoretical (risk-loving) distributions of termination probabilities to the data. However, we note that in principle this method is relatively simple. The time-adjustment in Section 6.1 no longer applies to the theoretical terminations because the utility function has changed. Regardless, we can limit our scope to items for which there are a large number of observations, so that v is constant. Particularly useful are fixed-price and free auctions, because the theoretical probabilities are constant for any utility function for these types of auctions.

We do offer a rough estimate of λ, simply to establish the order of magnitude suggested by the data. For fixed-price and free auctions, we derive the empirical probability of a bid winning for each of the most commonly auctioned items. We use this to provide a baseline estimate of λ. This method amounts to solve $pU(v-b) + (1-p)U(-b) = u(0)$ which amounts to

$$p\left(\frac{1-e^{-\lambda(v-b)}}{\lambda}\right) + (1-p)\left(\frac{1-e^{\lambda b}}{\lambda}\right) = 0 \tag{57}$$

where we have performed a transformation to our CARA utility so that we avoid the trivial solution $\lambda = 0$. Solving this condition for the 20 most commonly auctioned (fixed-price or free auctions) items, gives us an estimate of $\bar{\lambda} = -.0147$ as the sample mean. The solved values of λ took on values in the range $[-.0046, -.0489]$. We do not consider these estimations to provide any support for the risk-loving hypothesis.

Broadly, it seems apparent that there are lottery-like motivations that drive the behavior of agents in terms of their entry into auctions and their bidding decisions. Moreover, Figure 2 demonstrates what is both qualitatively and quantitatively difficult to ignore: the percentage profitability of an auction is extremely well-predicted by the level of bid-increment. This perfectly matches the heuristic we would anticipate a skew-seeking agent to employ: namely, identifying and pursuing auctions that have the most skewed payoff distributions (i.e., those with the lowest bid increments).

Still, we recall the central observation in Section 6.5. The risk or skew seeking effect, if existent, must be acting in concert with another effect, which is very strong in the latter stages of auctions (so that bids are increasingly unprofitable even when payoffs are less skewed). This naturally suggests a modal interpretation: agents may select themselves into auctions by using a skew-heuristic associated with risk-lovingness, but then continue to make increasingly unprofitable bids for some other reason.

7 Discussion

We conclude this paper much as we started it, with a long list of empirical questions in need of answers. The Swoopo problem is very likely to remain relevant for quite some time. There are still important empirical questions that can be easily answered, so researchers are likely to maintain interest. And there are no overwhelming signs of agent behavior changing, so bidders and Swoopo will continue to provide the fascinating motivation. A great deal of the satisfaction in posing questions and conjecturing answers about Swoopo data is related to the fact that so many observations are downright perverse.

One extremely crucial observation, which builds on our theoretical results in Section 4 and empirics in Sections 6.5 and 6.6.2, is that Swoopo revenues appear to be *extremely sensitive to population size*. This can be interpreted as sensitivity to perceived population, actual population, and certainly both. We are particularly happy with the intuition proposed in Section 4.6 and consider the formulation of these models a logical next step for our research.

It should be apparent that high quality bid-level data is a prerequisite for doing the sort of reliable data analysis we have in mind. In an ideal world, we would devise a method to determine how many users who are logged into Swoopo are currently viewing a given page. Although this alone is not the best definition of the population size, it establishes a good proxy. In the future, we would consider using measures of active bidders (number of bidders in the last k minutes) and try to derive a reliable estimator for the population parameter n.

As a separate justification, if we return to the simplest scenario and consider the mixed-strategy employed in equilibrium, we have symmetric agents bidding with probability

$$1 - \left(\frac{b}{v - kd}\right)^{\frac{1}{n-1}} \tag{58}$$

Although we took great pains to demonstrate reasons why b, v, or d might vary across agents, we recognize that these values can only be so different. However, the parameter n can vary immensely, and it is difficult to identify stories for how agents would keep their estimations of n close to its true value. Given that we are observing probabilities deviating systematically from prediction, we are inclined to believe strongly that the explanation is due in large part to n.

A topic left nearly untouched in our research is a fuller description of bidder profiles. For illustrative purposes, and largely to address a question raised by the revenues data, we uncovered the existence of an *early bidder* population that routinely bids in a concentrated manner in the very earliest time-steps only. We would be curious to see if a conscientious observer could conjecture a wider variety of agent profiles, and provide convincing work to categorize individual users. We hypothesize that the timer is very likely to be a useful tool of separating categories of users, and of course the BidButler should produce a bimodal distribution of those who use it regularly, and those who do not at all, with very little in between.

An area of research to become increasingly fruitful will involve cross-site comparisons, as Swoopo has developed a large following of *swoopoclones*, websites that imitate the basic functionality of Swoopo. Although [3] shows these clones tend to experience modest daily losses, we are curious as to the entry problem for *swoopoclones*. We recognize that explicit barriers to entry are essentially negligible, but it appears quite clear that the implicit barrier to entry is some minimum threshold of users needed to finance profitable auctions. We would also be curious to see if the daily losses experienced by *swoopoclones* show evidence of serving as loss leaders. That is, does it seem these websites are purposely oversupplying at certain times so as to cultivate a user base? In general, we look to these clones as further evidence of our general observation that the size of the bidder population is crucially important.

Also, we ignore the issues of collusion, shill bidding, and other concerns about dishonest behavior, primarily because we do not view these hypotheses as testable at this juncture. However, competitor firms, *swoopoclones* seem like the ideal place to test hypotheses about what collusion or shill bidding would look like.

Finally, we are interested to see in what ways FreeBids are used differently from regular, paid bids. It seems like a great environment to test theories about mental accounting or endowment effects. Once purchased, a paid bid and a FreeBid are indistinguishable, so evidence of markedly different profitability or usage of one type would suggest a behavioral peculiarity.

In general, we are pleased to observe the diversity of research appearing about Swoopo in recent months. We hope that this trend continues, as we consider it an extremely relevant and interesting dataset. To that end, we believe the most important next step is that a concerted effort be given to more thorough data collection. No existing research has utilized what we consider an adequate dataset, namely with a large number of bid-level observations, without omitted bids. This step, more than any other, will enable a wide variety of well-motivated research on the Swoopo problem.

References

[1] N. Ashraf, J. Berry, and J.M. Shapiro. *Can higher prices stimulate product use? Evidence from a field experiment in Zambia*. National Bureau of Economic research Cambridge, Mass., USA, 2007.

[2] T. Astebro, J. Mata, and L. Santos-Pinto. Preference for Skew in Gambling, Lotteries and Entrepreneurship.

[3] N. Augenblick. Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis. *Unpublished manuscript. Available at www.stanford.edu/ñed789*, 2009.

[4] S. Benartzi and R.H. Thaler. Risk aversion or myopia? Choices in repeated gambles and retirement investments. *Management Science*, pages 364–381, 1999.

[5] J.W. Byers, M. Mitzenmacher, and G. Zervas. Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank. *Arxiv preprint arXiv:1001.0592*, 2010.

[6] E. Eyster. Rationalizing the past: A taste for consistency. *Nuffield College, Oxford*, 2002.

[7] E. Haisley, R. Mostafa, and G. Loewenstein. Myopic risk-seeking: The impact of narrow decision bracketing on lottery play. *Journal of Risk and Uncertainty*, 37(1):57–75, 2008.

[8] M. Haynes. Programming Tech Blog. *http://matthaynes.net/blog/2008/10/05/money-to-burn-visit-swoopo/*, 2008.

[9] T. Hinnosaar. Penny Auctions. *Unpublished manuscript at http://toomas.hinnosaar.net*, 2009.

[10] B. Platt, J. Price, and H. Tappen. Pay-to-bid Auctions. *NBER Working Paper*, 2010.

[11] M. Shubik. The dollar auction game: A paradox in noncooperative behavior and escalation. *Journal of Conflict Resolution*, 15(1):109, 1971.

[12] R.H. Thaler, A. Tversky, D. Kahneman, and A. Schwartz. The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test*. *The Quarterly Journal of Economics*, 112(2):647–661, 1997.

[13] Richard H. Thaler. Paying a Price for the Thrill of the Hunt. *New York Times*.

A Auxiliary Figures

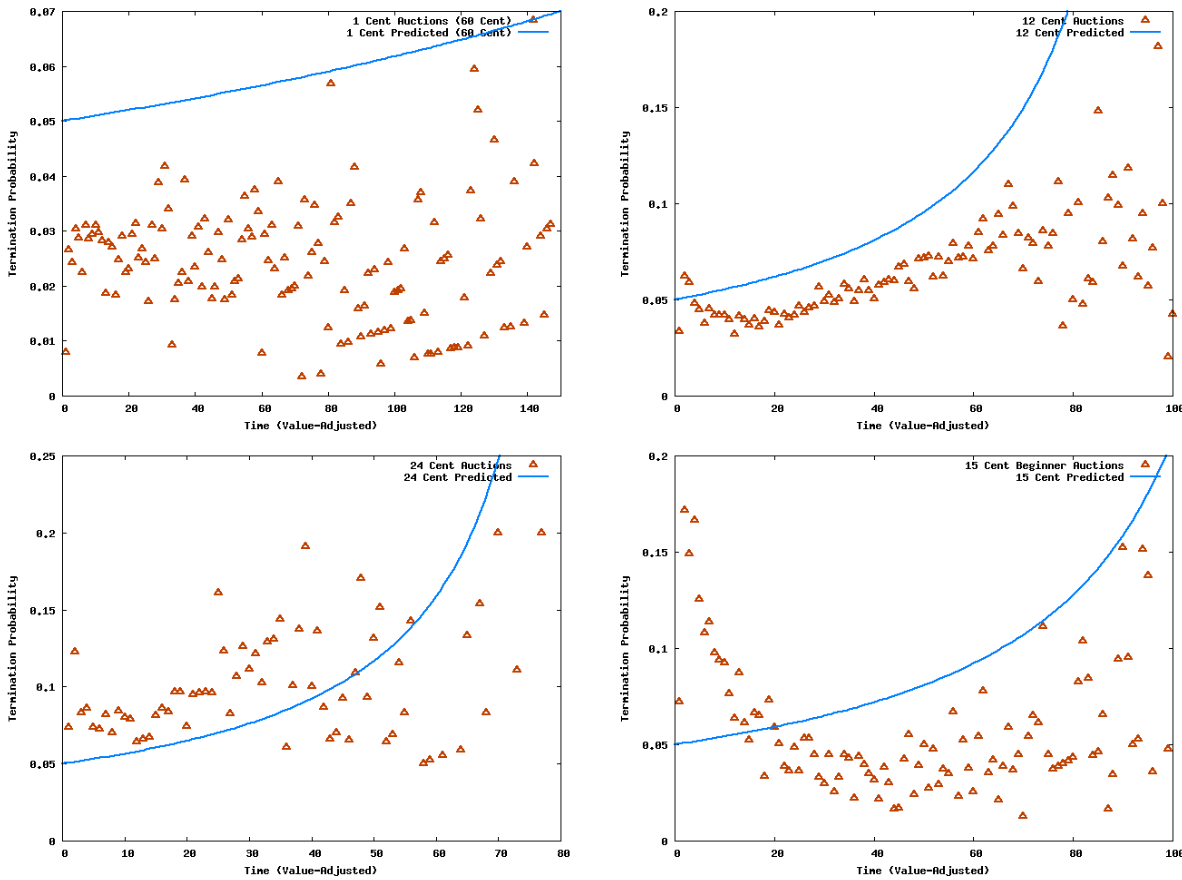

Figure 8: These figures demonstrate a comparison of equilibrium predictions and Swoopo data concerning the termination probability at a given time step. The graphs are normalized according to the procedure described in Section 6.1. The first two serve as independent confirmation of data in Figure 3. The second pair confirm some earlier observations. The 24 cent auction demonstrates negative revenue, while the Beginner Auction has a strong early incentive.