

Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis

Ned Augenblick*

December 28, 2009

JOB MARKET PAPER

Abstract

I theoretically and empirically analyze consumer and producer behavior in a relatively new auction format, in which each bid costs a small amount and must be a small increment above the current high bid. I describe the set of equilibrium hazard functions over winning bids and identify a unique function with desirable conditions. Then, I examine bidder behavior using two datasets of 166,000 auctions and 13 million individual bids, captured with a real-time collection algorithm from a company called Swoopo. I find that players overbid significantly in aggregate, yielding average revenues of 150% of the good's value and generating profits of €18 million over four years. While the empirical hazard rate is close to the predicted hazard rate at the beginning of the auction, it deviates as the auction progresses, matching the predictions of my model when agents exhibit a sunk cost fallacy. I show that players' expected losses are mitigated by experience. Finally, I estimate both the current and optimal supply rules for Swoopo using high frequency data, demonstrating that the company achieves 98.6% of potential profit. The analysis suggests that over-supplying auctions in order to attract a larger userbase is costly in the short run, creating a large structural barrier to entrants. I support this conclusion using auction-level data from five competitors, which establishes that entrants collect relatively small or negative daily profits.

Keywords: Internet Auctions, Market Design, Sunk Costs

JEL Classification Numbers: D44, L83

*Address: 579 Serra Mall, Stanford, CA 94305, ned.augenblick@stanford.edu. The author is grateful to Doug Bernheim, Jon Levin, and Muriel Niederle for advice and suggestions, and Oren Ahoobim, Aaron Bodoh-Creed, Tomas Buser, Jesse Cunha, Jakub Kastl, Fuhito Kojima. Carlos Lever, Scott Nicholson, Monika Piazzesi and Annika Todd for helpful comments. Orhan Albay provided invaluable research assistance. This research was funded by the George P. Shultz fund at the Stanford Institute for Economic Policy Research as well as the B.F. Haley and E.S. Shaw Fellowship.

1 Introduction

This paper theoretically and empirically explores consumer and producer behavior in the relatively new market for *penny auctions*, using two datasets collected from the largest auctioneer, Swoopo.¹ This auction provides an ideal field laboratory to study individual behavior as it is a reasonably transparent game (the rules are discussed below) played by many people (there are over 20,000 unique participants per week in my dataset) for relatively high stakes (the value of auctioned goods averages over €200) over a long period of time (the auctions have been run since 2005). Additionally, as the auctioneer must make important supply choices and many companies have recently begun to run the auctions, the analysis yields insights into producer optimization and market dynamics in a nascent auction market.

My analysis leads to three main contributions to the auction, market design, and behavioral economics literature. My first main result is that the auctioneer collects average revenues that exceed 150% of the value of the auctioned good, providing an unequivocal example of consistent overbidding in auctions. In order to better understand this behavior, I model the penny auction in a game-theoretic framework and solve for the equilibrium hazard rates of winning bids. My second main result is that, relative to equilibrium predictions, bidders overbid more and more as the auction continues. I establish that this behavior is consistent with the predictions of my model when agents exhibit a naive sunk cost fallacy. I show that higher levels of experience in the auction mitigate these losses. For my third main result, I use high-frequency data to separately estimate both the actual and optimal supply rules for Swoopo, demonstrating that the company captures nearly 99% of potential short term profits. The analysis suggests that over-supplying auctions for a given number of users leads the auctions to end prematurely, which can produce negative profits for the auctioneer. This effect implies a structural barrier to entrants, who must over-supply auctions in order to attract a larger userbase. This conclusion is consistent with findings from data on five major competitors, which establishes that entrants are not making significant daily profits in the medium term.

Before detailing my main results, it is useful to briefly describe the rules of the penny auction. First, players are restricted to bid a fixed *bid increment* above the current bid for the object, which is zero at the beginning of the auction. For example, if the current bid is €10.00 and the bid increment is €0.01, then the next bidder must bid €10.01, the next €10.02, and so on. The auction is commonly referred to as a "penny auction" as a result of the common use of a one penny bid increment. Second, players must pay a non-refundable fixed *bid cost* (€0.50 in my dataset) to place a bid. The majority of the auctioneer's final revenue is derived from the bid costs collected throughout the auction. Finally, the end of the auction is determined by a countdown timer, which increases with every bid (by approximately ten seconds in my dataset). Therefore, a player wins the object when her bid is not followed by another bid in a short period of time. Of traditionally studied auctions, the penny auction is most closely related to the War-of-Attrition (WOA).²

¹As of June 2009, Swoopo appeared to be the largest company running penny auctions by all important measures, such as revenue, number of daily auctions, and number of daily users.

²In both the WOA and penny auctions, players must pay a cost for the game to continue and a player wins when other players decide not to pay this cost. However, there are two main differences. First, in the

Using a theoretical analysis, I find that this auction format induces a mixed-strategy equilibrium in which the winning bid amount is stochastic. Moreover, by characterizing all equilibrium hazard rates, I show that any equilibrium in which play continues past the second period must be characterized by a unique global hazard function from that point onward. I demonstrate that this prediction is robust to a variety of changes and extensions to the game. I then test these predictions using auction-level data on 166,000 unique auctions and bid-level data on 13 million bids. The bid-level data was captured from Swoopo's website using a multi-server real-time collection algorithm that recorded nearly every bid from every person on every auction over a hundred day period.

My first main empirical result is that the average revenue from these auctions significantly exceeds the easily-obtainable value of the goods. Although Swoopo makes negative profit on more than half of the auctions in the dataset, the average profit margin is 52%, which has generated €18M in profits over a four year period. In an illustrative example, my dataset contains over 2,000 auctions for *direct cash payments*, in which the average profit margin is 104% of the face value of the prize. This finding contributes an unambiguous field example of overbidding in auctions to a large literature on the subject, which includes experimental work³ and a number of recent field studies.⁴ Perhaps the most convincing of the field studies is Lee and Malmendier (2008), who show that bidders pay 2% above an easily accessible "buy it now" in a second-price eBay auction on average, a much smaller effect than exhibited in this paper.

I then investigate players' strategies more closely by comparing the empirical and theoretically predicted survival and hazard functions. My second main finding is that the empirical hazard rate of these auctions starts at the rate predicted by equilibrium analysis, but deviates further and further below this rate as the auction continues. This decline implies that the expected return from bidding (and paying €50) drops as the auction continues, from nearly €50 in the beginning of the auction to only €12-€16 at later stages in the auction. Using a modified version of my theoretical model, I show that this behavior is consistent with that of agents who exhibit a naive sunk cost fallacy: as agents continue to play the game, they spend more money on bids, leading them to experience a higher psychological cost from leaving the auction (the modification follows Eyster (2002)). This finding provides suggestive empirical evidence of the existence and effects of sunk costs, which has been demonstrated in experiments (Arkes and Blumer (1985), Dick and Lord (1998)), but has been more difficult to observe empirically.⁵

penny auction, only one player potentially pays the bid cost in each period, unlike in a WOA, in which all players who continue to play must pay a cost in each period. Second, unlike in a WOA, the winning player must pay an additional cost of the winning bid, leading to a drop in the net value of the object as the game progresses.

³Overbidding has been documented in first-price common value auctions (see Kagel and Levin (2002)), second-price auctions (Kagel and Levin (1993), Heyman et al (2004), Cooper and Fang (2006)), and all-pay-like auctions (Millner and Pratt (1989), Murnighan(2002), Gneezy and Smorodinsky (2007)).

⁴These studies have suggested overbidding in a variety of contexts, ranging from real estate auctions (Ashenfelter and Genesove, 1992) to the British spectrum auctions (Klemperer, 2002). This literature often struggles with difficulties of proving the true value of the auctioned items, leading more recent studies to focus on online auction markets (see, for example, Ariely and Simonson (2003)).

⁵As noted in Eyster (2002), "Empirically testing for the sunk-cost fallacy is complicated by the fact that

The individual bid-level data, which contains approximately 96% of all bids made on the site for a three month period (13 million bids for over 129,000 users) allows me to describe the behavior of the market well. Participation in the auction is highly skewed: although half of the users stop playing after fewer than 20 bids, the top 10% of bidders in my dataset place an average of nearly 800 bids. Profits are gained largely from the most active users: the top 25% of users generate 75% of the profit. Interestingly, experience is associated with a significantly higher expected return from bidding. Users with little or no experience receive an expected return of only €18 on each €50 bid, while those with a large amount of experience (over 5000 bids) receive slightly over €50. By controlling for user fixed effects and allowing learning rates to vary across users, I show that this effect is partially due to learning, as opposed to selection bias. Finally, I determine the specific bidding strategies that drive the relationship between experience and profit by controlling for each strategy in the regression. This analysis suggests that two-thirds of the gain in expected profits associated with additional experience is due to the increased use of "aggressive bidding" strategies, in which the player bids immediately and continuously following other players' bids.

For my third main finding, I examine the supply side of the market. Using user and auction data at each point in time, I separately estimate Swoopo's *actual* and *optimal* supply rule: the number of auctions supplied for a given number of users on the site. For coordination purposes, Swoopo releases auction with very high initial timers, which gives them little ability to adjust to real-time changes in the number of users or auctions on the site. This process allows me to match the number of *expected* auctions at a point in the future with the number of *expected* users on the site at that time, which is Swoopo's *actual* supply curve. Then, I use natural deviations in the number of actual active users and active auctions from this expectation to estimate the *optimal* supply rule. These curves are extremely similar; Swoopo's supply curve obtains 98.6% of estimated potential profits given the empirical distribution of the number of users on the site. Both curves show that there are significant diminishing returns to the supply of auctions for a given number of users, because auctions are more likely to end prematurely when there are fewer users bidding on each auction. This finding suggests that there are high short-term costs for an entrant attempting to over-supply auctions in order to develop a userbase, creating a barrier to entry even though the initial startup cost of an Internet auction site is very cheap. Auction-level data that I collected from the top five competitors supports this conclusion: only one of Swoopo's five main competitors is making large daily profits, which still amount to only 6.6% of Swoopo's daily profits.

information-based explanations for behavior are often difficult to rule out." Eyster suggests that the most valid empirical paper on sunk costs Camerer and Weber (1999), which shows that NBA players are given more playing time than predicted if their team used a higher draft pick to acquire them. I partially surmount these difficulties by focusing on global, rather than individual, hazard rates.

There are only a small number of empirical⁶ and experimental⁷ papers on the WOA, as it is difficult to observe a real-life game in which the setup is the same as a WOA and there is a *known* bid cost and good value. Therefore, even though the games are different, this insights from penny auctions are potentially useful in understanding the way that people act in a WOA (and the closely related All-Pay Auction), which has been used to model a variety of important economic interactions.⁸ The results also contribute to the broader understanding of behavioral industrial organization, which studies behavior biases in the marketplace (see Dellavigna (2008) for a survey). Additionally, the paper relates to a set of recent papers that study other "innovative" auction formats, such as the lowest and highest unique price auction (Eichberger and Vinogradov (2008), Houba et al (2008), Östling et al. (2007), Rapoport et al. (2007)).⁹ Finally, the paper complements two other recent working papers on penny auctions. Using a subset of the Swoopo's American auction-level data, Platt et al. (2009) demonstrate that allowing flexibility in both risk-loving parameters and the perceived value of each good can explain the majority of bidding behavior. Hinnosaar (2009) studies the theoretical effect of imposing complementary assumptions to those in this paper, producing bounds on equilibrium behavior.

The paper is organized as follows: The second section presents the theoretic model of the auction and solves for the equilibrium hazard rates. The third section discusses the data and provides summary statistics. The fourth section compares the empirical survival and hazard rates with the equilibrium rates found in the second section. The fifth section outlines a theoretical model of sunk costs in this auction. The sixth section analyzes the effect of experience on performance of the bidders and discusses potentially profitable strategies. The seventh section focuses on the market for these auctions by providing an analysis of Swoopo's supply curve and its five top competitors. Finally, the eighth section concludes.

⁶Examples include Card and Olsen (1995) and Kennan and Wilson (1989), which only test basic stylized facts or comparative statics of the game. Hendricks and Porter (1996)'s paper on the delay of exploratory drilling in a public-goods environment (exploration provides important information to other players) is an exception, comparing the empirical shape of the hazard rate function of exploration to the predictions of a WOA-like model.

⁷See Horisch and Kirchkamp (2008), who find systematic *underbidding* in controlled experimental wars of attrition.

⁸For example: competition between animals (Bishop, Canning, and Smith, 1978), competition between firms (Fudenberg and Tirole, 1986), public good games (Bliss and Nalebuff, 1984), and political stabilizations (Alesina and Drazen, 1991). Papers with important theoretical variations of the WOA include Riley (1980), Bulow and Klemperer (1999), and Krishna and Morgan (1997).

⁹The paper also contributes to a growing literature that uses games and auctions to study behavior. There have been a large number of papers to study online auctions, which are surveyed by Bajari and Hortacsu (2004). More recently, Hartley, Lanot, and Walker (2005) and Post et al (2006) study such popular TV shows as "Who Wants to be a Millionaire?" and "Deal or No Deal?"

2 Auction Description and Theoretical Analysis

2.1 Auction Description

In this section, I briefly discuss the rules of the auction. There are many companies that run these auctions in the real-world and, while there are some small differences in the format, the rules are relatively consistent.

In a penny auction, there are multiple players bidding for one item. The bidding for the item starts at zero and rises with each bid. If a player chooses to bid, the bid must be equal to the current high bid plus the *bidding increment*, a small fixed monetary amount known to all players. For example, if the bidding increment is €0.01 and the high bid is currently €1.50, the next bid in the auction must be €1.51. Players must pay a small non-refundable *bid cost* every time that they make a bid. The auction ends when a commonly-observable timer runs out of time. However, each bid automatically increases the timer by a small amount, allowing the auction to continue as long as players continue to place bids. When the auction ends, the player who placed the highest bid (which is also the final bid) receives the object and pays the final bid amount for the item to the auctioneer.

The most similar commonly-studied auction format to the penny auction is the discrete-time war-of-attrition (WOA).¹⁰ In both auctions, players must pay a cost for the game to continue and a player wins the auction when other players decide not to pay this cost. However, there are two main differences. First, in a WOA, each player must pay the bid cost at each bidding stage in order to continue in the game. In the penny auction, only one player pays the bid cost in each bidding stage, allowing players to use more complex strategies because they can continue in the game without bidding. This difference also causes the game to continue longer on average in equilibrium, as agents spend collectively less in each period. Second, if the bidding increment is strictly positive, the *net value* of the good falls as bidding continues (and players' bids rise) in the penny auction, whereas the value of the good in the WOA stays constant.¹¹ This addition destroys the stationarity of the WOA model, as the agent's benefit from winning the auction changes throughout the auction.

The following section presents a theoretical model of the penny auction and provides an equilibrium analysis. In order to make the model concise and analytically tractable, I will make simplifying assumptions, which I will note as I proceed.

2.2 Setup

There are $n+1$ players, indexed by $i \in \{0, 1, \dots, n\}$: a non-participating auctioneer (player 0) and n bidders. There is a single item for auction. Bidders have a common value v for the

¹⁰In a WOA, each active player chooses to bid or not bid at each point in time. All player that bid must pay a bid cost. All players that do not bid must exit the game. The last player in the game wins the auction.

¹¹If the bidding increment is €0.00 (as it is in 10% of the consumer auctions in my dataset), the price of the object stays at €0.00 throughout the bidding.

item.¹² There is a set of potentially unbounded periods, indexed by $t \in \{0, 1, 2, 3, \dots\}$.¹³ The current high bid for the good starts at 0 and weakly rises by the *bidding increment* $k \in \mathbb{R}^+$ in each period, so that the high bid for the good at time t is tk (note that the high bid and time are deterministically linked).¹⁴ Each period is characterized by a publicly-observable current leader $l_t \in \{0, 1, 2, 3, \dots, n\}$, with $l_0 = 0$. To simplify the discussion, I often refer to the *net value* of the good in period t as $v - tk$.

In each period t , bidders simultaneously choose to bid or not bid. If any of the bidders bid, one of these bids is randomly *accepted*. In this case, the corresponding bidder becomes the leader for the next period and pays a non-refundable cost c . If none of the players bids, the game ends and the current leader receives the object and pays the final bid (tk). At the end of the game, the auctioneer's payoff consists of the final bid (tk) along with the total collected bid costs (tc).

I assume that players are risk neutral and do not discount future consumption. I assume that $c < v - k$, so that there is the potential for bidding in equilibrium. I assume that $\text{mod}(v - c, k) = 0$, for reasons that will become apparent (the alternative is discussed in the appendix). To match the empirical game, I assume that the current leader of the auction cannot place a bid.¹⁵ I refer to auctions with $k > 0$ as a (k) *declining-value auctions* and auctions with $k = 0$ as *constant-value auctions*.

The majority of the analysis focuses on the discrete hazard rate at each period, $\tilde{h}(t)$, which is the expected probability that the game ends at period t given that the game reaches period t . I use subgame perfection as my solution concept unless otherwise noted. For exposition purposes, I sometimes discuss a players' symmetric Markov strategy to avoid players conditioning their strategies on the past actions of other players in situations of indifference, although all of the results hold for non-symmetric and non-Markov strategies. Player i 's Markov strategy set consists of a bidding probability for every period given that they are a non-leader $\{x_0^i, x_1^i, x_2^i, \dots, x_t^i, \dots\}$ with $x_t^i \in [0, 1]$. Note that, in the case of Markov strategies, $\tilde{h}(t) = \prod_{i=1}^n (1 - x_t^i)$.

It is important to briefly discuss two subtle simplifications of the model. First, unlike in the real-life implementations of this auction, there is no "timer" that counts down to the end of each bidding round in this model. The addition of timer complicates the model without producing any substantial insights; any equilibrium in a model with a timer can be converted into an equilibrium without a timer that has the same expected outcomes and payoffs for each player. Second, when two agents make simultaneous bids, only one of the

¹²I assume that the item is worth $\underline{v} < v$ to the auctioneer. The case in which bidders have independent private values $v_i \sim G$ for the item is discussed in the appendix. As might be expected, as the distribution of private values approaches the degenerate case of one common value, the empirical predictions converge to that of the common values case.

¹³It is important to note that t does not represent a countdown timer or clock time. Rather, it represents a "bidding stage," which advances when any player makes a bid.

¹⁴The model takes place in discrete time so that each price point is discrete, allowing players to bid or not bid at each individual price. This matches the setup of the real-life implementation of the game.

¹⁵This assumption has no effect on the preferred bidding equilibrium below, as the leader will not bid in equilibrium even when given the option. However, the assumption does dramatically simplify the exact form of other potential equilibria, as I discuss in the appendix.

bids is counted. In current real-life implementations of this auction, both bids would be counted in (essentially) random order. Modeling this extension is difficult, especially with a large number of players, as it allows the time period to potentially "jump" from each period to many different future periods.¹⁶ However, the predictions of the models are qualitatively similar (numerical analysis suggests that the hazard rate of the equilibrium of the extended game is much more locally unstable, but globally extremely similar). These considerations are discussed in depth in the Appendix.

2.3 Equilibrium Analysis

In this section, I describe the equilibrium hazard rates for the penny auction, including a set of *bidding hazard rates*, which I use to make empirical predictions. The proofs for each proposition are in the Appendix.

I begin with the relatively obvious fact that no player will bid in equilibrium once the cost of a bid is greater than the net value of the good in the following period, leading the game to end with certainty if this period is reached.

Proposition 1 Define $F = \frac{v-c}{k} - 1$ if $k > 0$.

If $k > 0$, then in any equilibrium $\tilde{h}(t) = 1$ for all $t > F$

I refer to the set of periods that satisfy this condition as the *final stage* of the game. Note that there is no final stage of a constant-value auction, as the net value of the object does not fall and therefore this condition is never satisfied. However, as there is a final stage for all declining-value auctions, strategies in these periods are set, allowing us to use backward induction to determine potential equilibria. To begin, I identify the *bidding hazard rates* of the game:

Proposition 2 There is an equilibrium in which $\tilde{h}(t) = \left\{ \begin{array}{ll} 0 & t = 0 \\ \frac{c}{v-tk} & \text{for } 0 < t \leq F \\ 1 & \text{for } t > F \end{array} \right\}$.

In this equilibrium, some players use strictly mixed strategies for all periods after time 0 and up to (and including) period F . In these periods, players are indifferent between bidding and not bidding as their chance of winning the item in the following period (given that they bid) is $\frac{c}{v-(t+1)k}$, which is the cost of the bid divided by the benefit from winning the auction (the net value in the following period). Crucially, in a declining-value auction players in period F are indifferent given that players in period $F + 1$ bid with zero probability, which they must do by Proposition 1. This indifference allows players in period F to bid such that

¹⁶Hinnosaar (2009) models this extension, although with slightly different assumptions on the game and the possible parameters. He is able to provide bounds on some types of equilibrium behavior, which are consistent with the exact numerical results in the Appendix.

the hazard rate is $\frac{c}{v-Fk}$.¹⁷ Note that in a declining-value auction the hazard rate rises over time, with no bids placed after time period F . To understand this result intuitively, note that in order to keep players indifferent between bidding and not bidding, the probability of winning (which is equal to the hazard rate) must rise as the net value of the good is declining. In a constant-value auction, the net value and hazard rate stay constant throughout time as in the standard equilibrium of a WOA model.

Not surprisingly, there are continuum of other equilibria in this model that lead to hazard rates that are described completely in the following proposition:

Proposition 3 Consider any $p^* \in [0, 1]$ and $t^* \in \mathbb{N}$ with $t^* \leq F$. Let $\psi = \left\{ \begin{array}{ll} 0 & \text{if } p^* \leq \frac{c}{v-tk} \\ 1 & \text{if } p^* > \frac{c}{v-tk} \end{array} \right\}$.

Then, there is an equilibrium in which

$$\tilde{h}(t) = \left\{ \begin{array}{ll} 0 & \text{for } t < t^* \text{ and } \text{mod}(t + \psi + \text{mod}(t^*, 2), 2) = 0 \\ 1 & \text{for } t < t^* \text{ and } \text{mod}(t + \psi + \text{mod}(t^*, 2), 2) = 1 \\ p^* & \text{for } t = t^* \\ \frac{c}{v-tk} & \text{for } t^* < t \leq F \\ 1 & \text{for } t > F \end{array} \right\}$$

Furthermore, $\tilde{h}(t)$ in all equilibrium can be characterized by this form.

This proposition demonstrates that all equilibrium hazard rates can be described by two parameters, $p^* \in [0, 1]$ and $t^* \in \mathbb{N}$. To understand the intuition behind this result, recall that in any equilibrium that leads to the hazard rates in Proposition 2, players are indifferent at every period $t < F$ because of the strategies of players in the following period. Consider a situation in which players follow strategies that lead to the hazard rates in Proposition 2 for periods after t^* . If this is the case, players in period t^* are indifferent to bidding. Imagine that these players bid such that the hazard rate at t^* is some $p^* > \frac{c}{v-t^*k}$ (the rate for period t in Proposition 2). Given this strategy, players in period $t^* - 1$ will strictly prefer to bid in that period (as they now have a higher chance of winning in the following period), leading to a hazard rate of 0. As a result, players in period $t^* - 2$ will strictly prefer to not bid (as they have no chance of winning in the following period), leading to a hazard rate of 1. Consequently, players in period $t^* - 3$ will strictly prefer to bid, and so on, leading to alternating periods of bidding and not bidding (the argument is similar if players bid such that $p^* < \frac{c}{v-t^*k}$). All potential hazard rates are formed in this manner.

Note that these backward alternating periods of bidding and not bidding must occur in any equilibrium in which the hazard rate at some period t deviates from $\frac{c}{v-tk}$, which must lead to either $\tilde{h}(0) = 1$ or $\tilde{h}(1) = 1$, causing the game to end in either period 0 or period 1.

¹⁷For declining-value auctions, this requires the use of the assumption that $\text{mod}(v - c, y) = 0$. If this is not true, the players in the period directly before the final stage are not indifferent and must bid with certainty. As a result, the players in previous period must bid with zero probability (they have no chance of winning the object with a bid), causing players in the period previous to that to bid with certainty, and so on. This leads to a unique equilibrium in which the game never continues past period 1. However, as is discussed in the Appendix, there is an ε -equilibrium for an extremely small ε in which the hazard rates match those in Proposition 2.

Corollary 1, which is one of my main theoretical results, formalizes this intuition and notes that *if we ever observe bidding after period 1 in equilibrium, we must observe the hazard rates in Proposition 2* for all periods following period 1. Corollary 2 notes that Proposition 2 lists the unique set of equilibrium hazard rates for which the auctioneer’s expected revenue is v , which might be considered a natural outcome of a common value auction with many players.

Corollary 1 *In any equilibrium in which the game continues past period 1, $\tilde{h}(t)$ must match those in Proposition 1 for $t > 1$.*

Corollary 2 *For any $\alpha \in [\frac{c}{v}, 1]$, there exists a symmetric equilibrium in which the auctioneer’s expected payoff is αv . The equilibrium in Proposition 2 is the unique equilibrium of the game in which the auctioneer’s expected revenue is equal to v .*

In the following empirical sections of this paper, I restrict attention to the bidding equilibrium hazard rates in Proposition 2 for the reasons outlined in Corollaries 1 and 2. In addition to satisfying desirable theoretical characteristics, these hazard rates form the basis for clear and testable empirical predictions, which are unaffected by parameters such as the number of players in the game.

2.4 The Bidding Equilibrium - Equilibrium Predictions

As previously noted, I model the game in discrete time in order to capture important qualitative characteristics that cannot be modeled in continuous time (such as the ability to bid and not bid in each period). However, in order to make smooth empirical predictions about the hazard rates, I will now focus on the limiting equilibrium strategies when the size of the time periods shrinks to zero. While this does not significantly change any of the qualitative features of the discrete hazard rates, it creates smoother predictions, which allows the survival and hazard rates to be compared across auctions with different parameters.

Specifically, let Δt denote a small length of a time and modify the model by characterizing time as $t \in \{0, \Delta t, 2\Delta t, 3\Delta t \dots\}$ and changing the cost of placing a bid to $c\Delta t$. With this change in mind, define the non-negative random variable T as the time that an auction ends. I define the survival function $S(t; y, v, c)$, hazard function $h(t; y, v, c)$, and cumulative hazard function $H(t; y, v, c)$ for auctions with parameters y, v, c in the normal fashion (as $\Delta t \rightarrow 0$ and suppressing dependence on y, v, c):

$$S(t) = \lim_{\Delta t \rightarrow 0} \Pr(T > t) \tag{1}$$

$$h(t) = \lim_{\Delta t \rightarrow 0} \frac{S(t) - S(t + \Delta t)}{\Delta t \cdot S(t)} \tag{2}$$

$$H(t) = \int_0^t h(\tilde{t}) d\tilde{t} \tag{3}$$

Solving for these functions leads to the following proposition:

Proposition 4 *In the preferred equilibrium, $h(t) = \frac{c}{v-tk}$ for $t < \frac{v}{k}$.*

When $k = 0$, $H(t) = \frac{c}{v}t$ and $S(t) = e^{-(\frac{c}{v})t}$.

When $k > 0$, $H(t) = \frac{c(\ln(v) - \ln(v-tk))}{k}t$ and $S(t) = (1 - \frac{tk}{v})^{\frac{c}{y}}$ for $t < \frac{v}{k}$.

For the situation in which $k > 0$, note that each function approaches the corresponding function when $k = 0$ as $k \rightarrow 0$, which is reassuring:

$$\lim_{k \rightarrow 0} \frac{c(\ln(v) - \ln(v-tk))}{y} t = \frac{ct}{v} \quad (4)$$

$$\lim_{k \rightarrow 0} (1 - \frac{tk}{v})^{\frac{c}{y}} = e^{-(\frac{c}{v})t} \quad (5)$$

While Proposition 4 is useful to determine the hazard and survival rates for a specific auction, it is more useful to compare hazard and survival rates across auctions for goods with different values. To that end, define $\hat{t} = \frac{t}{v}$ as the *normalized time period*, define random variable \hat{T} as the (normalized) time that an auction ends, define the *normalized Survival and Hazard rates* in a similar way to above:

$$\hat{S}(\hat{t}) = \lim_{\Delta\hat{t} \rightarrow 0} \Pr(\hat{T} > \hat{t}) \quad (6)$$

$$\hat{h}(\hat{t}) = \lim_{\Delta\hat{t} \rightarrow 0} \frac{\hat{S}(\hat{t}) - \hat{S}(\hat{t} + \Delta\hat{t})}{\Delta\hat{t} \cdot \hat{S}(\hat{t})} \quad (7)$$

With this setup, it is easy to show that:

Proposition 5 *In the bidding equilibrium,*

$$\hat{h}(\hat{t}) = \frac{c}{1-\hat{t}k} \text{ for } \hat{t} < \frac{1}{k}$$

When $k = 0$, $S(t) = e^{-c\hat{t}}$

When $k > 0$, $\hat{S}(\hat{t}) = (1 - \hat{t}k)^{\frac{c}{k}}$

Proposition 5 forms the basis for my empirical analysis, as it establishes a way for to directly compare the survival and hazard rates of auctions for goods with different values. For example, the proposition indicates that an auction for a good with a value 50 will have the same probability of surviving to period 50 as an auction for a good with a value 100 surviving to period 100. Note that the proposition does not provide a way to compare the rates of auctions with different bidding increments.

3 Description of Data

3.1 Description of Swoopo

Swoopo is the largest and longest-running company that runs penny auctions (five of Swoopo’s competitors are discussed later in the paper). Swoopo was founded in Germany in 2005, and currently provides local versions of their website for other countries, such as the United Kingdom (started in December 2007), Spain (started in May 2008), the United States (started in August 2008). Nearly every auction is displayed simultaneously across all of these websites, with the current highest bid converted into local currency.¹⁸ Swoopo auctions consumer goods, such as televisions or appliances, as well as packages of bids for future auctions and cash payments. As of May 2009, Swoopo was running approximately 1,500 auctions with nearly 20,000 unique bidders each week.

The general format of auctions at Swoopo follows the description in Section 2.1: (1) players must bid the current high bid of the object plus a set *bidding increment*, (2) each bid costs a non-refundable fixed *bid cost*, and (3) each bid increases the duration of the auction by a small amount. While most companies that run penny auctions solely use a bidding increment of €0.01, Swoopo runs auctions with bidding increments of €0.10 (76% of the auctions), €0.01 (6%), and €0.00 (18%).¹⁹ The cost of a bid in Europe has stayed constant at €0.50 (the cost of a bid for most of my dataset was usually £0.50 and \$0.75 in the United Kingdom and the United States, respectively).²⁰

In most auctions, Swoopo allows the use of the *BidButler*, an automated bidding system available to all users. Users can program the BidButler to bid within a specific range of values and the BidButler will automatically place bids for the user when the timer nears zero.²¹ Certain auctions, called *Nailbiter Auctions* (9.5% of auctions), do not allow the use of the BidButler. As of November 2008, Swoopo also runs *Beginner Auctions* (10.8% of recent auctions), which are restricted to players that have never won an auction.

3.2 Description of Data

My data on Swoopo consists of two distinct datasets, both of which were captured using a multi-server website collection algorithm.

(1) Auction-Level Data

The auction-level data contains approximately 166,000 auctions for approximately 9,000 unique goods spanning from September 2005 to June 2009. The data was retrieved separately from Swoopo’s American, German, Spanish, and English websites. For each auction,

¹⁸For this paper, I will always refer to costs and prices in Euro.

¹⁹In July 2009, Swoopo changed the possible bidding increments to €0.01, €0.02, €0.05, and €0.20. My dataset ends before this change.

²⁰The cost of a bid in the United States was \$1.00 from September 2008 to December 2008. Note that incorporation of this change would only increase Swoopo’s estimated profit.

²¹If two players program a BidButler to run at the same time for the same auction, all the consecutive BidButler bids are placed immediately.

the dataset contains the item for auction, the item's value (see Section 3.3), the type of auction, the bidding increment, the final (highest) bid, the winning bidder, and the end time. From October 2007 on, the data also contains the final (highest) ten bidders for each auction. The summary statistics for this dataset are listed in the top portion of Table 1.

(2) Individual Bid-Level Data

The bid-level data contains approximately 13.3 million bids placed by 129,000 unique users on 18,000 auctions, and was captured in real-time from Swoopo's American website from late February 2009 to early June 2009.²² The algorithm has the ability to record new information every 2-3 seconds, depending on the Internet connection and the website's response time. As the Swoopo website posts a live feed of the last ten bids in each auction, the algorithm can capture the vast majority of bids even when bids are made very rapidly. The one exception is the situation in which multiple people use a BidButler at the same time in one auction. In this case, each player's automated bids are made instantaneously, causing the current high bid to jump immediately to the lowest of the high bounds of the two BidButlers. Fortunately, it is relatively simple to spot this issue in the data and infer the bids that are not listed. Using Nailbiter Auctions (in which BidButler use is forbidden), I estimate that the algorithm captures 96% of the bids in the periods when it is running. Each observation in this dataset contains the (unique) username of the bidder, the bid amount, the time of the bid, the timer level, and if the bid was placed by the BidButler.²³ Note that the auctions in this dataset are a subset of the auctions in the auction-level dataset. The summary statistics for this dataset are listed in bottom portion of Table 1.

In addition to data about Swoopo, I captured similar auction-level datasets for five of Swoopo's competitors: BidStick, RockyBid, GoBid, Zoozle, and BidRay. I will refer to this data briefly when I analyze the market for these auctions.

3.3 Value Estimation

Analyzing the auctions requires an accurate estimate of the value of the good. For each good, Swoopo publishes a visible "worth up to" price, which is essentially the manufacturer's recommended price for the item. This price is one potential measure of value, but it appears to be only useful as an upper bound. For example, Swoopo has held nearly 4,000 auctions involving 154 types of watches with "worth up to" prices of more than €500. However, the vast majority of these watches sell on Internet sites at heavy discounts from the "worth up to" price (20-40%). It is difficult, therefore, to justify the use of this amount as a measure of value if the auctioneer or participant can simply order the item from a reputable company at a far cheaper cost. That said, it is also unreasonable to search all producers for the lowest possible cost and use the result as a measure of value, as these producers could be disreputable or costly for either party to locate.

²²Due to various issues (including a change in the way that the website releases information), the capturing algorithm did not work from March 6th-March 8th and April 8th-April 11th. Furthermore, the efficiency of the algorithm improved with a change on March 18th.

²³The algorithm captures the time and timer level when the website was *accessed*, not at the time of the bid. The time and timer level can be imperfectly inferred from this information.

In order to strike a balance between these extremes, I estimate the value of items by using the price found at Amazon.com and Amazon.de for the exact same item and using the "worth up to" price if Amazon does not sell the item. I refer to this new value estimate as the *adjusted value* of the good.²⁴ As prices might have changed significantly over time, I only use Amazon prices for auctions later than December 2007 and scale the value in proportion to any observable changes in the "worth up to" price over time. Amazon sells only 20% of the unique consumer goods sold on Swoopo, but this accounts for 49% of all auctions involving consumer goods (goods that are sold in Amazon are likely to occur more in repeated auctions). For the goods that are sold at Amazon, the adjusted value is 74% of the "worth up to" price without shipping costs and 71% when shipping costs are added to each price (Amazon often has free shipping, while Swoopo charges for shipping). As the adjusted value is equal to the "worth up to" price for the 51% of the auctions for consumer goods that are not sold on Amazon, it still presumably overestimates the true value.²⁵

To test the validity of the measure of value, note that the equilibrium analysis (and general intuition) suggest that the winning bid of an auction should be positively correlated with the value of the object for auction. Therefore, a more accurate measure of value should show a higher correlation with the distribution of winning bids for the good. The correlation between the winning bids and the "worth up to" price is **0.521** (with a 95% confidence interval of (0.515,0.526)) for auctions with a €10 increment for the items I found on Amazon.²⁶ The correlation between the winning bids and the adjusted value is **0.708** (with a 95% confidence interval of (0.704,0.712)) for these auctions. A Fisher test of correlation equality confirms that the adjusted value is (dramatically) significantly more correlated with the winning bid, suggesting that it is a more accurate measure of value.

3.4 Profit Analysis

According to the equilibrium analysis above, one would not expect the auction format used by Swoopo to consistently produce more revenue than the market value of the auctioned good. *One of the main empirical findings of this paper is that this auction format consistently produces this level of revenue.* Averaging across goods, bidders collectively pay 52% over the adjusted value of the good, producing an average profit of €112. For the 166,000 auctions over four years in the data set, the auctioneer's profit for running the auction is €18M.²⁷ The distribution of monetary profit and percentage profit across all auctions is shown in Figure 1 (with the top and bottom 1% of auctions trimmed). Perhaps surprisingly, the auctioneer's profit is below the value of the good for a slight majority (51%) of the items. Table 2

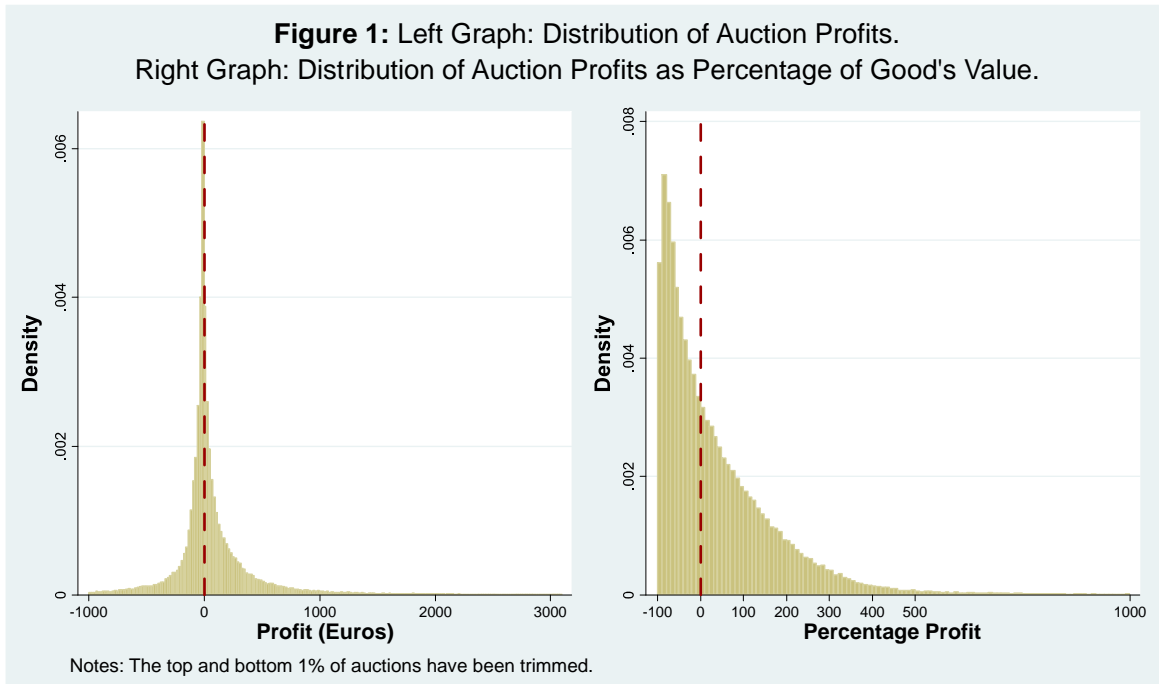
²⁴This is a somewhat similar idea to that in Ariely and Simonson (2003), who document that 98.8 percent of eBay prices for CDs, books, and movies are higher than the lowest online price found with a 10 minute search. My search is much more simplistic (and perhaps, realistic). I only search on Amazon and only place the exact title of the Swoopo object in Amazon's search engine for a result.

²⁵The main results of the paper are unchanged when run only on the subset of goods sold at Amazon.

²⁶Note that I cannot compare aggregate data across auctions with different bid increments for these correlations, as the distribution of final bids of auctions for the same item will be different. The results are robust to using the (less common) bid increments of €0.00 and €0.01.

²⁷This profit measure does not include the tendency for people to buy multi-bid packages but not use all of the bids ("breakage"). The bid-level data suggests that this is a significant source of revenue for Swoopo.

breaks down the profits and profit percentage by the type of good and the increment level of the auction. Notice that auctions involving cash and bid packages (items with the clearest value) produce profit margins of over 104% and 214%, respectively. Consumer goods, which are potentially overvalued by the adjusted value measure, still lead to an estimated average profit margin of 33%.



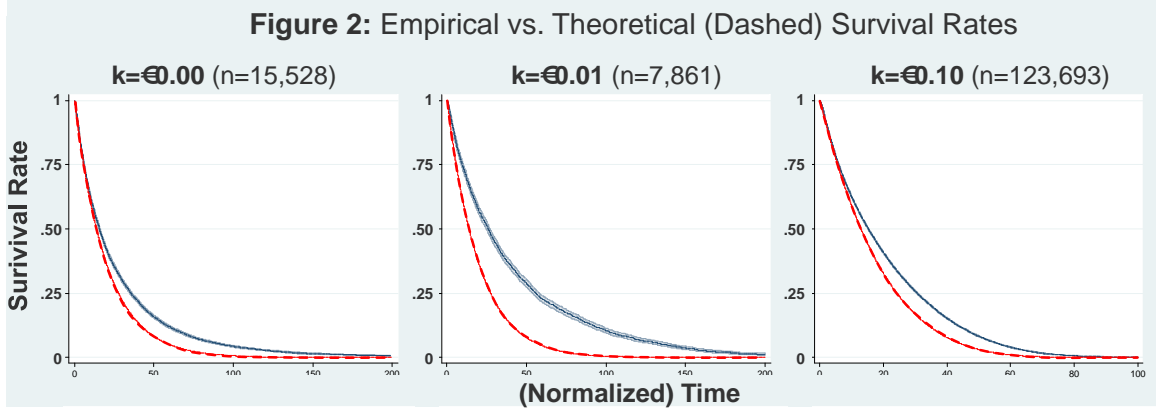
4 Empirical Tests of Model

4.1 Research Question and Empirical Strategy

The theoretical analysis above yields multiple empirical predictions about the survival and hazard functions of penny auctions. I am able to identify the empirical survival and hazard rates of the auctions as the final auction outcomes (the final and highest bid in auction) are stochastic. Recall from Section 2.4 that normalizing the time measure of the auctions by the value of the goods allows the comparison of these rates across auctions for goods with different values (given that they have equal bid increments).

4.2 Consumer Goods

Figure 2 displays the Kaplan-Meier Survival Estimates (and confidence intervals, which are extremely tight) of the normalized time measure of consumer goods auctions for the three increment levels, along with the survival rates predicted by the equilibrium hazard rates derived in Section 2. The survival functions are consistently higher than the equilibrium



survival functions for each normalized time measure, which is expected given the profit statistics above and the fact that the auctioneer receives more profit from auctions that last longer. For auctions with bidding increments of €0.10 and €0.00 (which represent 94% of the auctions in the dataset), it appears that the survival rates follow the equilibrium survival rates closely for early normalized time measures before rising consistently above the predicted survival rates.

Survival rates are difficult to interpret because they represent the *cumulative* effect of auction terminations up to each normalized time period. The empirical hazard functions are more illustrative of players' behavior at each point. Figure 3 displays the smoothed hazard rates with confidence intervals along with the hazard functions predicted by equilibrium strategies for each increment level.²⁸ Notice that the equilibrium hazard functions for the different increments are the same at the beginning of the auction (as the bids always start at zero), stay constant if the increment is €0.00 (as the current bid amount is always constant), and rise more steeply through time with higher increments (as the current bid rises faster with a higher increment). Most interesting, for auctions with bid increments of €0.00 or €0.10, the hazard function is very close to that predicted by equilibrium analysis in beginning periods of the auction. However, for all auctions, the deviation of the empirical hazard function below the equilibrium hazard function increases significantly over time.

While the hazard functions are suggestive of the global strategies of the players, it is difficult to interpret the economic magnitude of the deviations from the predicted actions. In order to determine this magnitude, note that the bidder at period $t - 1$ is paying the auctioneer a bid cost c in trade for a probability of $h(t)$ of winning the net value of the good $(v - tk)$ at time t . In other words (assuming risk-neutrality of both parties), the auctioneer is selling the bidder a stochastic good with an expected value of $h(t)(v - tk)$ for a price c at time t . Therefore, I define *percent instantaneous markup* as the percentage of the cost of this good above its value at the point it is sold:

Definition 3 *Percent Instantaneous Markup (at time t)* = $(\frac{c}{h(t)(v - tk)} - 1) * 100$

²⁸For this estimation, I used an Epanechnikov kernel and a 10 unit bandwidth, using the method described by Breslow and Day (1986) and Klein and Moeschberger (2003). The graphs are robust to different kernel choices and change as expected with different bandwidths.

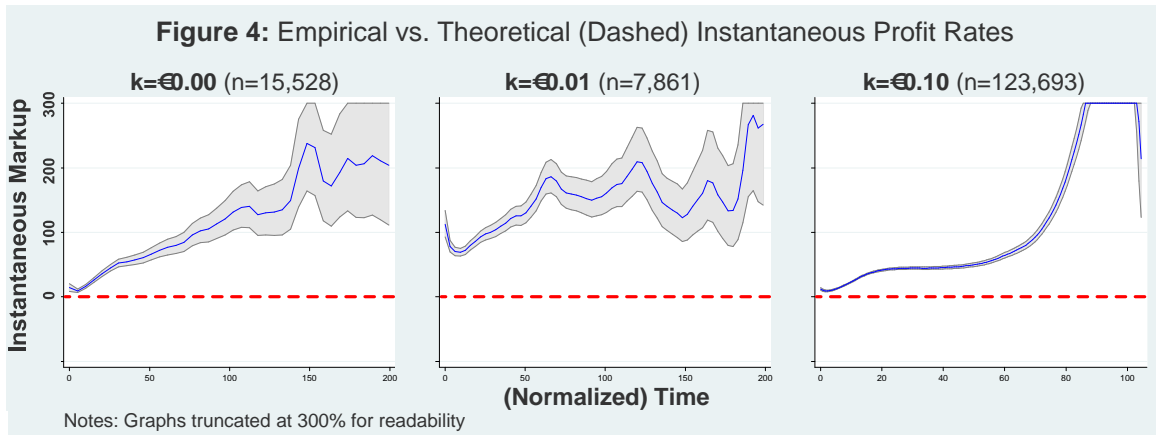
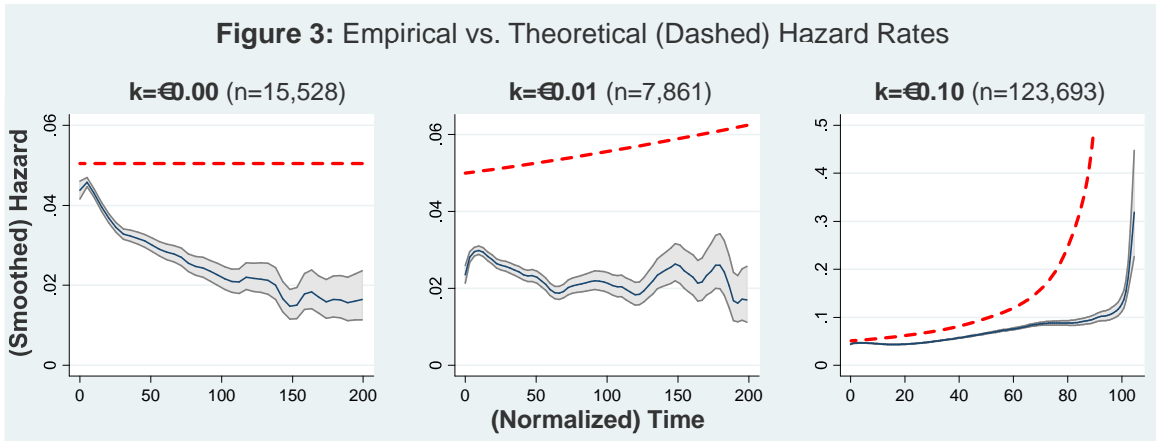
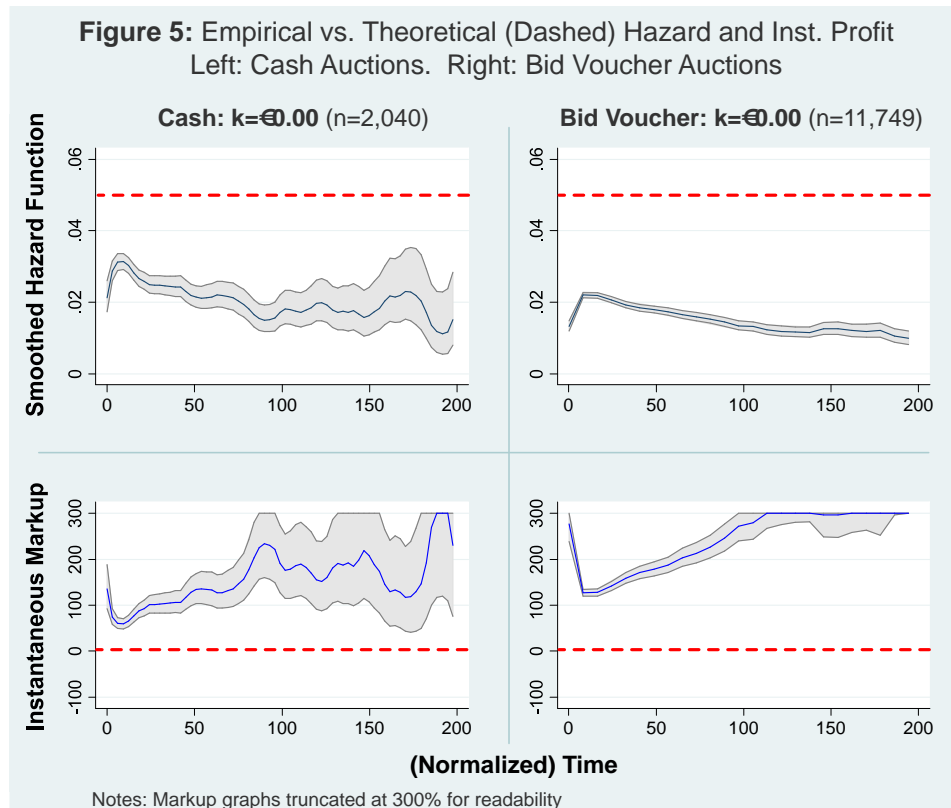


Figure 4 displays the percent instantaneous markup and confidence intervals along with the predicted percent markup for each increment level at each (normalized) time period. Note that the predicted markup is always zero because the expected equilibrium payoff from a bid is equal to the cost of the bid at each period. For auctions with bidding increments of $\text{€}0.10$ and $\text{€}0.00$, the empirical instantaneous markup starts near this level, but rises over the course of the auction to 200-300%. This estimate suggests that, if an auction survives sufficiently long, players are willing to pay $\text{€}0.50$ (the bid cost) for a good with an expected value of $\text{€}0.12$ - $\text{€}0.16$. Therefore, rather than making uniform profit throughout the auction, the auctioneer is making a large amount of instantaneous profit at the end of the auction.

4.3 Cash and Bid Vouchers

One concern about the results for the consumer good auctions is that the measure of the true value of the good is noisy and values presumably differ across users. Furthermore, for the auctions with positive bid increments, the net value of the good changes over time (as the current high bid rises through the game). To address these issues, this section focuses on the 18,790 auctions for cash payments and vouchers for bids. In these auctions, the value



of the item is arguably common across participants (bid vouchers are immediately available on Swoopo for a common fixed price) and, focusing on auction with a bid increment of €0.00 (73% of the auctions), the net value is constant throughout the auction.

As suggested earlier in Table 2, the profits for these auctions are significantly higher than those for consumer goods. The reasons for this result are unclear, although it might be that the measure of value is more accurate in these auctions, more people are attracted to these auctions, or a specific subset of players bid on these auctions. The hazard rates and percent instantaneous profits for these auctions are shown in Figure 5. As would be expected from the profit estimations for these auctions, the empirical hazard functions are dramatically lower than the equilibrium hazard functions, even at the beginning of the auction. However, consistent with the important qualitative features of the consumer good auctions results, there is still an upward trend in the empirical instantaneous profits for both types of auctions. In the cash auctions, the instantaneous profit is 50-100% during the first stages of the auction, but rises to around 150-300% in the later stages. In the bid voucher auctions, the instantaneous profit is between 100-150% during the first stages of the auction (minus a very short initial period) but steadily rises to over 300% in the later stages.

5 Modeling Sunk Costs

The previous section demonstrated that the deviation of the empirical hazard rates from the predicted hazard rates changes dramatically over the course of an auction. Therefore, any explanation for this behavior must include a factor that changes as an auction progresses. For example, a simple misunderstanding of the rules of the game or an consistent underprediction of the number of participants cannot explain these results alone as they would predict constant deviations from the predicted hazard rate. Furthermore, any potential explanation must account for behavior in auctions with different bidding increments. For example, if players do not account for the current bid over time, they will overbid in decreasing-value auctions, but this reasoning cannot account for the empirical decreasing hazard observed in constant-value auctions. Even with these constraints, there are still a variety of explanations that cannot be separately identified with my data. For example, the behavior could be explained by a tendency for players to make worse predictions about others' behavior as the auction continues, possibly because there are fewer learning opportunities at later stages in the auction (these stages are less likely to occur). In this section, rather than discussing all of the possibilities, I focus on the naive sunk cost fallacy as a potential explanation and demonstrate that modifying the model to include this effect leads to qualitatively similar hazard rates to those observed in the data.

Following Ashraf et al (2008), I use the framework proposed by Eyster (2002) to model sunk costs.²⁹ The reader is referred to that paper for technical details of the utility function. Applying Eyster's model and terminology, agents in the modified model desire "consistency" in their decisions and pay a psychological cost, which I call "regret", if they spend money on bids and do not win the auction, weighted by the parameter $\rho \in [0, \infty)$ in the utility function. As a result, agents receive less utility from exiting the auction as they pay for more bids, even though these costs are sunk. Note that this modification alone will cause agents to *underbid* as they will require a premium (in the form of a higher probability of winning) to continue at any stage to offset their (correctly predicted) future psychological losses from the sunk costs. Therefore, I follow Eyster in assuming that agents consider the effect of their current decisions on their future utility, but they naively believe that their weight on future regret will be $\rho(1 - \eta)$ with $\eta \in [0, 1]$.³⁰

In the interest of simplicity, I deviate from the Eyster's multiple period model in one substantial way. Rather than assuming that an agent feels regret for all decisions in the game, I assume that an agent simply feels regret from his initial decision (to play or not play in the game). To elucidate this difference, consider an agent who leaves the game after bidding 10 times, with bids costing 1 unit. In Eyster's model, the agent experiences regret from each past decision to stay in the auction for a total of 55ρ units (he would have saved 10 units had he exited instead of placing the first bid, 9 units if he had exited instead of placing the second bid, 8 units...). In my model, the agent simply experiences 10ρ units of regret as he would have saved 10 units from not playing the game. As one could just

²⁹In fact, Eyster considers a standard discrete war of attrition model as an example in his paper, producing similar results.

³⁰A note on terminology: I choose to use η instead of v (Eyster's parameter of naivety) to avoid confusion with the value of the object v .

rescale ρ to account for this difference, the substantial difference between the models lies in the growth of regret as the game continues. In Eyster's model, regret grows "triangularly" over time, from 1 to 3 to 6 to 10, etc. In my model, regret grows linearly over time, from 1 to 2 to 3 to 4, etc. I do not believe that there is a good reason to choose either model over the other in this application, so I proceed with the linear model in the interest of simplicity.

Specifically, consider an agent who has placed b bids up until time period t . The total utility of the agent from *never bidding again* becomes:

$$-bc - \rho bc \tag{8}$$

That is, the agent experiences the monetary loss ($-bc$) of the bids as well as regret ($-\rho bc$) from deciding to play the game in the first place.

Similarly, if an agent bids in period t , does not win in the next period, and never bids again, he will receive utility of:

$$-(b+1)c - \rho(b+1)c \tag{9}$$

However, due to naivety, he (mistakenly) perceives that his feeling of regret will be lower than it really is:

$$-(b+1)c - (1-\eta)\rho(b+1)c \tag{10}$$

The case in which an agent bids and wins the auction in the next period is slightly more complicated. The level of regret depends on the situation. If the net value of the item is weakly higher than the total cost the agent, the agent does not regret his decision to enter the auction. In this case, he simply receives the utility of:

$$v - ty - (b+1)c \tag{11}$$

Notice that bc (the monetary bid cost up to period t) occurs in equations 8, 10, and 11, which is consistent with bc as a sunk cost. However, the regret term only occurs if the person exits the auction, which is consistent with the notion of the sunk cost fallacy. If the person is naive, he believes that the weight on the regret will be lower in the future than today.

Alternatively, if the net value of the auction is higher than the value of the object, the agent *does* regret his decision to enter the auction. In this case, he receives utility:

$$v - ty - (b+1)c - \rho(b+1)c$$

Note that, in this situation, the regret term appears in the utility term in all situations, so the agent fully recognizes the sunk cost (as before, if the agent is naive, he perceives this term to be $v - ty - (b+1)c - (1-\eta)\rho(b+1)c$).

In order for this modification to affect equilibrium behavior, agents must be able to condition their strategies on the number of bids each player has made (because this now affects agents' payoffs). Following the general path of Eyster's solution (in which naive players correctly perceive other's true strategies, although they misperceive their own) yields the following *outcome* of the preferred equilibrium and the hazard rate, which is summarized in Proposition 6

Proposition 6 *There is an equilibrium of the modified game in which:*

$$h(t) = \left\{ \begin{array}{ll} \max\left\{0, \frac{c+c\rho-c\rho\eta(\frac{t}{2}+1)}{v-ty+(1-\eta)c\rho(\frac{t}{2}+1)}\right\} & \text{for } t \leq \frac{2(v-c)}{c+2y} \\ \max\left\{0, \frac{c+c\rho-c\rho\eta(\frac{t}{2}+1)}{v-ty}\right\} & \text{for } t > \frac{2(v-c)}{c+2y} \end{array} \right\}$$

The effect of the regret over spending fixed costs is slightly complicated. At the beginning of the auction, agents with regret are less likely to bid than agents without regret because they have no current mistakes to regret and they realize (to the extent that they are sophisticated) that they will have to pay regret costs in the future if they bet and lose. As the auction proceeds, this difference diminishes as agents amass larger sunk costs through bidding. At some point, if agents are naive, the game continues with higher probability than with normal agents because agents (incorrectly) believe that their amassed fixed costs will be lessened if they bid and then drop out in the following period. If agents are particularly naive, they can reach a point in which no one drops out, with bidders staying in the game only because they (incorrectly) believe that bidding and dropping out tomorrow will reduce the regret from their large fixed costs.

Figure 6 displays the equilibrium hazard rates for $\rho = .3$, $c = \text{€}.50$, for an increment of $\text{€}.00$ as η rises. Note that the curves with higher levels of naivety display the same qualitative features as those in the empirical data.

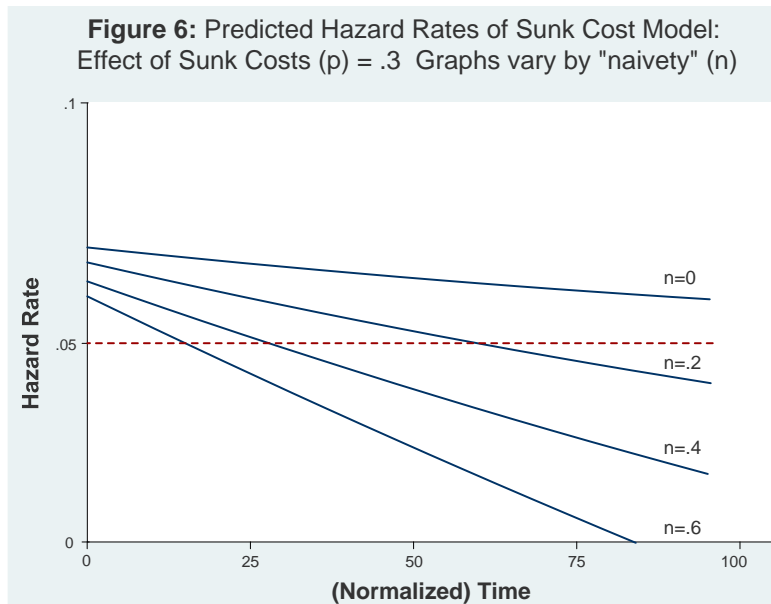
6 Experience, Strategies, and Profits

6.1 Effect of Experience on Profits

Research Question and Empirical Strategy

The empirical results above demonstrate that, regardless of the reason(s), players in aggregate are bidding in a way that consistently leads them to make (reasonably large) negative expected payoffs. This section addresses the effect of experience on the expected payoffs of the players.

Broadly, there are two strategic ways in which a player could improve their profits. The first is simply to stop bidding. If the empirical hazard rates seen in previous sections *were consistent across all auctions*, the best response of a monetary-maximizing player would be to never bid, as bidding leads to a negative expected value throughout the auction. However, to the extent that there is heterogeneity in auction hazard rates (in different items or at



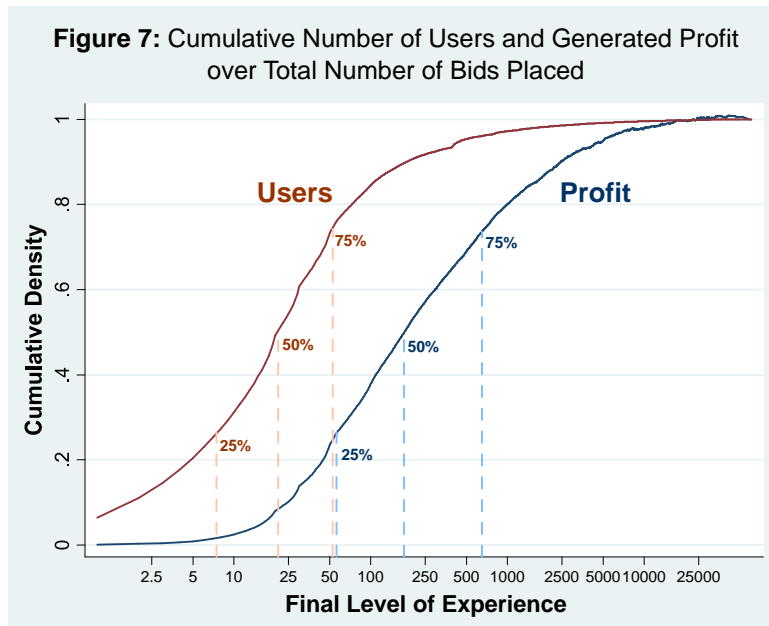
different times of day) and one's strategy can affect other player's actions (such as increasing the future hazard rate by playing very aggressively), there might exist bidding strategies that increase profits. I discuss the effect of experience on these two broad strategies.

There are multiple potential measures of "experience." For my analysis, I define the experience of a player at a point in time as the number of bids made by that player in all auctions *before* that point in time³¹ On potential concern with this measure is that, as the coverage of the bid-level data starts long after Swoopo began running auctions, some players enter this dataset with previous experience from past auctions. Fortunately, the auction-level data contains a list of the final ten bidders for the auctions before the start of the bid-level data, which I use to estimate the number of bids placed by each player before entering the bid-level dataset. To accomplish this, I first perform an OLS regression to estimate the relationship between the number of appearances in the top ten lists and the number of bids made by a player using the bid-level data.³² Second, given the number appearances in the top ten before the individual dataset, I use this estimate to predict the number of bids each player made in this time period. While this measure is imperfect (it does not capture players that bid before the bid-level dataset without finishing in the top ten and assumes that the relationship between bids and appearances in the top ten is consistent across time), it does provide a rough measure of the number of previous bids for players that enter with large amounts of experience.

Using this measure, the bid level data provides the necessary variation to identify the relationship between profits and experience. Specifically, the data contains natural variation in the profit from each bid (winning is dependant on other players' choices in the following period, which is not deterministic) as well as experience across users (some users enter with

³¹The qualitative results below are robust to using different experience measures, such as the number of auctions previously played or the total time previously spent on the site.

³²According to the regression estimates, an appearance in a top ten list is associated with 59.6 bids.



more experience and some users stay longer than others) and within users (the heaviest 10% of users place an average of nearly 800 bids in my dataset).

(Indirect) Evidence that Bidders Learn to Stop Bidding

Figure 7 displays the cumulative density of users against the users' experience at the end of the dataset (in log scale). The figure demonstrates that many players stop bidding after placing relative few bids. For example, 75% of users stop bidding before placing 50 bids and 86% stop bidding before placing 100 bids. While there are many reasons that users might leave the site, this statistic is indirect evidence that players learn that, on average, bidding is not a profitable strategy.

For reference, Figure 7 also displays the a "cumulative profit curve," which plots the percentage of the auctioneer's total profit produced by players at or below each level of experience. For example, while 75% of users make less than 50 bids, the cumulative profit curve shows that these users only generate 25% of the total profit. Conversely, the top 11% of bidders in terms of experience create more than 50% of the total profit.

Experienced Players Learn Strategies to Increase Profits

In the previous section, I demonstrated that players that make many bids account for a large percentage of Swoopo's profit. In this section, I examine the relationship between experience and expected bid profits for these players. First, I use a non-parametric regression to show a clear positive relationship between experience and the expected profit from each bid. Then, I parametrize the regression to demonstrate that this relationship is highly statistically significant. Finally, in order to control for potential selection effects, I add

user fixed effects and allow learning rates to vary across users, demonstrating that learning partially drives the relationship between experience and profits.

As discussed in Section 4.2, each bid can be interpreted as an independent "bet" in which the bidder pays the bid cost in exchange for a chance to win the net value of the good in the following round. Note that the expected profit from this bet lies between losing the bid cost ($-\text{€}0.50$) and the net value of the object $(v - (t + 1)k)$. There are two potential reasons that higher levels of experience might be associated with an increase in the expected profit of this bet. First, users might learn better bidding strategies (such as bidding on certain items or bidding "aggressively") through playing the game. Second, users that naturally use better strategies might be more likely to continue playing the game, thus leading to a correlation between experience and high expected value through selection.

Rearranging the dataset into an (incomplete) panel dataset in which users are indexed by i and the order of the bids that an individual places is indexed by t , let y_{it}^U be the payoff of user i 's t th bid.³³ Then, a general model of the effect of experience on profit is

$$y_{it}^U = L_i(t) + \varepsilon_{it} \tag{12}$$

where $L_i(t)$ is individual i 's learning function and ε_{it} is the error that arises from the stochastic nature of the auction. As a first step towards understanding the learning functions, Figure 8 displays a non-parametric regression constraining $L_i(t) = L(t)$ as well as a histogram of the number of bids made at each experience level. Clearly, there appears to be a positive relationship between the profit of a bid and the level of experience of the bidder. A player with no experience can expect to lose $\text{€}0.40$ per each $\text{€}0.50$ bid, while those with very high experience levels have slightly positive expected payoffs per each bid. However, note that this positive effect requires a relatively large amount of experience: raising the expected value of a bid to near zero requires an experience level of nearly 10,000 bids.

Following the quadratic shape of the non-parametric regression and still constraining $L_i(t) = L(t)$, I first parameterize the model as:

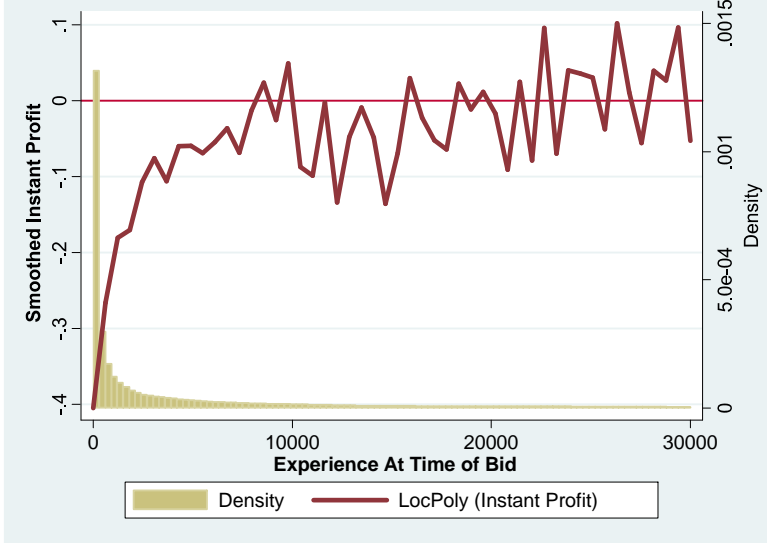
$$y_{it}^U = \beta_0 + \beta_1 t + \beta_2 t^2 + \varepsilon_{it} \tag{13}$$

with the results shown in the first column of Table 3 (with standard errors clustered on users). These estimates show that, on average, there is an economically and statistically significant concave relationship between experience and profits, which demonstrates that *there are strategies which consistently yield higher profits*. Specifically, for each 1000 bids, players initially increase the expected return from each $\text{€}0.50$ bid by $\text{€}0.03$ from a baseline of $\text{€}0.17$.

However, as discussed above, it is not clear that this result is due to individual learning. It is possible that individuals with larger coefficients continue in the game for longer, leading t to be positively correlated with the error term. To help mitigate this selection problem, I

³³Note that, confusingly, this is a different t than in the theoretical section. In that section, t represented a change in the bid level, while t here represents a level of experience for a user.

Figure 8: Local Poly. Regression of Profit on Experience.
Histogram of Experience at Each Bid Also Shown.



estimate the model with fixed effects for users:

$$y_{it}^U = \beta_1 t + \beta_2 t^2 + \phi_i + \varepsilon_{it} \quad (14)$$

with the results shown in the second column of Table 3. This specification suggests that, to the extent that the heterogeneity in learning functions is captured by an added constant, there is a selection effect, but learning alone does play a role in the positive association between experience and profits. However, this specification does not account for the possibility that learning *rates* differ across individuals. If these learning rates are correlated with high levels of experience, there will still be selection bias. To capture this effect, I define T_i as the highest level of experience achieved by a player and allow the learning rate to vary linearly with T_i :

$$y_{it}^U = (\alpha_0 + \alpha_1 T_i)t + \beta_2 t^2 + \phi_i + \varepsilon_{it} \quad (15)$$

with the results shown in the third column of Table 3. Interestingly, the coefficient estimate $\hat{\alpha}_1$ is negative, suggesting that there is a negative selection effect with respect to learning rates.³⁴ For example, the estimated linear effect on profits of each additional 1000 bids for a player that leaves the game with very little experience is nearly €0.035, while that for a player that leaves after making 10,000 bids is only €0.009. This relationship is consistent with the previous findings that more experienced players have high natural expected profits (on average) and that there are decreasing returns to experience as expected profits rise. Crucially, both specifications 15 and 14 maintain an economically and statistically significant positive estimate on the experience coefficient even when controlling for selection effects,

³⁴The coefficient estimate $\hat{\alpha}_2$ on a quadratic interaction effect $T_i^2 t$ is insignificant and does not change the results.

suggesting that *players do learn strategies that increase their expected profit (especially at low levels of experience)*.

6.2 Strategies Used by Experienced Players to Increase Profits

This section determines the specific strategies that allow players with higher levels of experience to produce higher expected profits. My basic empirical strategy (discussed in detail below) is to estimate the effect of experience on profit (as in specification 13), while controlling for the effect of using specific strategies. To the extent that the relationship between higher experience levels and profit is driven by the use of each strategy, the coefficient on experience will be reduced accordingly. I find experienced players are not making higher expected profits as a result of time-based strategies (bidding at certain times of the day or days of the week) or item-based strategies (bidding on certain items), but that "aggressive bidding" strategies (bidding immediately whenever possible) account for the majority of the profit associated with higher levels of experience.

Formally, note that any change the expected profits from a bid *must* be driven by the use of different bidding strategies. Specifically, letting variables s_1, s_2, \dots denote the level of use of each strategy (and allowing strategies to represent any level of interaction of multiple strategies), it must be that

$$y^U = S_1(s_1) + S_2(s_2) + \dots + \varepsilon \quad (16)$$

for some functions $S_1(\cdot), S_2(\cdot), \dots$, with ε again representing the natural stochastic nature of winning the auction. Therefore, any estimate of experience on expected profits (like those in Section 6.1) must be indirectly capturing the effect of experience on the use of these different strategies, so that

$$y_{it}^U = L_i(t) + \varepsilon_{it} = S_1(L_i^1(t)) + S_2(L_i^2(t)) + \dots + \varepsilon_{it} \quad (17)$$

for some learning functions $L_i^1(\cdot), L_i^2(\cdot), \dots$. To parameterize this model, I follow the first (quadratic, user-consistent) parametrization of the effect of experience above for each learning function, so that

$$s_l = L_i^l(t) = \delta_0^l + \delta_1^l t + \delta_2^l t^2 \quad (18)$$

for each strategy $l = 1, 2, 3, \dots$. Finally, I linearly parametrize the $S_l(\cdot)$ functions, so that:

$$y_{it}^U = \gamma_0 + \gamma_1(\delta_0^1 + \delta_1^1 t + \delta_2^1 t^2) + \gamma_2(\delta_0^2 + \delta_1^2 t + \delta_2^2 t^2) + \dots + \varepsilon_{it} \quad (19)$$

Using this interpretation, the estimated coefficient $\hat{\beta}_1$ in specification 13 is a consistent estimate of $\gamma_1 \delta_1^1 + \gamma_2 \delta_1^2 + \dots$, the total effect of t on the use of each strategy (δ_1^l) multiplied by the effect of that strategy on profits (γ_l). Crucially, note that if s_{1it} (the use of strategy 1 by person i at time t) is observable and included in the regression, the estimated coefficient on t becomes a consistent estimate of $\gamma_2 \delta_2^2 + \gamma_3 \delta_3^3 + \dots$, the total effect on profit of the use of

all strategies except strategy 1. Note that it is possible to produce a consistent estimate of $\gamma_1 \delta_1^1$ (the effect of profits of the linear change in the use of strategy 1 through experience) by differencing these estimates.

In order to focus on strategies in the simplest version of the game, I report results for Nailbiter Auctions, in which the BidButler is not allowed. The first column of Table 4 presents the results of a regression of experience on profits for Nailbiter Auctions. Following the logic above, the rest of the table presents the same regression, controlling for a variety of strategies. For example, the second column displays the regression while controlling for item fixed effects. While there are significant and positive coefficients on certain items (suggesting that bidding on these items leads to significantly higher profits), the coefficients on experience remain virtually unchanged, suggesting that the strategy of bidding on certain items is not driving the relationship between experience and profits. Similarly, the third column, which displays the regression controlling for time-of-day and day-of-week fixed effects, suggests that time-based strategies are also not driving this relationship.

There are strategies that have significant effects on the estimates of the experience coefficient. Broadly, the most important appears to be "bidding aggressively," in which a player bids extremely quickly following another player's bids (rather than wait until the timer runs to a few seconds) and bids repeatedly for a large number of bids. Column four of Table 4 presents the results of the regression controlling for these strategies, with the "seconds" categorical variables representing the number of seconds from the previous bid (zero seconds representing the most aggressive bid) and the "streak" categorical variables representing the number of bids made in the auction previous to this bid (with higher numbers representing more aggressive bids). First, note that *"aggressive bidding" strategies have a very significant effect on profitability of each bid*. Bidding extremely quickly after the previous bid (within one second) raises the expected profit by €0.58 over waiting over 20 seconds to bid. Having previously bid more than 20 times raises the expected profit by €0.25 over bidding for the first time in the auction. Second, note that the coefficient on experience has been reduced dramatically, suggesting that *increased use of aggressive bidding by experienced players drives the majority of the increase in profits from experience*. In fact, €0.02 of the €0.03 gain in expected profits associated with an additional 1000 bids arises from the increased use of aggressive bid strategies.

7 Supplier Behavior

In the previous sections, I examine the behavior of participants in penny auctions. In this section, I analyze the behavior of suppliers.³⁵ First, I calculate the optimality of Swoopo's behavior by separately estimating Swoopo's actual and optimal *supply rule* (the number of auctions to provide for a given number of participants). This analysis suggests that high

³⁵Some observers have commented on the potential for Swoopo to make shill (fake) bids in order to keep the auction running. Based on my analysis of patterns in the bid-level data, I do not find any evidence of the most obvious shill bidding techniques (creating fake bidders and using them nearly exclusively for shill bids). I cannot rule out the possibility that Swoopo is using more sophisticated techniques that are more difficult to identify.

supplier profits require a consistent and large group of users participating in the auction and that over-supplying auctions can be costly to the firm.³⁶ Interestingly, this finding suggests a potential barrier to new entrants to the market. While there is very little cost to creating a similar auction site, entrants must over-supply auctions in order to attract a larger userbase, but the attraction process takes time and consequently this leads to significant short-term losses. I then show that evidence from five of Swoopo’s top competitors is consistent with this conclusion, as none of the firms are making relatively large profits and three are making negative profits.

7.1 Supply Rules: Empirical Strategy

In the following sections, I study the optimal provision of auctions for a given number of users at Swoopo. First, I identify Swoopo’s *actual* supply rule, which is the number of auctions it attempts to supply for a given number of users on the site. Swoopo releases auctions with very high initial timers, so it must predict the number of users on the site in the future when the auction’s timer will become very low and players will start to bid (note that, as Swoopo is an automated website and faces no time-based constraints, the changes in the number of expected users across the day are independent of supply capacity). Therefore, it is possible to determine Swoopo’s supply rule by matching the expected number of auctions (given the release times) with the expected number of users on the site (given past user data). However, as Swoopo does not adjust the number of auctions in real-time, there is significant natural variation in the number of auctions and users once a point in time is actually reached, due to the natural stochastic nature of the ending times of auctions and the natural randomness in the number of users on the site. I use this variation to determine the profit curves from supplying a given number of auctions for a given number of users, which can be used to determine the *optimal* short-term supply curve.

Note that, unlike Swoopo, I cannot determine the precise number of users looking at the auctions at any instant in time. As a proxy for this measure, I calculate the number of bidders that placed a bid within 15 minutes of that time, which I call *active users*.³⁷ I define an *active auction* as one in which the timer of the auction shows less than one minute. To create the dataset for determining Swoopo’s optimal supply curve, I determine the average number of active users and auctions in each ten-second interval in my dataset.³⁸ As I mention in Section 3.2, the data capturing algorithm was improved on March 18th. In order to keep the measure of number of users consistent over different days, I only use data after that point.

³⁶Note that the theoretical analysis does not make this prediction. In a theoretical model, players adjust correctly to the number of players (or expected players) and the hazard rates remain consistent with the theoretical predictions given changes in the number of users.

³⁷The qualitative results are robust to changes in the intervals of time.

³⁸The qualitative results are robust to time windows from ten seconds to one hour (at which point there is not enough data to accurately estimate the coefficient).

7.2 Swoopo’s Actual Supply Rule

The goal of this section is to determine Swoopo’s chosen number of auctions for a given number of active users (the *supply curve*) at each point in time (each observation represents a ten second interval, as discussed above). Importantly, note that Swoopo releases auctions with many hours on the timer (12-24 hours) and that Swoopo is an automated website and therefore faces no time-constraints on supply.

In order to determine Swoopo’s expected number of auctions for a given point in time, one can use the initial timer, the value of the object, and the empirical survival rates in Section 4 to estimate a expected survival function $S(t; v)$ for each auction. $S(t; v)$ maps each time period t into the probability that the auction of an item with value v is active at that point.³⁹ Using this survival function, I estimate Swoopo’s desired supply at time t , \tilde{Q}_{St} , as

$$\tilde{Q}_{St} = \sum_{Auctions} S(v, t) \quad (20)$$

When Swoopo releases this auctions, it must estimate the number of users that will be on the site when the auction becomes active (the timer shows less than one minute). Luckily, the number of users Q_{Dt} varies predictably depending on the time of the day and the day of the week, as demonstrated by the fact that the regression of users on time:

$$Q_{Dt} = \beta_0 + \beta_1 TimeDummies + \varepsilon_t \quad (21)$$

yields an R^2 of .716.⁴⁰ The predicted values \hat{Q}_{Dt} from this regression are shown with the bold line in Figure 9.

The relationship between the expected supply and the expected demand is Swoopo’s actual supply rule (note that there is no concern of endogeneity in this regression as there is no way for the expected supply to impact the *expected* demand when an auction is released):

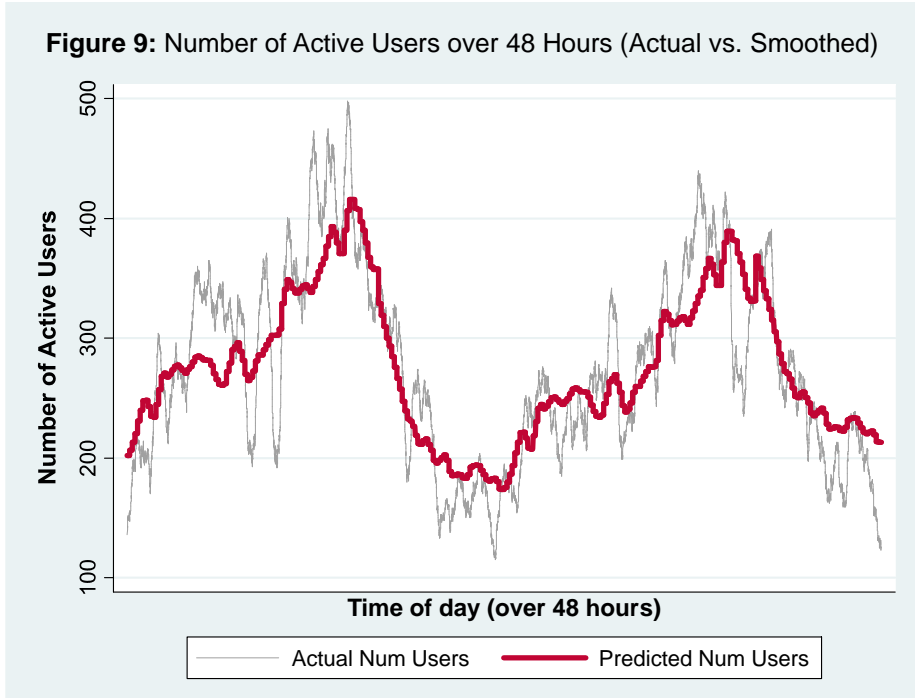
$$\tilde{Q}_{St} = \beta_0 + \beta_1 \hat{Q}_{Dt} + \varepsilon_t \quad (22)$$

with the results shown in the first column of Table 5. The results are highly significant and suggest that Swoopo attempts to supply a new auction for every 42 new active users. Given the nature of the data, there is a concern that the error term is serially correlated. Rather than attempt to correct for this correlation, which is potentially complicated in nature⁴¹, I run the same regression using only every 360th observation (each hour), with the results

³⁹Note that this is not the same as the survival function in the theoretical section. In that section, t represented a change in the bid level, while t here represents ten second intervals of clock time. As each rise in the bid level is associated with an average 9 second rise in the timer, these measures are related, but not the same.

⁴⁰My chosen specification uses dummies for time of day (broken into 10 minute sections) and individual day. A variety of specifications with different time-dummies yield nearly identical results.

⁴¹Given the shape of the survival function, it is not possible to perfectly match the expected supply at each point in time with an arbitrary desired supply, which leads to error, captured with ε_t . Modelling the effect of this constraint on ε_t is difficult.



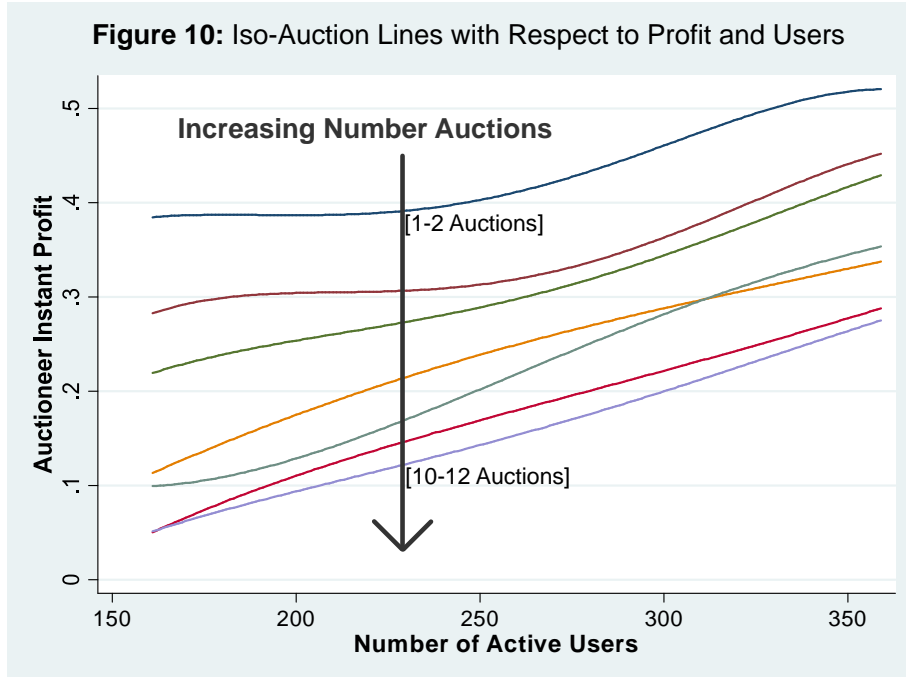
shown in the second column of table 5.⁴² The results remain largely unchanged. In order to check the potential for a non-linear supply, I run a regression with a quadratic term, with the results shown in the third column of Table 5. As I show in Section 7.4, the estimated supply curve is largely unchanged for the normal range of active users.

7.3 Swoopo’s Optimal (Short Term) Supply Rule

Even though Swoopo attempts to match the number of active auctions with the number of active users, there is significant variation in the actual number of active auctions for a given number of active users (and vice versa). Specifically, for a given number of users, the number of active auctions is nearly normally distributed around the desired number of auctions with a standard deviation of 2.83. This variation arises from the natural stochastic nature of the ending times of auctions and the number of users on the site. I use this variation to compare the (instantaneous) profit from supplying different number of auctions given a number of users to determine the optimal short-term supply curve. Note that as I have no exogenous variation in the long-term supply, I cannot identify the effect of changes in supply on long-term profit (presumably, a higher supply will attract more users to the site in the long-term, affecting profits). Therefore, my estimation of the optimal supply curve does not take into account any potential effects of supply on demand in the long-term.

Formally, let y_{jt}^A be the auctioneer’s payoff at time t from auction j , which must lie between the bid cost (€50) and the net value of the object. Then, a general model of

⁴²The results are consistent regardless of the chosen interval. I report results for one hour as this is the largest interval with enough observations to produce reasonably significant results.



the effect of the number of users Q_{Dt} , the number of auctions Q_{St} , and a vector of auction characteristics \vec{x}_j on profit is

$$y_{jt}^A = P(Q_{Dt}, Q_{St}, \vec{x}_j) + \varepsilon_{jt} \quad (23)$$

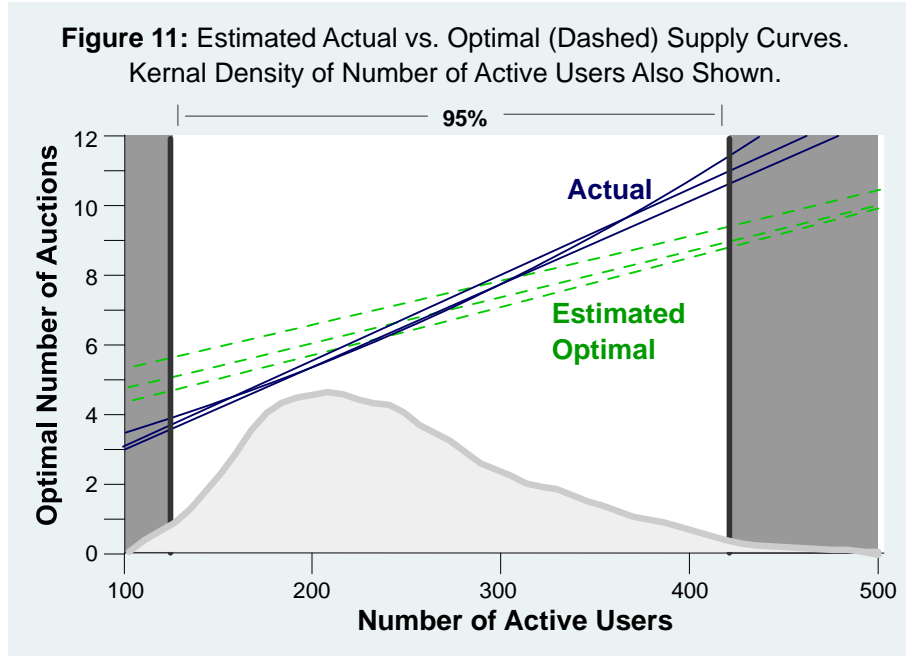
As a first step towards understanding the profit function $P(\cdot)$, Figure 10 displays non-parametric regressions of the number of active users on instantaneous profit for multiple equal sized groupings of the number of active auctions. First, note that an increase in the number of active users consistently *increases* the predicted instantaneous profit for each number of active auctions. Second, note that an increase in the number of active auctions consistently *reduces* the predicted instantaneous profit for each number of active users. Finally, note that these effects appear to be largely linear and independent of each other.

With these results in mind, I parametrize the model as

$$y_{jt}^A = \beta_0 + \beta_1 Q_{Dt} + \beta_2 Q_{St} + \varepsilon_{jt} \quad (24)$$

with the results shown in the first column of Table 6. This regression estimates that instantaneous profit rises by €0.088 for each additional hundred active users on the site, but falls by €0.034 for each additional auction. The results of a quadratic regression are shown in the second column of Table 6. The results of a regression with item fixed effects (to control for the heterogeneous effects of the item for auction) are shown in the third column of Table 6. As I show in Section 7.4, the estimated effects are largely unchanged across these regressions for the normal range of active users.

From these results, it is straightforward to determine the optimal number of active auc-



tions for a given number of active users. For example, using the first specification, the predicted instantaneous profit $\hat{\pi}^A$ to the auctioneer from running Q_S auctions given Q_D users is

$$\hat{\pi}^A(Q_S, Q_D) = Q_S(\hat{\beta}_0 + \hat{\beta}_1 Q_D + \hat{\beta}_2 Q_S) \quad (25)$$

Consider a situation in which there are 200 active users. Using the estimates from the first specification, the predicted profit from running one auction is €.368, from running two auctions is $2*(\text{€}.368 - \text{€}.034) = \text{€}.669$, from running three auctions is $3*(\text{€}.368 - \text{€}.068) = \text{€}.903$, and so on. It is easy to show that, given the estimates from the first specification, $\hat{\pi}^A(Q_S, Q_D)$ is a strictly concave in Q_S with a unique maximum $Q_S^*(Q_D) = \frac{\hat{\beta}_0 + \hat{\beta}_1 Q_D}{-2*\hat{\beta}_2} = \frac{.225 + 0.000883 Q_D}{2*-.0335}$. Solving this equation when $Q_D = 200$ leads to an estimated optimal supply of $Q_S^*(Q_D) \approx 6$ auctions.

Using the same logic for the other specifications, it is easy to calculate the optimal (short-term) supply curve for each specification. These curves are compared with the actual supply functions in the following section.

7.4 Comparison between Actual and Optimal Supply Rules

In the previous two sections, I estimated the actual and optimal short-term supply curve of auctions. Figure 11 displays these curves for each specification along with a kernel density estimation of the number of active users on the site. Clearly, the different specifications for both curves produce extremely similar qualitative results.

It is possible to compare the supply curves quantitatively by comparing the estimated profits from each curve. Specifically, given the estimated empirical distribution of users $f(Q_D)$ and the estimated profit function $\hat{\pi}^A(Q_S, Q_D)$, the estimated profit from following supply curve $Q_S(Q_D)$ is

$$\hat{\Pi} = \int f(Q_D) \hat{\pi}^A(Q_S(Q_D), Q_D) dQ_D$$

Using estimates from the first specifications of both models, *the optimal supply curve yields expected instantaneous profits of €1.48, while the estimated supply curve yields €1.46*, suggesting that Swoopo captures 98.6% of potential profits in its supply curve.

This section has produced a variety of results, with a three important points. First, it appears that Swoopo is efficient at profit-maximization with respect to its supply curve. Second, all else equal, the auctioneer's profit is increased by additional users and reduced by additional auctions. Therefore, for a given number of users, supplying the optimal number of auctions is important. Third, the auctions require a reasonable number of consistent active users (over 40, based on my measure) to run successfully.

7.5 Competitor Profits

As Swoopo makes a large amount of profit for running penny auctions, why would other companies not begin to offer these auctions? Swoopo holds no intellectual property and the cost of creating a nearly identical product is extremely cheap.⁴³ Furthermore, as bidders would presumably prefer to compete with fewer other bidders (there is a *negative* network externality), entrants could be potentially favored over an established firm. Consistent with this logic, there has been a large influx of competitor firms in this market.⁴⁴

However, the above results concerning the supply rules suggest a potentially important structural barrier to entrants in this market. As there are significant diminishing returns to the supply of auctions, over-supply of auctions for a given number of users can be costly. However, entrants must over-supply auctions in order to attract a larger userbase, leading to large costs until the userbase grows to match the supply.⁴⁵ This conclusion is bolstered by auction-level data compiled from competitor sites. Table 7 displays the (recent) use and profit statistics of Swoopo and five other major entrants to this market.⁴⁶ Only one of the five major competitors is making large daily profits, which are still an extremely small percentage (6.6%) of Swoopo's daily profits. The other four competitors are making small

⁴³In fact, companies that sell pre-designed website templates for these auctions allow a potential competitor to start a similar site in a few hours.

⁴⁴Simply searching google for "penny auction" or "swoopo" will reveal a variety of paid advertisements for other companies running penny auctions.

⁴⁵Another potential reason is switching costs. Although it would appear that switching costs are low, users appear to switch between companies very rarely. Only 24 (of over 129,000) Swoopo usernames in the bid-level dataset appear in the top ten lists of the five largest competitors.

⁴⁶Based on cursory research, these five companies were the top five competitors to Swoopo as of June 2009.

or negative daily profits. This analysis suggests that entrants will not immediately reduce the rents in this market, at least in the medium term.

8 Discussion and Conclusion

This paper presents a variety of theoretical and empirical results concerning the market for *penny auctions*, a relatively new auction format. My first result is that, in aggregate, players significantly overbid in these auctions, leading to large expected profits for the auctioneer. Comparing the empirical hazard function with theoretically predicted function yields my second result: players overbid more and more as the auction continues. I show that this behavior matches the predictions of a model with agents that exhibit a naive sunk cost fallacy. Interestingly, players with higher levels of experience have higher empirical returns from bidding, most notably through the increased use of "aggressive bidding" strategies. For my third main result, I demonstrate that Swoopo nearly optimally matches the expected number of active auctions to the expected number of users on the site at each point in time. The supply rules suggest that an entrant that attempts to over-supply auctions in order to attract a large userbase will incur large short term losses, creating a structural barrier to entry. I show that this conclusion is consistent with findings from data from five competitors, which shows that competitors lag significantly behind the sole market leader in terms of daily profit.

From a policy perspective, the conclusions of the paper raise the question of regulation for this type of auction. On one hand, the auction appears to resemble a lottery, with large numbers of participants losing relatively little, one participant winning a significant prize, and the auctioneer making large profits. This suggests that, to the extent that governments choose to regulate lotteries (which they often do, for moral, paternalistic, or revenue-generating reasons (Clotfelter and Cook(1990))), there is a role for regulation of these auctions. However, there are also some key differences which make the role of government regulation less clear: this auction possesses no exogenous source of randomness (all randomness arises from the strategies of other players); skill does play a role in the expected outcome; and there is no obvious deception or manipulation of the players of the game.

There are multiple avenues for further research in the market for penny auctions. For example, while my results are suggestive that players exhibit a sunk cost fallacy, the format of Swoopo's auction makes it difficult to fully differentiate between this and a small set of other potential explanations. Laboratory study would allow for the controlled variation needed to more precisely identify the appropriate model of players' behavior. Alternatively, it might be possible to achieve this goal empirically by using variation in the rules of the auctions run by new and future entrants. Finally, as the market for penny auctions is still relatively young, time will allow a more comprehensive study of the evolution of individual producer behavior and market dynamics as the market matures.

References

- Alesina, A., and A. Drazen, Why are stabilizations delayed?, *American Economic Review*, 81 (1991), 1170–1188.
- Anderson, S., Friedman, D., Milam, G. and Singh, N. Buy it now: A hybrid market institution, Working Paper (2007).
- Ariely, D. and I. Simonson, Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions, *Journal of Consumer Psychology*, 13 (2003), pp. 113-123.
- Ashraf, N., J. Berry and J. Shapiro. Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia. Harvard Business School Working Paper (2008).
- Bajari, P. and Hortacsu, A.. Economic Insights from Internet auctions, *Journal of Economic Literature*, 42 (2004), 457-486.
- Bishop, D. T., C. Canning, and J. M. Smith, The war of attrition with random rewards, *Journal of Theoretical Biology*, 74 (1978), 377–388.
- Bilodeau, M., and A. Slivinski, Toilet cleaning and department chairing: Volunteering a public service, *Journal of Public Economics*, 59 (1996), 299–308.
- Bliss, C., and B. Nalebuff, Dragon Slaying and Ballroom Dancing: The Private Supply of a Public Good, *Journal of Public Economics* 25 (1984), 1–12.
- Bramsen, J., A Pseudo-Endowment Effect in Internet Auctions, Institute of Food and Resource Economics Working Paper (2008)
- Breslow, N. E. and Day, N. E. *Statistical Methods in Cancer Research. Volume II - The Design and analysis of cohort studies.* Lyon: International Agency for Research on Cancer, 1987, 178-229.
- Bulow, J., and P. Klemperer, The generalized war of attrition, *American Economic Review* 89 (1999) 175–189.
- Card, D., and C. A. Olson, Bargaining Power, Strike Durations, and Wage Outcomes: An Analysis of Strikes in the 1880s, *Journal of Labor Economics*, 13 (1995), 32–61.
- Clotfelter, C. and Cook, P. On the Economics of State Lotteries. *The Journal of Economic Perspectives* 4 (4), 105-119.
- Eichberger, J., and D. Vinogradov, Least Unmatched Price Auctions, Working Paper (2007).
- Eyster, E. Rationalizing the past: A taste for consistency. Nuffield College Mimeograph (2002).
- Fudenberg, D., and J. Tirole, A Theory of Exit in Duopoly, *Econometrica*, 54 (1986), 943–960.
- Ghemawat, P., and B. Nalebuff, Exit, *Rand Journal of Economics* 16 (1985), pp. 184–194.
- Grubb, M., Selling to Overconfident Consumers, Working Paper (2008)

- Hendricks, K., and Porter, R. H., The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts, *American Economic Review*, 86 (1996), 388–407.
- Hendricks, K. and Wilson, C. A., The War of Attrition in Discrete Time." Mimeo, University of British Columbia (1995).
- Heyman, J., Orhun, Y., and Ariely, D., Auction Fever: The Effect of Opponents and Quasi-Endowment on Product Valuations, *Journal of Interactive Marketing*, 18 (2004), 7-21.
- Hinnosaar, Toomas. Penny Auctions. Working Paper - Northwestern University (2009).
- Hoerisch, H., and Horisch, O., Less Fighting than Expected – Experiments with Wars of Attrition and All-Pay Auctions, Working Paper (2008).
- Houba, H. van der Laan, D. and Veldhuizen, D., The Unique-Lowest Sealed-Bid Auction. Working Paper (2008).
- Kennan, J., and R. Wilson, Strategic Bargaining Models and Interpretation of Strike Data, *Journal of Applied Econometrics*, 4 (1989) S87–S130.
- Klein, J. P. and Moeschberger, M. L. *Survival Analysis Techniques fo Censored and Truncated Data* (2nd Edition). New York: Springer-Verlag, 2003, 166-177
- Krishna, V., and J. Morgan, An analysis of the war of attrition and the all-pay auction, *Journal Economic Theory* 72 (1997), 343-362.
- Malmendier, U., and Vigna, S.D., Paying Not to Go to the Gym, *American Economic Review*, Volume 96 (2006), 694-719.
- Milgrom, P., and R. J. Weber, Distributional Strategies for Games with Incomplete Information, *Mathematics of Operations Research*, 10 (1985), 619-632.
- Murnighan, K., A Very Extreme Case of the Dollar Auction, *Journal of Management Education*, 26 (2002), pp. 56-69.
- Platt, B., Price, J., and Tappen, H. Pay-to-Bid Auctions. Working Paper - Brigham Young University (2009).
- Rapoport, A., H. Otsubo, B. Kim, and W.E. Stein, Unique Bid Auctions: Equilibrium Solutions and Experimental Evidence, Working Paper (2007).
- Raviv, Y. and G. Virag, Gambling by auctions. Working Paper (2007).
- Riley, J.G., Strong Evolutionary Equilibrium and the War of Attrition, *Journal of Theoretical Biology* 82 (1980), 383-400.
- Roth, D., Rationalizable predatory pricing, *Journal of Economic Theory*, 68 (1996), 380–396.
- Shubik, M., The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation, *Journal of Conflict Resolution* 15 (1971), 109–111.

Standifird, S., Roelofs, M., and Durham, Y., The Impact of eBay's Buy-It-Now Function on Bidder Behavior, *International Journal of Electronic Commerce*, 9 (2004), 167-176.

Vigna, S.D., Psychology and Economics: Evidence from The Field, *Journal of Economic Literature*, forthcoming (2009).

Wolf, J., Arkes, H., and Muhanna, W., Is Overbidding in Online Auctions the Result of a Pseudo-Endowment Effect, Working Paper (2005).

Myerson, R.B., Population Uncertainty and Poisson Games, *International Journal of Game Theory*, 1998

Zimin, I. N., and J. P. Ivanilov, 1968, A macroeconomic model of a war of attrition, in *Theory of optimal solutions*, vol. 4, pp. 27-41, Kiev. Akad. Nauk Ukrain. SSR.(Russian).

A Appendix

A.1 Robustness of the model

A.1.1 $\text{mod}(y-k,c) \neq 0$

The results in the analytic section relied heavily on the assumption that $\text{mod}(v-k,c)=0$. If this assumption does not hold, there is no equilibrium in which the game continues past period 1. However, as the following proposition shows, strategies that lead to the hazard rates in Proposition 2 form an ϵ equilibrium with ϵ very small and limiting to 0 as the size of time periods shrinks to 0:

Proposition 7 *If $\text{mod}(v - c, k) \neq 0$, there is no equilibrium in which the game continues past period 1. Define $F^* = \max\{t | t < \frac{v-c}{k} - 1\}$. There is an ϵ -perfect equilibrium which yields the same (discrete) hazard rates as those in Proposition 2 with $\epsilon = \frac{1}{n}(1 - \frac{c}{v-F^*k})(v - (F^* + 1)k - c)[\prod_{t=1}^{F^*-1} (1 - \frac{c}{v-tk})]$. There is an contemporaneous ϵ^c -perfect equilibrium (Mailath (2003)) which yields the same (discrete) hazard rates as those in Proposition 2 with $\epsilon^c = \frac{1}{n-1}(1 - \frac{c}{v-F^*k})(v - (F^* + 1)k - c)$. There is a contemporaneous ϵ^c -perfect equilibrium which yields the same hazard and survival rates as those in Proposition 4 with $\epsilon^c \rightarrow 0$ as $\Delta t \rightarrow 0$.*

To give an idea of the magnitude of the mistake of playing this equilibrium in auctions in my dataset, consider an stylized auction constructed to make ϵ as high as possible, with $v = \text{€}9.95$, $c = \text{€}0.50$, $k = \text{€}0.10$, and $n = 20$. In this case, $\epsilon = \text{€}0.000000000224$ and $\epsilon^c = \text{€}0.00060$. That is, even in the most extreme case and using the stronger concept of contemporaneous ϵ^c -perfect equilibrium, players lose extremely little by following the proposed strategies. This is because their only point of profitable deviation is at the end of the game, where their equilibrium strategy is to bet with low probability, there is a small chance that their bet be accepted, and the cost of the bet being accepted is small (and, ex ante, there is an extremely small chance of ever reaching this point of the game).

A.1.2 Independent Values

In the model in the main paper, I assume that players have a common value for the item. The equilibrium is complicated if players have values v_i is drawn independently from some distribution G of finite support before the game begins or $v_i(t)$ is drawn independently from G at each time t . In these equilibria, players' behavior is dependant largely on the exact form of G , with very few clear results about bidding in each individual period (which is confirmed by numerical simulation). However, if players have independent values which tend to a common value, the distribution of hazard rates approaches the bidding hazard rates in the following way:

Proposition 8 *Consider if (1) v_i is drawn independently from G before the game begins or (2) $v_i(t)$ is drawn independently from G at each time t . For any distribution G , there is a unique set of hazard rates $\{\tilde{h}^G(1), \tilde{h}^G(2), \dots, \tilde{h}^G(t)\}$ that occur in equilibrium. Let the support of G_i be $[v - \delta_i, \bar{v} + \delta_i]$. For any sequence of distributions $\{G_1, G_2, \dots\}$ in which $\delta_i \rightarrow 0$ and the game continues past period 1 in equilibrium, $\tilde{h}^G(t) \rightarrow \tilde{h}(t)$ from Proposition 2 for $t > 0$. For any sequence of distributions G with $\delta \rightarrow v$ and $\Delta t \rightarrow 0$, there exist a sequence of corresponding contemporaneous ϵ^c -perfect equilibria with hazard and survival rates equal to those in Proposition 4 in which $\epsilon^c \rightarrow 0$.*

A.1.3 Leader can bid

Throughout the paper, I assume that the leader cannot bid in an auction. This assumption has no effect on the preferred equilibrium below, as the leader not bid in equilibrium even when given the option. However, the assumption does dramatically simplify the exact form of other potential equilibria, as shown below.

Consider a modified game in which the leader can bid. Now, a (Markov) strategy for player i at period t is the probability of betting both if a non-leader ($x_t^{i, NL*}$) and, for $t > 0$, when a leader ($x_t^{i, L*}$) (there is no leader in period 0).

Proposition 9 *In the modified game, Proposition 2 still holds.*

The equilibria in the situation in which non-leaders bid becomes significantly more complicated and finding a closed form solution becomes extremely difficult. For example, consider the equilibrium in which no player bids at period F (this is the equilibrium characterized by $t = F$ and $p = 0$ in Proposition 3). The (numerically) solved equilibrium hazard rates are shown for $v = 10$, $c = .5$, $k = .1$ and $n = 3$ in Figure 1. Notice the obvious irregularities in the equilibrium hazard rates in the modified game. This occurs because the ability of a leader to bid in period t distorts the incentives of non-leaders in previous periods. To see this, consider the situation in which $\tilde{h}(t+1) = 1$ and $\tilde{h}(t) = 0$. When leaders cannot bid, there is no benefit from a non-leader bidding in period $t-1$ as he will not win the object in period t (because the game will continue with certainty) or period $t+1$ (because he cannot bid in period t and therefore will never be a leader at $t+1$), at which point the game will

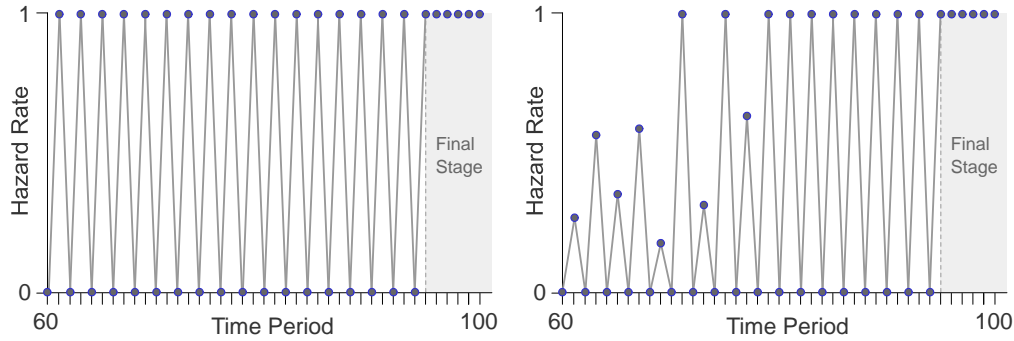


Figure 1: Change in non-bidding equilibria when the leader is allowed to bid.

end. However, when leaders can bid, non-leaders in period $t - 1$ can potentially benefit from bidding. Although there is still no chance that the non-leader in period $t - 1$ will win the object in period t by bidding, she will be able to bid (as a leader) in period t , leading to the possibility that she will win the object in period $t + 1$. Therefore, non-leaders will potentially bid in this situation in equilibrium not to win the object in the following period, but simply to keep the game going for a (potential) win in the future.

A.1.4 Allowing Multiple Bids to Be Accepted

Allowing multiple bids to be accepted significantly complicates the model, especially in a declining-value auction. Consider a player facing other players that are using strictly mixed strategies. If the player bids in period t , there is a probability that anywhere from 0 to $n - 2$ other non-leading players will place bids, leading the game to immediately move to anywhere from period $t + 1$ to period $t + n$. In each of these periods, the net value of the object is different, as is the probability that no player will bid in that period and the auction will be won (which is dependant on the equilibrium strategies in each of the periods).

It is possible to solve the model numerically, leading to a few qualitative statements about the hazard rates. Figure 2 shows the equilibrium hazard rates (with $k = .1, c = .5, n = 10$) given small changes in the value of the good ($v = 10, 10.25, 10.5, 10.75$), as well as the analytical hazard rates from Proposition 2. These graphs demonstrate three main qualitative statements about the relationship between the equilibria in the modified model and the original model:

1. The hazard rates of the modified model are more unstable locally (from period-to-period) than those from Proposition 2, especially in later periods. As n increases, this instability decreases (I do not present graphs for lack of space).
2. The hazard rates of the modified model closely match those from Proposition 2 when smoothed locally.




Figure 2: Numerical Analysis of the hazard rate of auctions for different values (solid lines) when the multiple bids are accepted at each time period vs. the predicted hazard rate (dotted red line) when only one bid is accepted.

3. The hazard rates of the modified model are more stable globally to small changes in parameters in the model. Recall that the hazard rates in Proposition 2 were taken from an equilibrium when $\text{mod}(y - c, k) = 0$. When $\text{mod}(y - c, k) \neq 0$, the hazard rates oscillated radically (although they were smooth in an ε -equilibrium with very small ε). The modified model is much more globally robust to these changes.

A.1.5 Timer

In the model in the paper, unlike that in the real world implementation of the model, there is no timer within each period. Consider a game in which, in each discrete period t , players can choose to place a bid at one sub-time $\tau \in [0, T]$ or not bid for that period. As in the original game, if no players bid, the game ends. If any players bid, one bid is randomly chosen from the set of bids placed at the smallest τ of all bids (the first bids in a period). Now, a player's (Markov) strategy set is a function for each period $\chi_t^i(\tau) : [0, T] \rightarrow [0, 1]$, with $\int_0^T \chi_t^i(\tau) d\tau$ equaling the probability of bidding at some point in that period.

Proposition 10 *For any equilibrium of the modified game, there exists an equilibrium of the original game in which the distribution of the payoffs of each of the players is the same.*

This proposition demonstrates that, while the timer adds complexity to the player's strategy sets, it does not change any of the payoff-relevant outcomes.

A.2 Proofs

Proposition 1

Proof: Assume that an equilibrium exists in which $\tilde{h}(t^*) < 1$ for some $t^* > \frac{v-c}{k} - 1$. There must be some player i with $x_{t^*}^i > 0$. Consider any proper subgame starting at period t^* . For each period in this game, there is some probability $a_t^i \in [0, 1]$ that player i has a bid accepted in that period and some probability q_t that the game ends at that period. As $x_{t^*}^i > 0$, $a_{t^*}^i > 0$ and $q_{t^*} < 1$. Player i 's continuation payoff starting at time t^* is $\sum_{t=t^*}^{\infty} a_t^i(-c + q_{t+1}(v - (t+1)k)) < \sum_{t=t^*}^{\infty} a_t^i(-c + q_{t+1}(v - (\frac{v-c}{k} + 1)k)) < \sum_{t=t^*}^{\infty} a_t^i(-c + q_{t+1}(c - k)) < 0$. But, player i could deviate to never bidding and receive a payoff of 0. Therefore, this can not be an equilibrium.

Proposition 2

Proof: Consider the following symmetric strategies: $x_t^{i*} = \left\{ \begin{array}{ll} 1 & \text{for } t = 0 \\ 1 - \sqrt[n-1]{\frac{c}{v-tk}} & \text{for } 0 < t \leq F \\ 0 & \text{for } t > F \end{array} \right\}$.

Note that: for $t = 0$, $\tilde{h}(t) = 0$; for $0 < t \leq F$, $\tilde{h}(t) = (1 - (1 - \sqrt[n-1]{\frac{c}{v-tk}}))^{n-1} = \frac{c}{v-tk}$; for $t > F$, $\tilde{h}(t) = 1$.

Claim: this is a subgame perfect equilibrium.

First, consider if $k > 0$. I will show that, for each period t , the following statement (referred to as statement 1) is true: there is no strictly profit deviation in period t and the continuation payoff from entering period t as a non-leader is 0. For the subgames starting in periods $t > F$, refer to the proof of Proposition 1 for a proof of the statement. For the subgames starting in period $t \leq F$, the proof continues using (backward) induction with the statement already proved for all periods $t > F$. In period t , player i will receive a continuation payoff of 0 from not betting at time t (she will receive 0 in period t and will enter period $t + 1$ as a non-leader, which has a continuation payoff of 0 by induction). By betting, there is positive probability her bid is accepted. If this is the case, she receives $-c$ in period t and receives $v - (t+1)k$ in $t+1$ with probability $\frac{c}{v-(t+1)k}$ and 0 as a continuation payoff in period $t+2$ with probability $1 - \frac{c}{v-tk}$. If she does not win the object, the continuation payoff from entering period $t+2$ must be 0 by induction. This leads to a total continuation payoff from her bid being accepted of $-c + \frac{c}{v-(t+1)k}(v - (t+1)k) = 0$. If the bid is not accepted, she enters period $t+1$ as a non-leader and must receive a continuation payoff of 0 by induction. Therefore, the continuation payoff from betting must be 0. Therefore, the statement 1 is true for all periods and this is a subgame perfect equilibrium.

Proposition 3

Proof: The proof is two steps.

(1) Each of the proposed set of hazard rates can occur in equilibrium. This closely follows the proof of Proposition 2.

Consider the following symmetric strategies:

$$x_t^{i*} = \left\{ \begin{array}{ll} 1 & \text{for } t < t^* \text{ with } \text{mod}(t + GT + \text{mod}(t^*, 2), 2) = 0 \\ 0 & \text{for } t < t^* \text{ with } \text{mod}(t + GT + \text{mod}(t^*, 2), 2) = 1 \\ 1 - \sqrt[n]{p} & \text{for } t = t^* \\ 1 - \sqrt[n]{\frac{c}{v-tk}} & \text{for } t^* < t \leq F \\ 0 & \text{for } t > F \end{array} \right\}.$$

Note that, as in the proof of Proposition 2. these strategies yield the hazard rates listed in the proposition.

Claim: This is a subgame perfect equilibrium.

First, consider if $k > 0$. The proof of Proposition 2 demonstrates that there is no strictly positive deviation from the strategies for any subgame starting in any period $t > t^*$. Consider period t^* . From the proof of Proposition 2, the continuation payoff from betting and not betting at t^* are the same (0) for all subgames and therefore there is no strictly profitable deviation from playing $1 - \sqrt[n]{p}$ in period t^* in any subgame. Now, consider period $t^* - 1$. Following the same logic as Proposition 2, the continuation payoff from betting and having a bid accepted is $\pi_A = -c + p(v - (t^*)k)$, and the continuation payoff from betting and not having the bid accepted is $\pi_{NA} = 0$. The continuation payoff from betting is a $\pi = \alpha\pi_A + (1 - \alpha)\pi_{NA}$, with the probability of having a bid accepted $\alpha \in (0, 1]$. Simple algebra shows that $p \stackrel{\leq}{\geq} \frac{c}{v-tk} \Leftrightarrow \pi \stackrel{\leq}{\geq} 0$. If $\pi > 0$ ($\pi < 0$), then players must strictly prefer to bid (not bid) in period $t^* - 1$, leaving no profitable deviation (note that if $\pi = 0$, players also do not strictly prefer to deviate). Similar logic shows that players must strictly prefer to not bid (bid) in period $t^* - 2$, bid (not bid) in period $t^* - 3$, and so on. This leaves no profitable deviation in any period and therefore the above strategies are a subgame perfect equilibrium.

(2) No other set of hazard rates can occur in an equilibrium.

Suppose there exists a different set of hazard rates $\hat{h}(t)$ that occur in some equilibrium which can not be characterized this set of hazard rates for some p and t^* . Proposition 1 shows that the hazard rates in this supposed equilibrium must be follow this characterization in periods $t > F$. Note that the hazard rates in Proposition 2 with any hazard rate p in period 0 can be characterized by p and $t^* = 0$. Therefore, there must be some $0 < t \leq F$ such that $\hat{h}(t) \neq \frac{c}{v-tk}$. Define $t^* = \max(t | \hat{h}(t) \neq \frac{c}{v-tk}, t \leq F)$ and $p = \hat{h}(t^*)$. There must be some $t < t^*$, in which $\hat{h}(t)$ differs from the listed set of hazard rates. However, by the proof above, players in periods $t < t^*$ strictly prefer to play the strategies defined above, leading to the a member of the listed set hazard rates, which is a contradiction.

Corollary 1

Proof: By Proposition 3, the hazard rates in this equilibrium must be a member of the listed set of hazard rates and characterized by some p and t^* . If $t^* > 1$, then the hazard rate in this equilibrium must be 1 in either period 1 or 0. Therefore, if an equilibrium continues past period 1, it must be that $t^* \in \{0, 1\}$. For all of these equilibria, the hazard rates match those in Proposition 2 for all $t > 1$.

Corollary 2

Proof:

(1) Claim: the auctioneer's expected revenue is equal to v in the equilibrium in Proposition 2: Note the item must be sold to some bidder as $\tilde{h}(0) = 0$. Therefore, the combined expected profit of the bidders must be v . Next, note that by the proof of Proposition 2, each player must have a continuation payoff of 0 at the start of the auction. Therefore, the auctioneer must have an expected continuation payoff of v .

(2) Claim: For any $\alpha \in [\frac{c}{v}, 1]$, there exists an equilibrium in which the auctioneer's expected payoff is αv . Consider the equilibrium hazard rates characterized by $t^* = 1$ and $p \in [\frac{c}{v-k}, 1]$. Given these hazard rates, $h(0) = 0$.

Proposition 4

Proof: Let $S(t) = p$. As $\tilde{h}(t) = \frac{c\Delta t}{v-(t)k}$ for $t \leq F$, $S(t + \Delta t) = (1 - \frac{c\Delta t}{v-(t+\Delta t)k})p$. Therefore, $S(t) = \lim_{\Delta t \rightarrow 0} \frac{S(t) - S(t+\Delta t)}{\Delta t \cdot S(t)} = \lim_{\Delta t \rightarrow 0} \frac{p(\frac{c\Delta t}{v-(t+\Delta t)k})}{\Delta t \cdot p} = \lim_{\Delta t \rightarrow 0} \frac{c}{v-(t+\Delta t)k} = \frac{c}{v-tk}$. As $H(t) = \int_0^t \frac{c}{v-tk} d\tilde{t}$, $H(t) = \frac{c(\ln(v) - \ln(v-tk))}{k}t$ if $k > 0$ and $t \leq F$, while $H(t) = \frac{c}{v}t$ if $k = 0$. Note that $H(t) = \int_0^t \lim_{\Delta t \rightarrow 0} \frac{S(\tilde{t}) - S(\tilde{t} + \Delta t)}{\Delta t \cdot S(\tilde{t})} d\tilde{t} = - \int_0^t \frac{1}{S(\tilde{t})} (\frac{d}{d\tilde{t}} S(\tilde{t})) d\tilde{t} = - \ln S(t)$. Therefore $S(t) = e^{-H(t)}$ and $S(t) = (1 - \frac{tk}{v})^{\frac{c}{k}}$ if $k > 0$ and $t \leq F$, while $S(t) = e^{-(\frac{c}{v})t}$ if $k = 0$.

Note that, as $\tilde{h}(F + \Delta t) = 1$, $\lim_{\Delta t \rightarrow 0} \Pr(T > F) = 0$ and therefore $S(F) = 0$.

Proposition 5

Proof: $S(tv_1; v_1) = (1 - \frac{tv_1y}{v_1})^{\frac{c}{y}} = (1 - ty)^{\frac{c}{y}} = (1 - \frac{tv_2y}{v_2})^{\frac{c}{y}} = S(tv_2; v_2)$ if $k > 0$ and $t \leq F$. $S(tv_1; v_1) = 0 = S(tv_2; v_2)$ if $k > 0$ and $t > F$. $S(tv_1; v_1) = e^{-(\frac{c}{v_1})tv_1} = e^{-ct} = e^{-(\frac{c}{v_2})tv_2} = S(tv_2; v_2)$ if $k = 0$. $S(t\frac{v_1}{c_1}; v_1, c_1) = e^{-(\frac{c_1}{v_1})t\frac{v_1}{c_1}} = e^{-t} = e^{-(\frac{c_2}{v_2})t\frac{v_2}{c_2}} = S(t\frac{v_2}{c_2}; v_2, c_2)$ if $k = 0$.

Proposition 7

Proof: Consider the strategies noted in the proof of Proposition 2 with $F = F^*$. For the standard ϵ -perfect equilibrium, we consider the ex ante benefit of deviating to the most profitable strategy, given that the other players continue to follow this strategy. Following the proof of Proposition 2, it is easy to show that there is no profitable deviation in periods $t > F^*$ and $t < F^*$. Therefore, the only profitable deviation is to not bet in $t = F^*$. This will yield a continuation payoff of 0 from period F^* . The ex ante continuation payoff from

betting is $\epsilon = \frac{1}{n}(1 - \frac{c}{v-F^*k})(v - (F^* + 1)k - c)[\prod_{t=1}^{F^*-1} (1 - \frac{c}{v-tk})]$. (To see this, note that

there is a $\prod_{t=1}^{F^*-1} (1 - \frac{c}{v-tk})$ change that the game reaches period F^* . In period F^* , there is a $(1 - \frac{c}{v-F^*k})$ probability that at least one player bets. As strategies are symmetric, this means that, ex ante, a player has a $\frac{1}{n}(1 - \frac{c}{v-F^*k})$ probability of her bet being accepted in this period, given that the game reaches this period. If the bet is accepted, the player will receive $(v - (F^* + 1)k - c)$. Therefore, the ex ante benefit from deviating to the most profitable strategy is ϵ . For the contemporaneous ϵ^c -perfect equilibrium, we consider the benefit of deviating to the most profitable strategy once period F^* is reached, given that the other players continue to follow this strategy. This is $\epsilon^c = \frac{1}{n-1}(1 - \frac{c}{v-F^*k})(v - (F^* + 1)k - c)$ (To see this, note that in period F^* , there is a $(1 - \frac{c}{v-F^*k})$ probability that at least one

player bets. As strategies are symmetric, this means that, ex ante, a non-leader has a $\frac{1}{n-1}(1 - \frac{c}{v-F^*k})$ probability of her bet being accepted in this period (as there are only $n-1$ non-leaders). If the bet is accepted, the player will receive $(v - (F^* + 1)k - c)$.

Proposition 8

Proof: In case 1, I will refer to $v_i(t) = v_i$. The proof is simple (backward) induction on the statement that there is a unique hazard rate that can occur in each period in equilibrium. By the same logic in the proof to Proposition 1, $\tilde{h}(t) = 1$ for all $t > \frac{v+\delta_i-c}{k} - 1$. Consider periods $t \leq F^* = \max\{t | t \leq \frac{v+\delta_i-c}{k} - 1\}$ where $\tilde{h}(t+1)$ is unique in equilibrium by induction. If $\tilde{h}(t+1) = 0$, then $\tilde{h}(t) = 1$ as any player with finite $v_i(t)$ strictly prefers to not bid. If $\tilde{h}(t+1) > 0$, a player with cutoff type $v^*(t) = \frac{c}{\tilde{h}(t+1)} + (t+1)k$ is indifferent to betting at time t given $\tilde{h}(t+1)$.

Therefore, $\tilde{h}(t) = G(\max(\min(v^*, v + \delta), v - \delta))$ and the statement is true.

Suppose G_i is such that the game continues past period 1.

Claim 1: If $\delta < k$, then (1) $\tilde{h}(t) = 0$ for $t \leq F \Rightarrow \tilde{h}(t-1) = 1$ and (2) $\tilde{h}(t) = 1$ for $t \leq F \Rightarrow \tilde{h}(t-1) = 0$.

>(1) This is true as a bidding leads to $-c$, a lower payoff than not bidding.

>(2) If $\tilde{h}(t) = 1$, then the payoff of bidding for a player with value \tilde{v} at period $t-1$ is $\Pr[\text{Bid Accepted}](\tilde{v} - tk - c)$. Note that $\Pr[\text{Bid Accepted}] > 0$ if a player bids. Note that $t \leq F \Rightarrow t \leq \frac{v-c}{k} - 1 \Rightarrow 0 \leq v - c - (t+1)k \Rightarrow 0 < v - \delta - c - tk$ as $\delta < v$. Therefore, $0 < \Pr[\text{Bid Accepted}](\tilde{v} - (t+1)k - c)$ for every $\tilde{v} \in [v - \delta, \bar{v} + \delta]$ and therefore $\tilde{h}(t) = 0$ and the claim is proved.

Claim 2: If $\delta < k$, $\tilde{h}(t) \in (0, 1)$ for every $0 < t \leq F$. Suppose that $\tilde{h}(t) = \{0, 1\}$ for some $0 < t \leq F$. If $\tilde{h}(1) = 1$, then game ends at period 1, leading to a contradiction. If $\tilde{h}(1) = 0$, then $\tilde{h}(0) = 1$, and game ends at period 0, leading to a contradiction. If $\tilde{h}(t) = 1$ (alt: 0) for $0 < t \leq F$, then $\tilde{h}(t-1) = 0$ (alt: 1), $\tilde{h}(t-2) = 1$ (alt: 0) by claim 1. But, then $\tilde{h}(1) = \{0, 1\}$, which leads to a contradiction as above.

Claim 3: By the same logic in the proof to Proposition 1, $\tilde{h}(t) = 1$ for all $t > \frac{v+\delta_i-c}{k} - 1$. Therefore, $\tilde{h}^G(t) = 1$ for $t > \frac{v-c}{k} - 1 = F$ as $\delta_i \rightarrow 0$. For $0 < t \leq F$, note that for some i^* , $\delta_i < k$ for all $i > i^*$ and therefore claim 1 holds for all $i > i^*$. If claim 1 holds, $\tilde{h}(t-1) \in (0, 1)$ implies a cutoff value $v^*(t) \in (v - \delta, v + \delta)$ from above, which by the definition of $v^*(t)$ implies that $\tilde{h}(t) \in (\frac{c}{v-\delta-(t+1)k}, \frac{c}{v+\delta-(t+1)k})$, and therefore $\tilde{h}^{G_i}(t) \rightarrow \frac{c}{v-tk}$ for periods $0 < t \leq F$ as $\delta_i \rightarrow 0$. Therefore, $\tilde{h}^G(t) \rightarrow \tilde{h}(t)$ from Proposition 2 for $t > 0$.

Proposition 9

Proof: Set $x_t^{i,NL^*} = x_t^{i^*}$ from the proof of Proposition 2 and set $x_t^{i,L^*} = 0$ for all i and all t . Note that, as in the proof of Proposition 2, these strategies yield the hazard rates listed in the Proposition 2. The same proof for Proposition 2 shows that, if strategies are followed, the continuation payoff from entering period t as a non-leader is 0 and there is no profitable deviation for a non-leader. Now, consider if there is a profitable deviation for a

leader. For the subgames starting in periods $t > F$, refer to the proof of Proposition 1 for a proof that there is no profitable deviation for a leader in these periods. For the subgames starting in period $0 < t \leq F$, the proof continues using (backward) induction with the lack of profitable deviation already proved for all periods $t > F$. In period t , by not bidding, the leading player will receive v with probability $\frac{c}{v-tk}$ (with the game ending) and 0 as a continuation probability as a non-leader in period $t + 1$ with probability $1 - \frac{c}{v-tk}$, yielding an expected payoff of $v(\frac{c}{v-tk}) > 0$. By bidding, the game will continue to period $t + 1$ with certainty, with some positive probability that her bid is accepted. If her bid is accepted, she receives $-c$ in period t and receives $v - (t + 1)k$ in $t + 1$ with probability $\frac{c}{v-(t+1)k}$ and 0 as a continuation probability as a non-leader in period $t + 2$ with probability $1 - \frac{c}{v-tk}$, leading to a continuation payoff of $-c + (v - (t + 1)k)(\frac{c}{v-tk}) = 0$. If her bid is not accepted, she will receive a continuation probability of 0 as a non-leader in period $t + 1$. Therefore, the payoff from not bidding in period t is strictly higher than the payoff from bidding.

Proposition 10

Proof: Consider a vector of bidding probabilities $x = [x^1, x^2, \dots, x^n] \in [0, 1]^n = X$ in some period. Let $\Psi : X \rightarrow \Delta^n$ be a function that maps x_t into a vector of probabilities of each player's bid being accepted, which I will denote $a = [a^1, a^2, \dots, a^n]$. Claim: For any $a^* \in \Delta^n$, $\exists x \in X$ such that $\Psi(x) = a^*$.

Consider the following sequence of betting probabilities, indexed by $j = \{1, 2, 3, \dots\}$. Let $x^i(1) = 0$. Define $a(j) = \Psi(x(j)) = \Psi([x^1(j), x^2(j), \dots, x^n(j)])$. Define $\tilde{a}_i(j) = \Psi([x^1(j - 1), x^2(j - 1), \dots, x^i(j), \dots, x^n(j - 1)])$ and let $x^i(j)$ be chosen such that $\tilde{a}_i^i(j) = a^{i*}$. Claim: $x(j)$ exists, is unique, $x(j - 1) \leq x(j)$ and $a_i(j) \leq a^*$ for all j . This is a proof by induction, starting with $t = 2$. As $x(1) = 0$, $x(2) = a^*$ by the definition of $\tilde{a}_i(j)$. Therefore, $x(2)$ exists, is unique, $x(1) \leq x(2)$ and $a_i(2) \leq a^*$ as $\frac{\partial \Psi_i}{\partial x^k} < 0$ for $k \neq i$. Now, consider $x^i(j)$. Note (1) $x^i(j) = 0 \Rightarrow \tilde{a}_i^i(j) = 0$, (2) $x^i(j) = 1 \Rightarrow \tilde{a}_i^i(j) \geq 1 - \sum_{k \neq i} a^k(j - 1) \geq 1 - \sum_{k \neq i} a^{k*} \geq a^{i*}$ where $1 - \sum_{k \neq i} a^k(j - 1) \geq 1 - \sum_{k \neq i} a^{k*}$ follows by $a_i(j - 1) \leq a^*$, which follows by induction (3) $\tilde{a}_i(j)$ is continuous in $x^i(j)$ and $\frac{\partial \Psi_i}{\partial x^i} > 0$. Therefore, there is a unique solution $x^i(j)$ such that $\tilde{a}_i^i(j) = a^{i*}$. As $a_i(j - 1) \leq a^*$ by induction, it must be that $x_i(j) \geq x_i(j - 1)$ as $\frac{\partial \Psi_i}{\partial x^i} > 0$. Finally, note that if $\tilde{a}_i^i(j) = a^{i*}$, then as $a_i^i(j) \leq a^{i*}$ as $\frac{\partial \Psi_i}{\partial x^k} < 0$ for $k \neq i$ and $x_k(j) \geq x_k(j - 1)$ for $k \neq i$.

Set $x^* = \lim_{j \rightarrow \infty} x(j)$. Claim: x^* exists and $\Psi(x^*) = a^*$. First, $\lim_{j \rightarrow \infty} x(j)$ must exist as $x^i(j)$ is bounded above by 1 and weakly increasing. Next, note that $\sum_i x^{i*}$ must also exist with $\sum_i x^{i*} \leq n$. Now, suppose that $\Psi(x^*) \neq a^*$. Then, as $\Psi(x(j)) = a(j) \leq a^*$ for all j , $\Psi(x^*) \leq a^*$ and there must be some i such that $\Psi^i(x^*) - a^{i*} = z > 0$. Choose L such $|\sum_i x^{i*} - \sum_i x^i(j)| < \frac{z}{2}$ for all $j > L$. By the definition of $x^i(j + 1)$, it must be that $x^i(j + 1) \geq x^i(j) + z$. But, as $x(j + 1) \geq x(j)$, then $\sum_i x^i(j + 1) \geq \sum_i x^i(j) + z \geq \sum_i x^{i*} + \frac{z}{2}$, which is a contradiction of L . Therefore, the claim is proved.

Table 1: Descriptive Statistics of the Auction-Level and Bid-Level Datasets

Auction-Level Data	Number of Observations	Mean	Standard Deviation	Fifth Percentile	Ninety-Fifth Percentile
<i>Auction Characteristics</i>					
Worth Up To Value	166,379	287.62	380.90	25.00	1099.00
Adjusted Value (Sec. 3.3)	166,379	252.29	353.58	18.95	999.00
Nailbiter Auction	166,379	.039	-	-	-
Beginner Auction	166,379	.095	-	-	-
<i>Bidding Increment</i>					
€0.00	166,379	.176	-	-	-
€0.01	166,379	.064	-	-	-
€0.10	166,379	.759	-	-	-
<i>Types of Good</i>					
Consumer Goods	166,379	.883	-	-	-
Bid Vouchers	166,379	.100	-	-	-
Cash	166,379	.013	-	-	-
<hr/>					
Bid-Level Data (on subset of Auctions above)	Number of Observations	Mean	Standard Deviation	Fifth Percentile	Ninety-Fifth Percentile
<i>Auction Characteristics</i>					
Worth Up To Value	17,951	244.50	310.03	25.35	978.07
Adjusted Value (Sec. 3.3)	17,951	225.20	292.92	19.32	963.77
Nailbiter Auction	17,951	.292	-	-	-
Beginner Auction	17,951	.001	-	-	-
Number Unique Bidders	17,951	53.81	90.01	5	218
<i>Bid Characteristics</i>					
Used BidButler	13,315,570	.625	-	-	-
Timer < 20 seconds	13,315,570	.634	-	-	-
<i>User Characteristics</i>					
Number of Bids	129,329	102.96	592.35	1	285
Number of Auctions	129,329	7.47	16.37	1	23
Number of Wins	129,329	.138	1.05	0	0

Note: Categories represented by dummy variables.

For example, mean(nailbiter auction)=.039 implies 3.9% of all auctions are nailbiter auctions

Table 2: Descriptive Statistics of Profit

	Number of Observations	Average Profit	Average Profit Margin
<i>All</i>	166,379	€112.68	52.20%
<i>Website</i>			
Post-Web	119,617	€129.56	66.20%
Pre-Web (Phone)	46,762	€69.50	16.40%
<i>Bidding Increment</i>			
€10	123,693	€38.64	28.41%
€01	7,861	€729.38	102.92%
€00	15,528	€224.38	35.65%
<i>Types of Prizes</i>			
Consumer	147,589	€95.18	33.21%
Bid Vouchers	16,603	€250.18	214.16%
Cash Voucher	2,187	€431.90	104.06%

Table 3: Regressions of profit on experience

	(1) OLS Instant. Profit	(2) FE Instant. Profit	(3) FE + slope FE Instant. Profit
Experience (1000)	0.0300*** (14.07)	0.0107*** (3.27)	0.0355*** (8.52)
Experience (sq)	-0.000474 *** (7.51)	-0.0000421 (.55)	0.00146 *** (8.78)
User FE	NO	YES	YES
Max Exp * Exp			-0.00264 *** (9.59)
Constant	-.328*** (75.37)	-.291*** (27.88)	2.215*** (119.02)
Observations	13,315,570	13,315,570	13,315,570

t statistics in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
(Errors clustered on users when no user FEs)

Table 4: Regressions of profit on experience, controlling for strategies

	(1) No strategy controls Instant Profit	(2) Bid on certain items Instant Profit	(3) Bid at certain times Instant Profit	(4) Bid "aggressively" Instant Profit
Experience (1000)	0.0296*** (6.99)	0.0261*** (6.03)	0.0298*** (7.02)	0.0105*** (2.86)
Experience (sq)	-.000546*** (4.55)	-.000444*** (3.82)	-.000551*** (4.63)	-.000222*** (2.80)
Item FE	NO	YES	NO	NO
Time FE	NO	NO	YES	NO
Streak (1-5)				0.057*** (3.50)
Streak (5-20)				0.122*** (5.54)
Streak (>20)				0.254*** (3.43)
Seconds (0-1)				0.581*** (17.41)
Seconds (1-5)				0.203*** (6.69)
Seconds (5-20)				0.097*** (3.43)
Observations	1,265,108	1,265,108	1,265,108	1,265,108

Note: t statistics in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: (Second stage) IV Regressions of desired supply on (predicted) demand

	(1) [10 seconds] Auctions	(2) [one hour] Auctions	(3) [quadratic] Auctions
Predicted Num Users	0.0238*** (776.68)	0.0246*** (40.03)	0.00982*** (62.90)
Predicted Num Users (sq)			0.000029*** (91.06)
Constant	.638*** (90.14)	.659*** (4.63)	2.215*** (119.02)
Observations	595718	1655	595718

Note: t statistics in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Statistics calculated using robust standard errors

Table 6: Regressions of instantaneous profit on number of users and auctions

	(1) OLS Instprofit	(2) Quadratic Instprofit	(3) OLS with FE Instprofit
Auctions	-0.0335*** (-20.13)	-0.0335*** (-20.08)	-0.0359*** (-20.72)
Users	0.000883*** (15.68)	0.000782** (3.07)	0.00101*** (16.74)
Users (sq)		1.74·10 ⁻⁷ (0.41)	
Item Fixed Effects	NO	NO	YES
Constant	0.225*** (17.42)	0.238*** (6.77)	0.206*** (15.03)
Observations	10,157,214	10,157,214	10,157,214

Note: t statistics in parentheses * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$
 Statistics calculated using robust standard errors

Table 7: Descriptive Profit Statistics of competition (from October 2008)

Company	Active Since	Auctions Per Day	Profit Per Day	Average Profit Margin
Swoopo	10/2005	271.77	€42,215.98	62.74%
BidStick	10/2008	38.22	€2,812.60	51.76%
RockyBid	03/2009	9.98	€-483.63	-11.9%
GoBid	02/2009	9.53	€-146.79	-0.13%
Zoozle	02/2009	6.64	€126.36	3.31%
BidRay	04/2009	1.75	€55.82	62.31%

Note: Statistics from Oct 2008-June 2009