

The Sunk-Cost Fallacy in Penny Auctions

Ned Augenblick*

July 2015

Abstract

This paper theoretically and empirically analyzes behavior in penny auctions, a relatively new auction mechanism. As in the dollar auction or war-of-attrition, players in penny auctions commit higher non-refundable costs as the auction continues and only win if all other players stop bidding. I first show that, in any equilibria that does not end immediately, players bid probabilistically such that the expected profit from every bid is zero. Then, using two large datasets covering 166,000 auctions, I calculate that average profit margins actually exceed 50%. To explain this deviation, I incorporate a sunk cost fallacy into the theoretical model to generate a set of predictions about hazard rates and player behavior, which I confirm empirically. While players do (slowly) learn to correct this bias and there are few obvious barriers to competition, activity in the market is rising and concentration remains relatively high.

Keywords: Internet Auctions, Market Design, Sunk Costs

JEL Classification Numbers: D44, D03, D22

*Address: Haas School of Business, 545 Student Services #1900, Berkeley, CA 94720-1900 ned@haas.berkeley.edu. This paper was previously circulated with the title "Consumer and Producer Behavior in the Market for Penny Auctions: A Theoretical and Empirical Analysis". 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, Ernesto Dal Bo, Stefano DellaVigna, Jakub Kastl, Fuhito Kojima, Carlos Lever, David Levine, Scott Nicholson, Monika Piazzesi, Matthew Rabin, and Annika Todd for helpful comments. Thanks to seminar participants at the Stanford Department of Economics, UC Anderson School of Management, UC Berkeley Haas School of Business, University of Chicago Booth School of Business, Northwestern Kellogg School of Management, Harvard Business School, MIT Sloan School of Management, Brown University Department of Economics, UC Davis Department of Economics, the London of Economics, Royal Holloway University, and SITE for suggestions and 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

A *penny auction* is a relatively new auction mechanism run by multiple online companies. In the simplest form of this dynamic auction, players repeatedly choose to pay a non-refundable fixed *bid cost* (\$0.75 in my empirical dataset) to become the *leader* in the auction, and win a good if no other player chooses to bid within a short period of time. Not surprisingly, theory suggests that the auctioneer's expected revenue should not exceed the value of the good. However, in a dataset of more than 160,000 auctions run by a company over a four year period, I show that average auctioneer profit margins empirically exceed 50%. In an illustrative example, my dataset contains more than 2,000 auctions for direct cash payments, in which the average revenue is 204% of the face value of the prize. This paper theoretically and empirically explores these deviations, as well as analyzing the evolution of the market for these auctions over time.

One potential explanation for high auctioneer profits comes from the dollar auction (Shubik 1971), which shares many characteristics with the penny auction. In the dollar auction, two players sequentially bid slowly escalating amounts to win a dollar bill, but are *both* required to pay their last bid. The dollar auction is known as a "prototypical example" of the *irrational escalation of commitment* (also known as *the sunk cost fallacy*), in which players become less willing to exit a situation as their financial and mental commitments increase, even if these commitments do not increase the probability of success (Camerer and Weber 1999). This suggests that the sunk-cost effect also could be playing a role in penny auctions, as players make similarly escalating financial commitments (in the form of bid costs) as the auction continues.

To better understand if sunk costs are driving high auctioneer profits, I first theoretically analyze the penny auction. Not surprisingly, there are multiple equilibria of this game, including asymmetric equilibria in which the game ends after one bid. However, any equilibrium in which play continues past the second period must be characterized by a unique set of aggregate hazard rates and individual strategies from that point forward. In all of these equilibria, players bid probabilistically such that the expected profit from every bid is zero. This equilibrium is similar to the symmetric equilibrium of the dollar auction (and another similar auction, the war-of-attrition). In each of these games, the players can be seen as playing a lottery every time they place a bid in equilibrium, with the probability of winning determined endogenously by the other players' mixed strategies.

Note that, under this interpretation, there are many reasons that we might expect players to overbid (or bid too often), such as risk-seeking preferences or a simple joy of winning.

To understand how to differentiate these explanations from the sunk cost fallacy, I augment the theoretical model such that a player perceives the value to winning the good as rising in her previously (sunk) bid costs. The core prediction of this alternative model is that bidding behavior, hazard rates, and auctioneer profits will start at the equilibrium level of the standard model, but will deviate farther from the standard model as the auction continues. That is, the crucial difference is not that players are willing to bid in this endogenous lottery, but that this willingness increases over time as sunk costs increase. In this sense, the penny auction (and the dollar auction) presents an ideal place to find the sunk cost fallacy given this constant repetition of a decision accompanied by the slow escalation of sunk costs.

I then turn to the data, which consist of a auction-level dataset of nearly 100 million bids on 166,000 auctions and a bid-level dataset of 13 million bids on 18,000 auctions from more than 129,000 users from a large online auctioneer called Swoopo. As predicted by the theory, the ending time of the auction is highly stochastic, with the auctioneer suffering losses on more than one-half of the auctions. However, as previously noted, revenues are far above theoretical predictions, generating \$26 million in profits over a four-year period. To determine if sunk costs are playing a role, I examine the difference between the empirical and theoretically predicted hazard functions. The hazard rates suggest that auctions end with the probability slightly under that predicted by the standard model in the early stages, but deviate farther and farther below as the auction continues. Economically, this leads to a bidder return of only 18 to 24 cents from each 75-cent bid at later stages of the auction, which suggests that bidders are willing to accept worse expected returns as the auction continues, which fits the predictions of the sunk-cost model.

To control for potential selection issues that may drive this result and provide a calibration of the level of the sunk-cost effect, I regress the outcome of an auction ending at any point on a large set of controls, the value of the good at the time of bidding, and the aggregate amount of sunk bid costs (the total of all players' individual sunk costs). Across all specifications, the coefficient on log aggregate sunk costs is highly significant and is 6-11% the coefficient on log value. In the online appendix, I run a structural estimation to control for the alternative explanations of risk-seeking and joy-of-winning preferences. The analysis similarly suggests that a dollar increase in aggregate sunk costs has an impact equal to 8% of a dollar increase in value, controlling for these alternative hypotheses (both of which are also estimated to play a statistically significant role in behavior). As I roughly estimate an average of 16 active players (and a median of 13 players) in auctions, a back-of-the-envelope calculation assuming equal distribution of sunk costs suggests that a player with an additional dollar of individual sunk costs acts as if the value of the good has been increased by just more than a dollar.

As this estimate is potentially biased upwards due to heterogeneity across auctions, I use the bid-level dataset to control for individual heterogeneity by using individual fixed effects. The majority of users attend multiple auctions (often involving the same item), allowing for the observation of changes in individual behavior given changes in individual sunk costs over auctions and time. I find that the probability that a player leaves an auction is highly significantly decreased as the player's individual sunk costs increase. Depending on the specification, the coefficient on log individual sunk costs is between 50-95% the coefficient on log value. Interestingly, sunk bid costs from other recent auctions play a very small role in players' decisions, suggesting that players are largely focused on the sunk costs in the current auction. Furthermore, experience largely mediates the sunk cost fallacy. That is, the coefficient on sunk costs falls as an individual's experience rises, reaching nearly zero for the players with the highest levels of experience in my dataset. This finding suggests that more experienced players might have higher expected profits from each bid. In fact, there is a significant positive (and concave) relationship between a user's experience and profits, even when controlling for user fixed effects, with the most experienced players collecting slightly positive profits in expectation.

Shifting to the larger market for these auctions, there are two main reasons to believe that large auctioneer profits are not sustainable in the long term. On the demand side, players can learn better strategies or to avoid the auction altogether. On the supply side, as it is extremely cheap to perfectly replicate the market leader's auction site and auctions, many other companies will likely enter the market. In fact, the early market leader and source of my dataset, Swoopo, declared bankruptcy in 2011. However, this event does not appear to be representative of the general trends in the market. Using auction-level data from 2009 for five competitors and Alexa Internet visitor data from 2008-2012 for 115 competitors, I show that market activity has generally increased over time, with the total of all sites reaching 0.01% of global Internet traffic. Furthermore, I show that the Herfindahl index of pageviews over the four-year period has remained above the Department of Justice cutoff for "moderate concentration" and commonly rises above the cutoff for "high concentration." These findings suggest that profits for penny auctions are not dying, at least in the medium term.

While the penny auction is an abstract and simple game, the basic strategic decision - determining when to exit given escalating sunk costs and opponents facing the same decision - is common in the real world. For example, the dollar auction was originally used to model escalating tensions in bargaining between firms or nations. Similarly, the war-of-attrition (WOA), which shares the same basic structure, has been used to model competition between firms (Fudenberg and Tirole 1986), public good games (Bliss and Nalebuff 1984), and political

stabilizations (Alesina and Drazen 1991), as well as being theoretically explored extensively (Bulow and Klemperer 1999; Krishna and Morgan 1997). While the game has been studied in the laboratory (Horisch and Kirchkamp 2010), there are only a small number of empirical papers on the WOA, as it is difficult to observe a real-life situation which transparently maps to the game.¹ Therefore, penny auctions provide a large natural experiment that closely mirrors the WOA, suggesting that sunk costs can cause people’s strategies to differ from the predictions of a rational choice model, even with high stakes and over long periods of time.

While the sunk cost fallacy is commonly implicated in a variety of contexts (Thaler 1990), the empirical evidence is relatively thin. Arkes and Blumer (1985) give unexpected price discounts to a randomly selected group of people who are buying season theater tickets, finding that those who pay full price attend more shows than those who receive the discount. Ho, Png, and Reza (2014) find that Singaporeans who pay more for a government license to purchase a car (the price of which varies widely over time) drive the car more. However, Ashraf, Berry, and Shapiro (2010) give unexpected price discounts to a randomly selected group of Zambians who are purchasing a chemical that cleans drinking water and find no effect on the use of the chemical. Experiments on the sunk cost fallacy (Friedman, Pommerenke, Lukose, Milam, and Huberman 2007) have also not found an effect, potentially because it is difficult to exogenously assign a sunk cost to experimental subjects.

The results contribute to the broader understanding of behavioral industrial organization, which studies firm reactions to behavior biases in the marketplace (see DellaVigna (2009) for a survey). The paper also complements a set of three concurrent papers on penny auctions. Hinnsaar (2013) analyzes the auctions theoretically, following a similar approach to this paper. The main difference between the analyses is that my model assumes that if multiple players submit a bid at the same time, only one is counted, while Hinnsaar’s model counts these simultaneous bids in random succession. Hinnsaar’s model leads to similar hazard rates, which imply stochastic end times and no expected profits for the auctioneer, but more complicated individual bidding behavior. Importantly, the major comparative statics are largely the same across the models. Using a subset of Swoopo’s American auction-level data, Platt, Price, and Tappen (2013) demonstrate that a model that incorporates both risk-loving parameters and flexibility in the perceived value of each good cannot be rejected by the observed auction-level ending times. Consequently, they conclude that risk-seeking plays an

¹Furthermore, most situations do not present a *known* bid cost and good value. Empirical studies include Card and Olson (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.

important role in the auctioneer’s profits. In the online appendix, I structurally estimate the sunk-cost model while controlling for the possibility of risk-seeking preferences. While the estimate of the sunk-cost parameter does not change significantly with this addition, I also find that risk preferences play some role in behavior, supporting Platt et al.’s conclusions. Finally, Byers, Mitzenmacher, and Zervas (2010) discuss the use of aggressive strategies and use a non-equilibrium theoretical model to show that misperceptions, such as underpredicting the number of users, can lead to higher-than-zero auctioneer profits. This model is difficult to test empirically, especially as it is difficult to estimate the true number of players who are participating in a given auction at a given time. Note that the misprediction model does make different predictions than the sunk-cost model as long as players’ misperceptions do not change as sunk costs rise.

Multiple working papers have followed this first wave of analysis. On the demand side, Wang and Xu (2011) use bid-level data to further explore bidder learning and exit from the market. Goodman (2012) uses bid-level data to explore bidder reputation using aggressive bidding strategies. Caldara (2012) uses an experiment to determine the effects of group size and timing, finding that timing does not matter but that more participants leads to higher auctioneer profits. On the supply side, Zheng, Goh, and Huang (2011) use a small field experiment to explore the effect of restricting participation of consistent winners, finding that restrictions can increase revenue. Anderson and Odegaard (2011) theoretically analyze a penny auctioneer’s strategy when there is another fixed price sales channel.

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 discusses auctioneer profits, and analyzes empirical hazard rates and individual behavior. The fifth section describes the evolution of market activity and supplier concentration over time. Finally, the sixth section concludes.

2 Auction Description and Theoretical Analysis

2.1 Auction Description

In the introduction, I discuss the simplest version of the penny auction and loosely compare the auction with the dollar auction and the war-of-attrition. This section expands the explanation and comparison.

There are many companies that run penny auctions, which largely follow the same rules

during the time covered by my dataset.² In the auction, multiple players bid for one item. When a player *bids*, she pays a small non-refundable *bid cost* and becomes the *leader* of the auction. The leader wins the auction when a commonly-observable countdown timer hits zero. However, each bid automatically increases the timer by a small amount, allowing the auction to continue as long as players continue to place bids. Therefore, players win when they place a bid and no other player places a bid in the next *period*. To complicate matters slightly, the winner also pays an additional *bid amount* to the auctioneer, which starts at zero and rises by a small commonly-known *bidding increment* with every bid (the bidding increment is commonly a penny, giving rise to the name of the auction). That is, as the auction continues, the *net value* of the good for the player is slowly dropping.

To understand the main differences between this game and the dollar auction (DA) or discrete-time dynamic war-of-attrition (WOA), consider the simplest version of the penny auction in which the bidding increment is set to zero (so that the players only pay bid costs to the auctioneer).^{3,4} As with a penny auction, players in a WOA and DA must pay a non-refundable cost for the game to continue and a player wins the auction when other players decide not to pay this cost. However, in a WOA, players must pay the cost at each bidding stage and are removed if they fail to pay the cost at any point in the auction. In the penny auction, only *one* player pays the bid cost in each bidding stage and players are free to bid as long as the auction is still running. The multi-player DA lies between these two extremes. Players are free to bid in each period regardless of their previous bids, but bidders who return after not bidding are required to repay the costs of the current highest participant (as the new bid must be higher the previous highest bid). Consequently, a player who wins the WOA or DA must have paid the auctioneer the largest amount, while this is not the case in the penny auction.

Another important difference arises when the bidding increment is strictly positive. In this case, the net value of the good is linearly declining as the penny auction continues. In contrast, the net value of the good is constant over time in the WOA and DA. This addition is theoretically troublesome as it destroys the stationarity used to solve the DA and WOA models.

The following section presents a theoretical model of the penny auction and provides an

²As of 2013, allpennyauctions.com held the most comprehensive source of information about penny auction sites and rules.

³In the discrete-time WOA, each active player chooses to bid or not bid at each point in time. All players who 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. The rules are less defined in the situation where all players exit in the same period, which is why the continuous-time version of the game is often used in theoretical papers.

⁴The bidding increment is \$0.00 in 10% of the consumer auctions in my dataset.

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 item.⁵ There is a set of potentially unbounded periods, indexed by $t \in \{0, 1, 2, 3, \dots\}$.⁶ Each period is characterized by a publicly-observable current leader $l_t \in \{0, 1, 2, 3, \dots, n\}$, with $l_0 = 0$. In each period t , bidders simultaneously choose $x_t^i \in \{Bid, Not\ Bid\}$. If any of the bidders bid, one of these bids is randomly *accepted*.⁷ 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 at period t* and the current leader receives the object.⁸

In addition to the bid costs, the winner of the auction must pay a *bid amount*. The bid amount starts at 0 and weakly rises by the *bidding increment* $k \in \mathbb{R}^+$ in each period, so that the bid amount for the good at time t is tk (note that the bid amount and time are deterministically linked). Therefore, at the end of the game, the auctioneer's payoff consists of the final bid amount (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. To match the empirical game, I assume that the current leader of the auction cannot place a bid.⁹ I often refer to the *net value* of the good in period t as $v - tk$. I consequently refer to auctions with

⁵I 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 online 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.

⁶It 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.

⁷In current real-life implementations of this auction, two simultaneous 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." Hinosaar (2013) theoretically analyzes this change (combined with other changes to model) and finds a multiplicity of very complicated equilibria. In the online appendix, I show that the predictions of my model become much more complicated, but remain qualitatively similar when this assumption is changed in isolation.

⁸Note that, unlike the real-world implementation, there is no "timer" that counts down to the end of each bidding round in this model. As discussed in the online appendix, the addition of a 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.

⁹This assumption has no effect on the bidding equilibrium in Proposition 2 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 online appendix.

$k > 0$ as (k) *declining-value auctions* and auctions with $k = 0$ as *constant-value auctions*.

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 or not bid in each individual period regardless of past choices). However, the discreteness of the game requires an additional technical assumption for declining-value auctions that $\text{mod}(v - c, k) = 0$. If this condition is not satisfied, the game unravels and there is no equilibria in which play continues past the first period.¹⁰

For simplicity, I will focus on Markov-Perfect Equilibria.^{11,12} Bidder i 's Markov strategy set consists of a bidding probability for every period given that she is a non-leader $\{p_0^i, p_1^i, p_2^i, \dots, p_t^i, \dots\}$ with $p_t^i \in [0, 1]$. I will commonly make statements about the discrete hazard function, $h(t, l_t) \equiv P[x_t^i = \text{Not Bid for all } i \neq l_t | \text{Reaching period } t \text{ with leader } l_t]$, which is a function that maps each *state* (a period and potential leader) to the probability that the game ends, conditional on the state being reached. Note that $h(0, 0) = \prod_i (1 - p_0^i)$ and $h(t, l_t) = \prod_{i \neq l_t} (1 - p_t^i)$.

Finally, for expositional purposes, I define two measures of profit for the auctioneer throughout the game. To understand these concepts, note that the bidder i at period $t - 1$ is paying the auctioneer a bid cost c in exchange for a probability of $h(t, i)$ of winning the net value of the good $(v - tk)$ at time t . In other words, the auctioneer is selling bidder i a stochastic good with an expected value of $h(t, i)(v - tk)$ for a price c at time t . Therefore, I define the *instantaneous profit* of the auctioneer at time t with leader l_t as $\pi(t, l_t) = c - h(t, l_t)(v - tk)$ and the *instantaneous percent markup* as: $\mu(t, l_t) = \left(\frac{\pi(t, l_t)}{h(t, l_t)(v - tk)}\right) \cdot 100$.

2.3 Equilibrium Analysis

While there are many hazard functions and strategy sets that can occur in equilibrium, I argue that it is appropriate to focus on a particular function and set (identified in Proposition 2) as these must occur in any state that is reached on the equilibrium path after period 1.

To begin the analysis, Proposition 1 notes the relatively obvious fact that no player will

¹⁰I discuss this issue in detail in the online appendix. While the equilibrium in Proposition 2 no longer exists if the condition does not hold, the strategies constitute a contemporaneous ϵ -perfect equilibrium for an extremely small ϵ (on the order of hundredths of pennies) given the observed empirical parameters.

¹¹A Markov-Perfect Equilibrium (Maskin and Tirole 2001) is a refinement of subgame perfection in which players are restricted to condition strategies only on payoff-relevant outcomes. In penny auctions, this removes seemingly-odd equilibria in which players coordinate bidding strategies depending on the identity of the current leader (i.e. player 10 bids if player 1 is the leader, while player 9 bids if player 2 is the leader).

¹²As I show in the online appendix, the statements for hazard rates all hold true when non-Markovian strategies are used.

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 in any history when this time period is reached.

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

If $k > 0$, then in any Markov Perfect Equilibrium, the game never continues past period F . That is, $h(t, l_t) = 1$ if $t > F$.

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. With this constraint in mind, Proposition 2 establishes the existence of an equilibrium in which bidding occurs in each period $t \leq F$:

Proposition 2 There exists a Markov Perfect Equilibrium in which players' strategies, the hazard rate, and auctioneer profits over time are described by

$$(A) p_i^t = \begin{cases} 1 & \text{for } t = 0 \\ 1 - \sqrt[n]{\frac{c}{v-tk}} & \text{for } 0 < t \leq F \text{ for all } i, \\ 0 & \text{for } t > F \end{cases}$$

$$(B) h(t, l_t) = \begin{cases} 0 & t = 0 \\ \frac{c}{v-tk} & \text{for } 0 < t \leq F \text{ for all } l_t, \text{ and} \\ 1 & \text{for } t > F \end{cases}$$

$$(C) \pi(t, l_t) = \begin{cases} 0 & \text{for any } t \text{ for all } l_t. \end{cases}$$

In an equilibrium with this hazard function, players bid symmetrically such that the hazard rate in all histories after time 0 and up to period F is $\frac{c}{v-tk}$. This hazard rate at time period t causes the expected value from bidding (and the auctioneer's profit) in all histories at period $t-1$ to be zero, leading players in these histories to be indifferent between bidding and not bidding. This allows players in $t-1$ to use strictly mixed behavioral strategies such that the hazard rate in all histories in period $t-1$ is $\frac{c}{v-(t-1)k}$, which causes the players in period $t-2$ to be indifferent, and so on. Crucially, in a declining-value auction, there is no positive deviation to players in period F , who are indifferent given that players in period $F+1$ bid with zero probability, (which they must do by Proposition 1).

Note that, in the hazard function in Proposition 2, $h(0, 0) = 0$ is a (arbitrarily) assumption. This choice does not change any of the results in the paper, but simply implies that some bidding always occurs in equilibrium. This is the only choice in which the auctioneer's

expected revenue is v , which might be considered the natural outcome in a common-value auction.¹³

For a constant-value auction, the strategies are equivalent to those in a symmetric discrete-time war-of-attrition (WOA) when n players remain in the game. However, the hazard rate for the WOA is higher as play only continues if more than one player bids, whereas play in a penny auction continues if any player bids.¹⁴

Not surprisingly, there is a continuum of other equilibria in this model. In some of these equilibria, players (correctly) believe that some player will bid with very high probability in period 1 or 2, respectively, which leads them to strictly prefer to not bid in the previous period.¹⁵ Consequently, the auction always ends at period 0 or period 1. Surprisingly, Proposition 3 notes that if we ever observe bidding past period 1, *we must observe the hazard rates in Proposition 2 for all periods following period 1*. If additionally all n players meaningfully participate in the start of the auction (bid with some probability in the initial two periods), *players must be following the individual strategies in Proposition 2 for all periods following period 1*.

Proposition 3 *For declining-value auctions ($k > 0$), in any Markov Perfect Equilibrium*

- (A) *Any observed hazard rate $h(t, l_t)$ follows Proposition 2 for $t > 1$.*
- (B) *Individual strategies p_i^t follow Proposition 2 for $t > 1$ if $p_0^i > 0$ and $p_1^i > 0$ for all i .*

For constant-value auctions ($k = 0$), these statements are true when restricting to symmetric strategies.

To understand the intuition for statement (1) when $k > 0$, consider some period t with $1 < t \leq F$ in which $h(t, l_t) \neq \frac{c}{v-tk}$. As a result of this hazard rate, player l_t must strictly

¹³Furthermore, if the auctioneer values the item at less than v , then he strictly prefers that bidding occurs in period 0, while the bidders are indifferent. If the auctioneer can select the equilibrium (or repeat the auction until some bidder bids in period 0), he would effectively select the particular equilibrium in Proposition 2.

¹⁴Another intuitive explanation for this difference is that, in equilibrium, the expected total costs (the bid costs of all players) spent in each period must equal the expected total benefit (the hazard rate times the value of the good). In the penny auction, only *one* player ever pays a bid cost at each period, whereas in the WOA, there is a chance that more than one player must pay the bid cost. Therefore, the benefit (determined by the hazard rate) must be higher in the WOA.

¹⁵There do not exist similar asymmetric equilibria in which the auction always ends in period 2 (or later). If this occurred, all non-leaders would strictly prefer to bid in period 1. Therefore, the auction would never end in period 1. But then all bidders in period 0 could never win the auction and would strictly prefer to not bid, causing the auction to never reach period 1.

prefer either to bid or not bid in period $t - 1$. If she prefers not to bid, then (t, l_t) will not be reached in equilibrium (and will never be observed). Alternatively, if she prefers to bid, then it must be that $h(t - 1, l_{t-1}) = 0$ for any $l_{t-1} \neq l_t$, leading all players other than i to strictly prefer to not bid in period $t - 2$. Therefore, player l_t cannot be a non-leader in period $t - 1$ in equilibrium, so (t, l_t) will not be observed in equilibrium. Proposition 3 can also be interpreted as an "instantaneous zero-profit" condition on the equilibrium path. The expected hazard rate $\frac{c}{v-tk}$ leads to zero expected profits. If this condition is violated, players in $t - 1$ or $t - 2$ bid in a way that keeps the state off of the equilibrium path.

Statement (2) requires the additional constraint that each player bids with some probability when $t = 0$ and $t = 1$. The constraint excludes equilibria in which one player effectively leaves the game after period 0 (leading to $n - 1$ players in the game) and in which some player is always the leader in a specific period (allowing her strategy for that period to be off-the-equilibrium path and therefore inconsequential). For intuition as to why strategies must be symmetric, consider the case in which players i and j choose strategies such that $p_t^i \neq p_t^j$ for some $t > 1$. Then, it must be that the players face different hazards as the leader in period t : $h(t|l_t = i) \neq h(t|l_t = j)$, leading one of these hazards to not equal $\frac{c}{v-tk}$, which leads to the issues discussed above.

Finally, note that the statements when $k = 0$ require the additional assumption of symmetric strategies.¹⁶ Unlike in declining-value auctions, there is a somewhat complicated non-symmetric equilibrium in which three players alternatively cycle between bidding with relatively low probability and certainty. While players still expect to make zero profits from each bid over time, the hazard rate oscillates around $\frac{c}{v-tk}$ between periods. I choose not to focus on this type of equilibrium because this behavior requires heavy coordination among players and I do not observe these oscillations empirically. Additionally, in the majority of my auction-level data and all of my bid-level data, $k > 0$.

2.4 The Sunk Cost Fallacy

As I will show in Section 4, the predictions of zero profits from the model above are strongly empirically violated. Therefore, in this section, I preemptively present an alternative model that better matches the patterns in the data. In this simple alternative model, players suffer from a *sunk cost fallacy*, in that they become more willing to bid as their bid costs rise, even though these costs are sunk. This model is a simplified and modified version of

¹⁶There are symmetric Markov equilibria that do not lead to the same hazard rates as those in Proposition 2. For example, consider the equilibrium in which all players always bid in odd (even) periods and never bid in even (odd) periods. In this equilibrium, the game always ends after period 0 (period 1).

the sunk-cost model introduced in Eyster (2002), in which players desire to take present decisions (continuing to invest in a bad project) to justify their past decisions (investing in the project initially).

To capture sunk costs in the most parsimonious and portable way, I simply assume that each player's perception of the value of the good rises as she spends more money on bid costs, capturing an additional benefit from justifying her sunk investment. Specifically, a player i who has placed s_i bids has *sunk costs* $s_i c$ and perceives the value of the good as $v + \theta s_i c$, with $\theta \geq 0$ defined as the *sunk cost parameter*. As this parameter rises, the player's *sunk costs* cause her to bid with a higher likelihood in the auction. If this parameter is zero, the model reverts to the standard model above.

I assume that the player is *naive* about this sunk-cost effect, in the sense that she is unaware that her perception of value might change in the future and that she is unaware that other players do not necessarily share her value perception. Without the first type of naivety, players would be aware that they will bid too much in the future and consequently require a compensating premium to play the game initially, leading to zero profits for the auctioneer (and violating the empirical observations). Without the second type of naivety, each player would have very complicated higher-order beliefs, being personally unaware of her own future changes in value perception, but being aware of other players' changing perceptions and being aware of other players' (correct) beliefs about her own changing perceptions. Furthermore, due to the mechanics of mixed strategy equilibria, each player's (non-Markovian) bidding probability would largely be determined by the sunk costs of *other players* rather than her own sunk costs.¹⁷ With this naivety assumption, a player simply plays the game as if the value of the good matches her perceived value, which includes a portion of her own sunk costs.¹⁸

The sunk costs faced by a player at a specific time t depend on the realizations of the player's own mixing decisions, the mixing decisions of the other players, and the realization of the leader selection process. Define s_i^t as the sunk bids placed by player i at time t for a particular realization of the game. Define \vec{s} as the vector containing all of the player's sunk bids and extend p_i^t to be dependent on s_i^t and $h(t, l_t)$ and $\pi(t, l_t)$ to be dependent on \vec{s} . Given this adjustment, Proposition 4 mirrors Proposition 2:

¹⁷In a mixed strategy equilibrium, each player's probability of bidding in the following period is chosen to make the other players indifferent between bidding in the current period. A player is still affected by her own sunk costs as she will not bid if the current bid amount is above her own perceived value.

¹⁸More specifically, following other examples of naivety in the literature, an equilibrium requires that each player's actions weakly maximizes her own payoff given the player's (potentially mistaken) beliefs of other players' actions. In turn, the player's beliefs about other players' actions must constitute an equilibrium for the game that the player perceives she is playing.

Proposition 4 *With sunk costs, there exists a Markov Perfect Equilibrium in which players' strategies, the hazard rate, and auctioneer profits over time are described by*

$$(A) p_i^t(s_i^t) = \begin{cases} 1 & \text{for } t = 0 \\ 1 - n^{-1} \sqrt{\frac{c}{v-tk+\theta s_i^t c}} & \text{for } 0 < t \leq F \text{ for all } i, \\ 0 & \text{for } t > F \end{cases}$$

$$(B) h(t, l_t, \vec{s}) = \begin{cases} 0 & t = 0 \\ \prod_{i \neq l_t} n^{-1} \sqrt{\frac{c}{v-tk+\theta s_i^t c}} & \text{for } 0 < t \text{ for all } l_t, \text{ and} \end{cases}$$

$$(C) \pi(t, l_t, \vec{s}) = \{ c - h(t, l_t, \vec{s})(v - tk) \text{ for any } t \text{ for all } l_t.$$

If $\theta = 0$, (A), (B) and (C) matches those in Proposition 2.

These formulas depend on the specific distribution of sunk costs across the players in each game. For expositional purposes, consider the simplifying assumption that $s_i^t c = \frac{1}{n} t c$ (that is, sunk costs are distributed equally across players). In this case, $h(t, l_t, \vec{s}) = \frac{c}{v-tk+\theta \frac{1}{n} t c}$. While this formula will likely not be satisfied in an individual realization of the game, it is helpful in understanding the comparative statics of the hazard rate and to provide a rough interpretation of the results when the individual distribution of sunk costs is unknown.

2.5 Summary of Theoretical Predictions

Propositions 2 and 4 predict a variety of comparative statics about the hazard rate of the auction and bidding behavior, both with and without a sunk cost fallacy.

If players do not suffer from a sunk cost fallacy, the hazard rate is $\frac{c}{v-tk}$ and players bid with probability $1 - n^{-1} \sqrt{\frac{c}{v-tk}}$ when $0 < t \leq F$, and the auctioneer's profits remain constant at 0. A few comparative statics are of note. First, none of the parameters affect the auctioneer's instantaneous profits, which remain at zero throughout the auction. Second, for constant-value auctions ($k = 0$), the hazard rate and individual bidding probabilities remain constant throughout the auction. For declining-value auctions ($k > 0$), individuals bid less in the auction as it proceeds (and the net value of the good is falling), leading to a higher hazard rate. This effect is strengthened as the bid increment k rises. Third, as the number of players increases, each player's equilibrium bidding probability drops, but the hazard rate and profits *stay constant*.¹⁹ Intuitively, the specific hazard rate in Proposition

¹⁹In the model, the exact number of players in the auction is common knowledge. More realistically, the number of players could be drawn from a commonly-known distribution. In this case, players will bid such that expected auctioneer profits are still zero. However, when the specific realization of the number of players is low (high), the auctioneer will make negative (positive) profits.

2 can be interpreted as a zero profit condition, which must hold regardless of the number of the players. This is useful empirically, as I cannot directly observe the number of players in the auction-level data. Finally, as the value of the good rises, individuals bid with higher probability and the hazard rate consequently decreases. As a result, auctions with higher values continue longer in expectation.

The final comparative static warrants a short digression. As the empirical data consist of many goods that take many values, the auctions are not predicted to share the same survival rates. This divergence creates a challenge in creating a visual representation of the predicted and empirical hazard rates. However, as I discuss in detail in the online appendix, this problem can be solved by using the concept of *normalized time* $\hat{t} = \frac{t}{v}$. The basic intuition is that, given a constant bidding increment k , an auction with a good of value v is approximately as likely to survive past time t as an auction with a good of value $2v$ surviving past time $2t$, with the relationship approaching equality as the length of periods approaches zero.²⁰ That is, all auctions have approximately the same survival rates in normalized time. As a result, hazard rates in normalized time are approximately the same for these auctions, allowing auctions with different values v to be compared. Note that the use of normalized time does not equalize survival rates across auctions with different bidding increments k , which consequently must be grouped into different visual representations.²¹

When players suffer a sunk cost fallacy, the hazard rate is $\prod_{i \neq l_t} n^{-1} \sqrt{\frac{c}{v + \theta s_i^t c - tk}}$ and players bid with probability $1 - n^{-1} \sqrt{\frac{c}{v + \theta s_i^t c - tk}}$ when $0 < t \leq F$, and the auctioneer's profits are $c - h(t, l_t, \vec{s})(v - tk)$. There are a few important changes in the comparative statics from the standard model. For the hazard rate and bidding probabilities, the effect is most easily seen for a constant-value auction (when the hazard is $\prod_{i \neq l_t} n^{-1} \sqrt{\frac{c}{v + \theta s_i^t c}}$). Rather than remaining constant over time, the hazard rate starts at the point predicted by the standard theory, but falls farther from this baseline as the auction continues. This occurs because individuals start with no sunk costs, but bid with higher probability as their personal sunk costs rise from paying for past accepted bids. This gradual deviation from the standard predictions also occurs in declining-value auctions, although it is possible that bidding probabilities rise due to the effect of the bid amount (which rises over time) outweighing the sunk-cost effect. This ambiguity does not occur when focusing on instantaneous profits (or profit margins),

²⁰For example, the probability that a constant-value auction ($k = 0$) with bid cost $c = 1$ and value $v = 100$ survives to time $t = 50$ is $(1 - \frac{1}{100})^{50} \approx 0.605$, while the corresponding probability with value $v = 200$ is $(1 - \frac{1}{200})^{100} \approx 0.606$. The comparable survival probabilities for these auctions given $k = 1$ are 0.495 and 0.497.

²¹It is less clear how to construct a similar normalized time measure to compare auctions with different bidding increments. Most notably, a constant-value auction (with $k = 0$) has a non-zero survival rate at every period, while the survival rate is always zero after the final stage of a declining-value auction ($k > 0$).

which start at zero but rise as aggregate sunk costs rise, regardless of the bidding increment. Finally, this effect of sunk costs is stronger as the number of players decreases, because the sunk costs become more concentrated in fewer players.²² Although interesting, this prediction is less empirically useful as I do not directly observe the number of players in an auction.

Given these comparative statics, the standard model predicts that both the probability that an individual does *not* place a bid and the aggregate hazard rate will either remain constant or rise as the auction continues, while the sunk cost model implies that these probabilities can fall. Furthermore, the sunk cost model predicts that deviations from the predictions of the standard model about hazard rates will grow larger as the auction continues and individual sunk costs rise. This prediction differentiates the hypothesis from alternative models, most transparently in a constant-value auction. If players have a constant *joy-of-winning* from winning the auction, are risk-seeking, or under-predict the number of players in the auction, they will bid with the *same probability* (above that of the standard model) *throughout the auction*. Therefore, my main empirical results involve regressing the likelihood of the auction ending (or a bidder leaving) on the aggregate (or individual) sunk costs, testing the null hypothesis that the coefficient is zero or positive, with an alternative hypothesis that it is negative. Although these specifications test the reduced form predictions of the sunk-cost model in isolation, a structural estimation in the online appendix confirms the reduced form evidence supporting a sunk cost fallacy when these other hypotheses are taken into account.

3 Data and Background

3.1 Description of Swoopo

Founded in Germany in 2005, Swoopo was the largest and longest-running company that ran penny auctions (five of Swoopo's competitors are discussed later in the paper) in 2010.²³ Swoopo auctioned 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 followed the description in Section 2.1: (1)

²²Consequently, as the number of players drops, the auctioneer's profits increase.

²³After Germany, Swoopo spread to the United Kingdom (December 2007), Spain (May 2008), the United States (August 2008). Nearly every auction was displayed simultaneously across all of these websites, with the current highest bid converted into local currency.

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 ran auctions with bidding increments of \$0.15 (76% of the auctions), \$0.01 (6%), and \$0.00 (18%). The cost of a bid stayed mostly constant at \$0.75, €0.50, and £0.50 in the United States, Europe, and the United Kingdom, respectively.²⁴

In the majority of auctions, Swoopo allowed the use of the *BidButler*, an automated bidding system available to all users. Users could program the BidButler to bid within a specific range of values and the BidButler would automatically place bids for the user when the timer neared zero.²⁵ Certain auctions, called *Nailbiter Auctions* (10% of all auctions, 26% of auctions in 2009), did not allow the use of the BidButler. While the ability to use the BidButler does not obviously change the theoretical predictions, the major regression tables additionally report results when restricting to Nailbiter Auctions.

3.2 Description of Data

I will refer to five datasets in this paper, all collected using algorithms that "scraped" the respective websites: The data for the Swoopo auctions consist of two distinct datasets, one that contains auction-level data for all auctions and another that contains more specific bid-level data for a subset of these auctions. To obtain an accurate estimate of the value of the good, I collected a third dataset on pricing from the Amazon website. Finally, I collected two distinct datasets about Swoopo's competitors for the market analysis, which will be discussed in Section 5:

(1) Swoopo Auction-Level Data: The auction-level dataset contains approximately 166,000 auctions for approximately 9,000 unique goods spanning from September 2005 to June 2009. These data represent more than 108.5 million bids. For each auction, the dataset contains the item for auction, the item's value, the type of auction, the bidding increment, the final (highest) bid, the winning bidder, and the end time. From October 2007, the data also contain the final (highest) 10 bidders for each auction. The summary statistics, many

²⁴A few deviations are of note. From September 2008 to December 2008, the cost of a bid in the United States was briefly raised to \$1.00. More significantly, Swoopo introduced a *Swoopo-It-Now* feature in July 2009 in which a player can use the money spent on bid costs in an auction as credit to buy that item from Swoopo. As this rule dramatically changes the game, all analysis in this paper occurs with data captured before July 2009.

²⁵If two players program a BidButler to run at the same time for the same auction, all the consecutive BidButler bids are placed immediately. Other players do observe that a player used the BidButler, but do not observe the bound set by that player.

of which have been previously referenced, are listed in the top portion of Table 1.

(2) Swoopo Individual Bid-Level Data: The bid-level dataset contains approximately 13.3 million bids placed by 129,000 unique users on 18,000 auctions, and was captured every 2-3 seconds from Swoopo's American website from late February 2009 to early June 2009.²⁶ 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.²⁷ 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 the bottom portion of Table 1.

(3) Amazon Price Data: For each good, Swoopo publishes a visible "worth up to" price, which is essentially the manufacturer's recommended price for the item and is commonly higher than the easily obtainable price of the good. In order to create a more accurate value measure, which I call the "adjusted value," I use the price of the exact same item at Amazon.com and Amazon.de (and use the "worth up to" price if Amazon does not sell the item). Sixty percent of auctions use an item that is sold by Amazon, and the adjusted value is 79% of the "worth up to" price. The correlation between the winning bids and the adjusted value (0.699) is much higher (Fisher p-value<0.0001) than between the winning bids and the "worth up to" value (0.523), suggesting that the adjusted value is a more accurate measure of perceived value.²⁸ As I use the "adjusted price" as the estimate for Swoopo's procurement costs, all profits are *underestimated* as Swoopo's costs might be much less than "adjusted value" as a result of standard supplier discounts.

(4) Competitors' Auction-Level Data: In addition to data about Swoopo, I captured similar auction-level datasets for five of Swoopo's competitors: BidStick, RockyBid, GoBid, Zozle, and BidRay. I will refer to these data briefly when I analyze the market for these auctions.

(5) Competitors' Daily Website Visitor Data: To capture the concentration statistics of the market over time, I collected daily website visitor data to 115 penny auction sites from Alexa Internet, which tracks Internet usage.

²⁶Due 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 on March 18th, capturing an estimated 96% of bids.

²⁷The 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.

²⁸More information about the value measure appears in the online appendix.

4 Empirical Results

The theoretical model makes a variety of clear predictions about bidder behavior in penny auctions. The most basic prediction is that auctioneer revenues will not exceed the easily obtainable value of the good. In this section, I will first show that Swoopo's revenues are, on average, more than 150% of the value of the auctioned good. This aggregate deviation could, of course, be driven by a variety of potential explanations.

As discussed in Section 2.5, a model of the sunk cost fallacy makes a set of unique predictions that differentiate it from other explanations. Essentially, the model predicts that deviations in hazard rates, profits, and individual behavior will become larger as aggregate and individual sunk costs accumulate. To test these predictions, I compare the theoretically-predicted with the empirical-observed hazard rates. Then, I examine how the auction-level hazard rate changes with aggregate sunk costs, using a reduced-form and structural estimation. Finally, I examine how the probability that individual players leave an auction changes as they incur larger sunk costs in that auction.

4.1 Auctioneer Profits

According to the equilibrium analysis above, one would not expect the auction format used by Swoopo to consistently produce more revenue than the easily obtainable price of the auctioned good. The first empirical finding of this paper is that this auction format consistently produces revenue above the market value. Averaging across goods, bidders collectively pay 51% over the adjusted value of the good, producing a conservative average profit of \$159. For the 166,000 auctions that span four years in the dataset, the auctioneer's profit for running the auction is conservatively estimated at 26 million dollars.²⁹ 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 negative for a slight majority of the items. Table 2 breaks down the profits and profit percentages 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 more than 103% and 199%, respectively. Consumer goods, which are potentially overvalued by the adjusted value measure, still lead to an estimated average profit margin of 33%. Given that the other auction-types are rare and dramatically differ in profit margins,

²⁹This 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 suggest that this is a significant source of revenue for Swoopo.

I focus on the auctions for consumer goods for the rest of the paper.³⁰

4.2 Auction-Level Hazard Rate

Recall from Section 2.5 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). Figure 2 displays the smoothed hazard rates over normalized time with 95% confidence intervals along with the hazard functions predicted by the standard model for each increment level.³¹ As noted in Section 2.5, 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 \$0.15 (which represent 93% of the observed auctions), the hazard function is very close to that predicted by equilibrium analysis in the 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. This matches the predictions of the sunk-cost model. Note that the sunk-cost model cannot explain the fact that empirical hazard rates for auctions start lower than the predicted hazard rate (particularly when the bid increment is \$0.01, which represents 7% of the auctions).

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. For this, recall the theoretical prediction that the *instantaneous percent markup* remains zero throughout the auction in the standard model, but rises as the auction continues in the sunk-cost model. Figure 3 displays the markup derived from the hazard rates. For auctions with bidding increments of \$0.15 and \$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 \$0.75 (the bid cost) for a good with an expected value of \$0.18-\$0.24. Therefore, rather than making a uniform profit throughout the auction, the auctioneer is making a large amount of instantaneous profit at the end of the auction.

³⁰The quantitative results are very consistent in auctions for bid packages and cash, as shown in a previous version of this paper.

³¹For this estimation, I used an Epanechnikov kernel and a 10-unit bandwidth, using the method described by Klein and Moeschberger (2003). The graphs are robust to different kernel choices and change as expected with different bandwidths.

To more formally test the alternative model of sunk costs, I run a set of regressions regarding the probability of an auction ending at a given time, while controlling for a variety of auction characteristics. To do this, I expand the auction-level dataset into a larger anonymous bid-level dataset by determining all of the implied bids in the auction. That is, if an auction has a bid increment of \$0.01 and the winning bid amount is \$1.00, there must have been 100 additional failed bids in the auction at bid amounts \$0.00, \$0.01, ..., \$0.99. This leads to a dataset of more than 94 million bids in auctions on consumer goods, which has the same structure as the detailed 13.3 million observation bid-level dataset, except that it does not contain information on the identity of the individual bidders.

With this dataset, I regress the binary variable of an auction ending after each bid time on the log of the aggregate amount of sunk costs incurred at that point, the log of the net value of the good at that point, and a large set of auction characteristic fixed effects (including bid-increment, item value, time-of-day, time-of-year, etc.).³² Columns (1)-(4) of Table 3 present the results of this regression without fixed effects, with fixed effects, limiting to nailbiter auctions, and limiting to the time period captured in the more detailed bid-level dataset. In all of the regressions, standard errors are clustered at the item level. First, note that the coefficient on sunk costs is highly significantly negative in each regression ($p < 0.0001$), capturing the notion that auctions are less likely to end as aggregate sunk costs increase. While the coefficient on the sunk costs appears small (-0.000120 in the first specification), note that the baseline rate of auctions ending at a given point is also very small (0.0145). A more appropriate comparison is the coefficient on net value, which represents the change in the probability that an auction ends given log changes in the net value of the good. For the four regressions, the coefficient on aggregate sunk costs is 6%, 7%, 8%, and 11% the coefficient on net value, respectively.

To understand the rough meaning of these ratios given the theoretical model of sunk costs, consider the model in which sunk costs are distributed across all individuals equally. Given this simplification, every dollar increase in aggregate sunk costs amounts to a $\frac{1}{n}$ dollar increase of individual sunk costs, which leads to individuals to perceive that the value of the good has increased by $\frac{1}{n}\theta$ dollars. The ratios noted above do not precisely correspond to $\frac{1}{n}\theta$ as they represent the relative effect of a increase in *log* dollars. However, in the online appendix, I perform a more comprehensive structural estimation of $\frac{1}{n}\theta$ controlling for joy-of-winning and risk aversion effects and find a very similar estimate of 8%.

Using the more detailed bid-level dataset, I generate a rough average estimate of $n = 16$

³²I use a linear probability model as I will run similar regressions on the individual data and need to accommodate a (very) large number of fixed effects.

active players (with a median of $n = 13$) in an auction at each bid.³³ However, even when n is known, the aggregate estimate does not control for unobserved heterogeneity in player composition, which can drive a selection effect that produces biased estimates. Particularly, imagine that there are some players who always bid too much. The auctions that contain these players will have lower hazard rates than other auctions, which will cause these auctions to be more likely to last longer. Therefore, the estimated hazard rate at later time periods will take only these auctions into account, consequently appearing lower than if we were to observe all auctions reaching that point. To correct for these shortcomings, I turn to the detailed bid-level data.

4.3 Individual Behavior

The detailed bid-level data allow for the observation of each bidder’s identity, which allows for the calculation of individual sunk costs over an auction and the ability to control for individual heterogeneity. Unfortunately, I cannot infer an individual’s bidding probability from the data, as I never can observe if the player *would have* made a bid at each stage if another player bids before her. However, I can observe the probability that the individual exits an auction after each bid, which I call the *psuedo-hazard rate*. Recall that the standard theory predicts that the probability that a player does not bid as the game progresses should stay constant (in constant-value auctions) or rise slightly (in declining-value auctions) if the number of users in the auction stays constant, while the sunk-cost model predicts that this probability will potentially decline (as the player is accumulating sunk costs over the course of the auction).

Figure 4 shows the local polynomial estimation of the pseudo-hazard rate (with 95% confidence intervals), aggregated across individuals. Rather than staying constant or rising, the pseudo-hazard rate declines significantly as the number of bids placed in the auction increases. For example, a player who has placed only a few bids has more than a 10% chance of leaving the auction in the next bid, whereas a player with hundreds of bids has less than a 1% chance of leaving the auction in the next bid.³⁴ The smoothed number of active users

³³To create this estimate, I assume that a player is an active participant in an auction for all of the time between her first and last bid in the auction and take the average number of players active at each bid level (removing bid levels when two or more players’ place many instantaneous bids due to the automatic BidButler). Note that the estimation could be biased downward (as some players might be active but have not yet placed a bid) or upward (as some players might not be active for all of the time between bids). This estimate is relatively sensitive to assumptions: the average including auctions with the BidButler is 22. The average number of players active at a given point in time in the auction (rather than a given bid level) is 10.

³⁴Interestingly, note the spikes at 20,30,50 and 100 bids - Swoopo sells bid packages in these precise amounts. There is a concern that the tendency for bidders to not bid at the end of a bid pack could effect

in the auction at the time of the bid is also included in the figure, in order to demonstrate that a decline in the number of active users is not driving the effect.

As these results are aggregated over multiple players, there is still a concern that heterogeneity across individuals is driving the result. However, given the size of the dataset, it is possible to regress the probability of leaving an auction on the log of the individual amount of sunk costs incurred at that point and the log of the net value of the good at that point, controlling for auction characteristic fixed effects as well as user fixed effects. Columns (1)-(5) of Table 4 present the results of this specification without any fixed effects, without user fixed effects, with all fixed effects, focusing on nailbiter auctions, and with an interaction of a measure of user experience and the sunk costs.³⁵ In all of the regressions, standard errors are double clustered at both the user and item level.

In the regressions without user fixed effects, the coefficient on individual sunk costs is negative (-0.0434 and -0.0445) and strongly significant ($p < 0.0001$). As expected, this coefficient is reduced when adding user fixed effects (to -0.0271), although it is still economically and statistically significant ($p < 0.0001$). The result implies that, as individual sunk costs double, the probability of leaving the auction is reduced by 0.019 ($= .7 \cdot 0.0271$). As with the analysis of the aggregate statistics, it is useful to compare the coefficients on individual sunk costs and the net value. Here, the reaction of bidders to a log increase in sunk costs is nearly 95% of the effect of a log increase in net value. When focusing on nailbiter auctions (Column (4)), in which players cannot place bids using an automated bid proxy and must actively place each bid, this ratio falls to 50%.

In the theoretical model, I assume that sunk costs accumulated in *other auctions* do not affect bidding behavior. However, it is conceivable that people consider all recent sunk bid costs when making bidding decisions. To explore this possibility, I determine the (log) additional number of sunk costs accumulated by a given player in other auctions within different time periods (0-30, 30-60, 60-90, 90-120, 120-300, 300-720, and 720-1,440 minutes) of making a bid in a given auction. I do not report the full results of this specification for space constraints. Adding these variables into the regression in Column (3) does not meaningfully change the coefficient on sunk costs within the auction (from -0.0271 to -0.0274). The sign on the coefficients on the other-auction sunk cost variables is mixed, with the highest-magnitude negative coefficient equalling -0.0014, which is only 5% of the effect

the results. As all bid pack sizes are divisible by ten, it is possible to estimate the set of these end bids for each player and control for this effect. This has virtually no impact on the results.

³⁵It is possible to directly control for the number of users in the auction using a noisy estimate, which does not change the results. However, given the noise in this estimate and the desire for consistency with the aggregate section, I use the same fixed-effects as in the aggregate section.

of a sunk cost accumulated within the current auction.³⁶ This finding suggests that players are narrow-bracketing, in that they largely only consider the sunk costs within the current auction when making decisions.

Finally, the column (5) allows the sunk cost coefficient to vary depending on the experience of the user at the time of the bid (experience is defined as the number of prior bids placed in any auction and is discussed more in the online appendix). This result suggests that the effect has a large magnitude (a coefficient of -0.0803) for inexperienced bidders and reverts to zero as experience increases, so that the players with the highest levels of experience in my dataset (30,000-60,000 bids) have a coefficient of nearly 0. Interestingly, this suggests one reason that more experienced players may do better in these auctions. In fact, as I show in the online appendix, there is a very significant positive (concave) relationship between user experience and user instantaneous profits, even controlling for user fixed effects. Specifically, a player with no experience can expect to lose \$0.60 per each \$0.75 bid, while those with very high experience levels have slightly positive expected payoffs per each bid.

While the aggregate and individual results are consistent with the predictions of the sunk-cost model, it is important to note that the results rely on non-experimental variation of sunk costs. This begs the question: What is driving this variation? In the theoretical model, sunk costs accrue randomly because the equilibrium strategy involves randomization. Reality is likely closer to the incomplete information setting in Harsanyi's purification theorem (Harsanyi 1973), in which players experience small private random shocks to payoffs and only bid if the shock is above some threshold.³⁷ Under this view, bidders with higher values for the good (potentially arising from accumulated sunk costs) have a lower threshold and consequently bid with higher probabilities. There is some suggestive evidence for this interpretation. For example, bidders sometimes take small breaks from bidding, presumably to answer the door or go to the bathroom. After these breaks, bidders are significantly more likely to return to an auction given higher sunk costs.³⁸ Similarly, bidders commonly switch between auctions, presumably due to lack of interest. Bidders are less likely to switch auctions given higher sunk costs in the initial auction and, conditional on having switched to a new auction, are more likely to return to the initial auction given higher sunk costs.³⁹

³⁶The full set of coefficients are, respectively: 0.0006, -0.0012, -0.0009, -0.0009, -0.0014, -0.0010, -0.0015 -0.0011.

³⁷In this environment, shocks could come in the form of outside events (bidders need to go to the bathroom, get back to work, walk the dog, etc.), internal events (such as a transient thought) or in the auction-related events (such as facing a new bidder).

³⁸Defining a break as leaving the auction site for more than 10 minutes, the regression of the likelihood of returning to an auction after a break on the log of individual sunk costs yields a coefficient of 0.0186*** (0.00070), controlling for auction-characteristic and user fixed effects.

³⁹The regression of the likelihood of switching to another auction (within 10 minutes) on the log of

While this evidence on mechanisms for random assignment of sunk costs is suggestive, it does not fully rule out the endogeneity problem. In fact, the sunk cost fallacy has been difficult to identify in empirical settings precisely because it is virtually impossible to observe exogenous assignment of sunk costs in a field setting: Taking on an initial investment is inherently a choice, and people who make initial investments are presumably more likely to make later investments.⁴⁰ My field environment somewhat circumvents this issue, as I can observe the *same user* making different investments for the same items in a relatively clean environment. However, if the same user experiences large changes in value perception over time for the same item (but mistakenly continues to enter the auction when her value is low), the endogeneity problem might still exist.

5 Market Size and Competition

The previous section establishes that penny auctions are highly profitable for the auctioneer, in part due to a naive sunk cost fallacy. There are two reasons to believe that these profits are not sustainable in the long run. First, on the demand side, consumers might learn to either modify their bidding behavior such that they do not lose money or avoid these auctions all together. The last results of 4.3 note that (much) more experience does appear to mediate the sunk cost fallacy and that more experienced players have higher expected profits from each bid. Second, on the supply side, competition might reduce each firm's profits as there are very few obvious barriers to entry in this market. There are no intellectual property barriers and the cost of creating a nearly identical product is extremely cheap. In fact, there are companies that sell pre-designed penny auction website templates that allow any potential competitor to start a similar site in a few hours. This view is supported by the fact that, in March 2011, Swoopo's parent company filed for bankruptcy, shutting down the auction website. Internet forums and articles cite a variety of sources for this event, including competitive forces, poor management, over-hiring, and a disappearing market.

A detailed analysis provides a complex picture. First, consider the supply side. In 2009, four years after Swoopo was founded and more than a year after entering the United States, the market was still highly concentrated. Table 5 displays the use and profit statistics of Swoopo and five other major entrants to this market in 2009.⁴¹ Each company produced

individual sunk costs yields a coefficient of -0.0271*** (0.00057), controlling for auction-characteristic and user fixed effects. Given that a person does switch auctions, the same regression with the likelihood of returning to the original auction yields a coefficient of 0.0505***(0.00074).

⁴⁰With experimental control, one could randomly "reject" bids or randomly assign a bid cost after a bid was placed, leading to certainly-exogenous variation in sunk costs.

⁴¹Based on cursory research, these five companies were the top five competitors to Swoopo as of June

a very small number of auctions in comparison to Swoopo. Furthermore, only one of the five major competitors was making large daily profits, which were still a small percentage (6.6%) of Swoopo's daily profits. The other four competitors were making small or negative daily profits. Although there was a clear opportunity for profits in this industry and it was not difficult to perfectly replicate Swoopo's website, these companies were not particularly successful, at least in the medium-term.

By 2011, there were hundreds of competitor sites. To quantify the structure of the market at this time, I collect visitor data on 115 penny auction websites that were active at some point from 2008-2012. The site list comes from two sources. I use the set of 97 sites that were tracked at some point in time by the largest penny auction tracking service, Allpennyauctions.com. I append 18 sites that operated prior to the tracking service, such as those in Table 5. For each of these sites, I collect visitor data (the daily unique pageviews per million views) from Alexa Internet, a company that tracks visitors to websites.⁴² I then construct a monthly concentration index using the visitor data by creating a Herfindahl index over the average pageviews. The results are shown with the dotted line in Figure 5, which also highlights the point of Swoopo's exit.

In early 2008, when penny auctions are introduced to the United States, there is an extremely high level of concentration (Swoopo was essentially a monopoly). As more competitors enter the market, the level of concentration is reduced. However, the Herfindahl index stays firmly above 0.15 (the Department of Justice cutoff for "moderate concentration") and often rises above 0.25 (the cutoff for "high concentration"). After Swoopo exits, concentration stays above .2 and rises as high as 0.4, with a new site (quibids.com) receiving around one-half of all penny auction traffic over this time.

Figure 5 also plots the total daily number of pageviews per million pageviews for all sites, a metric of activity for the entire market: a level of 100 suggests that the total pageviews of all penny auctions sites accounted for an average of 0.01% of global Internet traffic in that month. Although the growth is not monotonic (including a sharp drop following Swoopo's exit), the number of visitors is generally rising, reaching nearly 0.01% of Internet traffic. This growth is also reflected in the auction and profit statistics. At the peak of my dataset in 2009, Swoopo was running nearly 2,000 auctions a week with an estimated profit of around \$250,000 from selling \$425,000 worth of goods. For comparison, the current market leader

2009.

⁴²One might prefer another measure of usage, such as the number of auctions on the website. Unfortunately, historical data for this statistic is not available. However, auction data from November 2012 was available for 52 sites from Allpennyauctions.com. In November 2012, my measure of use, the combined number of daily unique pageviews per million users, is highly correlated (.987) with the number of auctions for these sites.

(quibids.com) in 2012 runs nearly 17,000 auctions a week with an estimated profit of \$550,000 from selling nearly \$1,250,000 worth of goods. These statistics suggest that, in fact, it does not appear that activity is falling or competition is lowering profits in the medium term.

The fact that concentration remains high in the face of apparently increasing demand implies that there is a significant barrier to competition in this market. However, as discussed above, it is possible to replicate the market leader's technology with very few upfront costs. 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. Finally, while there are presumably search costs and switching costs, these appear relatively small as there are well-known aggregator sites that list all penny auction sites (including reviews and profit statistics) and joining a new site takes a few minutes.

In informal discussions, small penny auction site owners point to a different structural barrier.⁴³ From their perspective, users choose penny auction sites based on the number of active auctions at any given time. While it is technically easy for another company to *perfectly match* the market leader's supply of auctions at every point in time, the auctions will continually end quickly without a large user base, leading to large immediate losses. If these temporary losses are high enough, companies are forced to grow slowly. The Alexa user data suggest relatively slow movements of shifting market power, which provides indirect evidence in support of this view. For example, it took nearly two years for quibids.com to overtake Swoopo in site rankings.

6 Discussion and Conclusion

The penny auction is a relatively new auction format, which is similar to the dollar auction or dynamic war-of-attrition. The auction provides an ideal field environment to test for the existence of the sunk cost fallacy, as it is a simple real-stakes strategic game in which players continuously make a simple binary decision as they slowly commit measurable, non-recoverable, and escalating monetary amounts to win a retail good with a well-known value. Furthermore, because the game is played by the same players many times, it is possible to observe the same individual's behavior given different levels of individual sunk costs, which vary naturally throughout the auction. A theoretical analysis suggests that a player who exhibits the sunk cost fallacy will become less willing to leave the auction as sunk costs accumulate, even though these past expenditures do not improve the player's chances of

⁴³I briefly talked to six penny auction site owners by phone in 2009-2010. This is a non-representative and contaminated sample as all instigated conversations with me after a version of this paper was circulated.

winning. This prediction is confirmed in the empirical analysis of aggregate and individual behavior, a conclusion that is robust to a variety of specifications, including controlling for other potential explanations.

As a result of bidder behavior, auctioneer profit margins from the original penny auction site empirically exceeded 50% across more than 166,000 auctions. This level of profits has led commentators to question the virtue of the auction format, suggesting that it is designed entirely to exploit players with behavioral biases (Thaler 2009; Stone 2009). Given this view, one might expect players to learn either not to play the auction or not to be affected by sunk costs, eliminating auctioneer profits in the long run. While there is evidence of learning in my dataset, auctioneer profits and traffic to other penny auction sites remains high nearly ten years after the introduction of the format, likely due to the ability to attract a constant stream of inexperienced new bidders on the Internet. Therefore, it appears that this business model is oddly robust in the medium term.

Under the interpretation that the auction is taking advantage of consumers, some have suggested the need for state regulation of penny auctions (King 2012). This debate points to a deeper issue in behavioral industrial organization: Should regulation exist to "protect" consumers from their own behavioral biases? In the case of penny auctions, critics have pointed to the similarity between the auction and a lottery, which governments commonly regulate for moral, paternalistic, or revenue-generating reasons (Clotfelter and Cook 1990). While there are similarities with a lottery (large numbers of participants losing relatively little and one participant winning a significant prize with small probability), there are also important differences (no exogenous source of randomness and skill playing a role in the expected outcome). Consequently, this issue of regulation is largely unresolved, although the debate continues given the introduction of other similar lottery-like auctions, such as Price Reveal Auctions (Gallice 2012) and Unique Price Auctions (Rapoport, Otsubo, Kim, and Stein 2007; Raviv and Virag 2009).


References

- Alesina, A. and A. Drazen (1991). Why are stabilizations delayed? *The American Economic Review* 81(5), 1170–1188.
- Anderson, C. K. and F. Odegaard (2011). Retail selling with all-pay auctions. *Working Paper*.
- Anderson, S., D. Friedman, G. Milam, and N. Singh (2008). Buy it now: A hybrid internet market institution. *Journal of Electronic Commerce Research* 9(2).
- Arkes, H. R. and C. Blumer (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes* 35(1), 124–140.
- Ashraf, N., J. Berry, and J. M. Shapiro (2010). Can higher prices stimulate product use? Evidence from a field experiment in Zambia. *American Economic Review* 100(5), 2383–2413.
- Bajari, P. and A. Hortacsu (2004). Economic insights from Internet auctions. *Journal of Economic Literature* 42(2), 457–486.
- Bilodeau, M. and A. Slivinski (1996). Toilet cleaning and department chairing: Volunteering a public service. *Journal of Public Economics* 59(2), 299–308.
- Bliss, C. and B. Nalebuff (1984). Dragon-slaying and ballroom dancing: The private supply of a public good. *Journal of Public Economics* 25(1-2), 1–12.
- Bramsen, J.-M. (2008). A pseudo-endowment effect in internet auctions. *Institute of Food and Resource Economics Working Paper*.
- Bulow, J. and P. Klemperer (1999). The generalized war of attrition. *American Economic Review* 89(1), 175–189.
- Byers, J. W., M. Mitzenmacher, and G. Zervas (2010). Information asymmetries in pay-per-bid auctions: How swoopo makes bank. *arXiv:1001.0592 [cs]*.
- Caldara, M. (2012). Bidding behavior in pay-to-bid auctions: An experimental study. *Working Paper*.
- Camerer, C. F. and R. A. Weber (1999). The econometrics and behavioral economics of escalation of commitment: a re-examination of Staw and Hoang’s NBA data. *Journal of Economic Behavior & Organization* 39(1), 59–82.
- Card, D. and C. A. Olson (1995). Bargaining power, strike durations, and wage outcomes: An analysis of strikes in the 1880s. *Journal of Labor Economics* 13(1), 32–61.

- Clotfelter, C. T. and P. J. Cook (1990). On the economics of state lotteries. *Journal of Economic Perspectives* 4(4), 105–119.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2), 315–372.
- Eyster, E. (2002). Rationalizing the past: A taste for consistency. *Nuffield College Mimeo-graph*.
- Friedman, D., K. Pommerenke, R. Lukose, G. Milam, and B. A. Huberman (2007). Searching for the sunk cost fallacy. *Experimental Economics* 10(1), 79–104.
- Fudenberg, D. and J. Tirole (1986). A theory of exit in duopoly. *Econometrica* 54(4), 943–960.
- Gallice, A. (2012). Price reveal auctions. *Working paper*.
- Goodman, J. (2012). Reputations in bidding fee auctions. *Working Paper*.
- Harsanyi, P. J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory* 2(1), 1–23.
- Hendricks, K. and R. H. Porter (1996). The timing and incidence of exploratory drilling on offshore wildcat tracts. *The American Economic Review* 86(3), 388–407.
- Heyman, J. E., Y. Orhun, and D. Ariely (2004). Auction fever: The effect of opponents and quasi-endowment on product valuations. *Journal of Interactive Marketing* 18(4), 7–21.
- Hinnosaar, T. (2013). Penny auctions are unpredictable. *Working Paper*.
- Ho, T., I. Png, and S. Reza (2014). Sunk cost fallacy in driving the world’s costliest cars. *Working Paper*.
- Horisch, H. and O. Kirchkamp (2010). Less fighting than expected: Experiments with wars of attrition and all-pay auctions. *Public Choice* 144(1/2), 347–367.
- Kennan, J. and R. Wilson (1989). Strategic bargaining models and interpretation of strike data. *Journal of Applied Econometrics* 4(S1), S87–S130.
- King, M. (2012). How penny auction websites can leave you with a hole in your pocket. *the Guardian*.
- Klein, J. P. and M. L. Moeschberger (2003). *Survival analysis: techniques for censored and truncated data*. New York: Springer.
- Krishna, V. and J. Morgan (1997). An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory* 72(2), 343–362.


- Maskin, E. and J. Tirole (2001). Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory* 100(2), 191–219.
- Murnighan, J. K. (2002). A very extreme case of the dollar auction. *Journal of Management Education* 26(1), 56–69.
- Platt, B. C., J. Price, and H. Tappen (2013). The role of risk preferences in pay-to-bid auctions. *Management Science* 59(9), 2117–2134.
- Rapoport, A., H. Otsubo, B. Kim, and W. E. Stein (2007). Unique bid auctions: Equilibrium solutions and experimental evidence. *Working Paper*.
- Raviv, Y. and G. Virag (2009). Gambling by auctions. *International Journal of Industrial Organization* 27(3), 369–378.
- Shubik, M. (1971). The dollar auction game: A paradox in noncooperative behavior and escalation. *The Journal of Conflict Resolution* 15(1), 109–111.
- Stone, B. (2009). Sites ask users to spend to save. *The New York Times*.
- Thaler, R. H. (1990). Anomalies : Saving, fungibility, and mental accounts. *The Journal of Economic Perspectives* 4(1), 193–205.
- Thaler, R. H. (2009). Paying a price for the thrill of the hunt. *The New York Times*.
- Wang, Z. and M. Xu (2011). Learning and strategic sophistication in games: Evidence from penny auctions on the internet. *Working paper*.
- Zheng, H., K. Goh, and K. Huang (2011). Evaluating the impacts of auction bidding restrictions on consumer surplus and behaviors - an empirical study of penny auctions. *ICIS 2011 Proceedings*.

Figure 1: Auction Profits




Notes: Left Graph: Histogram of auction profits in dollars. Right Graph: Histogram of auction profits as a percentage of good's value. Dotted lines represents zero profits. The top and bottom 1% of profit observations have been excluded for readability.

Figure 2: Hazard Rates as Aggregate Sunk Costs Rise




Notes: Auctions are separated by bid increment. The dashed line is the theoretical prediction of the auction hazard rate (likelihood that the auction ends at a given point conditional on reaching that point) in normalized time. The solid line is the empirical hazard rate (with 95% confidence intervals) calculated using the method described by Klein and Moeschberger (2003) with an Epanechnikov kernel given a 10-unit bandwidth.

Figure 3: Instantaneous Profits as Aggregate Sunk Costs Rise



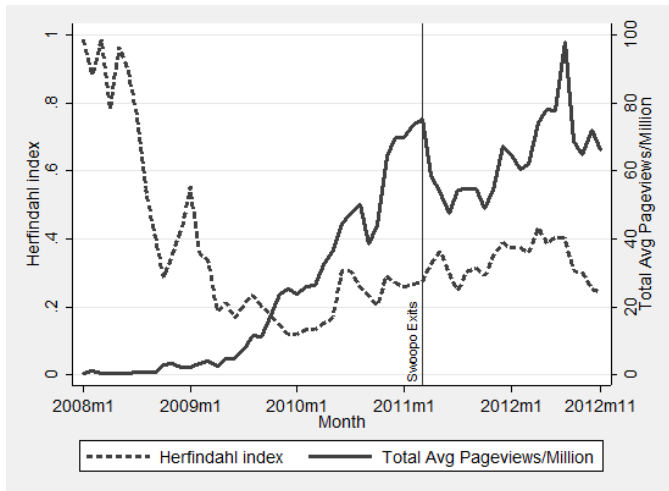
Notes: Auctions are separated by bid increment. The dashed line is the theoretical prediction of the instantaneous profit margin, which is always zero. The solid line is the empirical instantaneous profit margin (with 95% confidence intervals), calculated using the hazard rates in the previous figure.

Figure 4: Probability Player Leaves an Auction Given Sunk Costs



Notes: The line shows the local polynomial estimation of likelihood that a user leaves an auction as a function of the number of bids placed in that auction (with 95% confidence intervals). The dashed line shows the number of estimated users in the auction at the time of the bid to demonstrate that changes in this variable are not driving the effect.

Figure 5: Concentration and Demand Over Time



Notes: Solid line shows the total pageviews / million views of 115 penny auction websites (statistics from Alexa Internet) from 2008-2012, where 100 represents 0.01 of total global Internet traffic. The dashed line shows the Herfindahl index calculated using the same individual pageview metric for the 115 firms.

Table 1: Descriptive Statistics of 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 | 382.21 | 509.80 | 35.67 | 1455.55 |
| Adjusted Value | 166,379 | 342.86 | 477.64 | 23.99 | 1331.09 |
| Nailbiter Auction | 166,379 | .095 | - | - | - |
| <i>Bidding Increment</i> | | | | | |
| \$0.00 | 166,379 | .176 | - | - | - |
| \$0.01 | 166,379 | .064 | - | - | - |
| \$0.15 | 166,379 | .759 | - | - | - |
| <i>Types of Good</i> | | | | | |
| Consumer Goods | 166,379 | .887 | - | - | - |
| 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 | 18,063 | 334.63 | 423.79 | 34.57 | 1331.29 |
| Adjusted Value | 18,063 | 282.81 | 374.37 | 19.99 | 1259.30 |
| Nailbiter Auction | 18,063 | .29 | - | - | - |
| Number Unique Bidders | 18,063 | 53.53 | 90.01 | 4 | 218 |
| <i>Bid Characteristics</i> | | | | | |
| Used BidButler | 13,363,931 | .625 | - | - | - |
| <i>User Characteristics</i> | | | | | |
| Number of Bids | 129,403 | 103.27 | 594.65 | 1 | 285 |
| Number of Auctions | 129,403 | 7.47 | 16.37 | 1 | 23 |
| Number of Wins | 129,403 | .139 | 1.05 | 0 | 0 |

Notes: The bid-level dataset covers a subset of the auction-level dataset. For binary characteristics, such as *Used BidButler*, the mean represents the likelihood of an observation having that characteristic. *Adjusted Value* refers to the price at Amazon at the time of the auction (when available).

Table 2: Descriptive Statistics of Profit

| | Number of Observations | Average Adjusted Value | Average Profit | Average Profit Margin |
|-------------------|---------------------------|---------------------------|-------------------|--------------------------|
| All | 166,379 | \$342.85 | \$159.40 | 50.6% |
| Bidding Increment | | | | |
| \$0.15 | 126,328 | \$273.93 | \$58.21 | 29.0% |
| \$0.01 | 10,709 | \$671.55 | \$866.60 | 181.6% |
| \$0.00 | 29,342 | \$519.65 | \$336.93 | 95.9% |
| Types of Prizes | | | | |
| Consumer | 147,589 | \$359.85 | \$135.38 | 33.1% |
| Bid Vouchers | 16,603 | \$181.72 | \$313.14 | 199.4% |
| Cash Voucher | 2,187 | \$419.27 | \$612.70 | 103.8% |

Notes: "Average Profit Margin" refers to the unweighted average of profit margins and therefore does not match "Average Profit" divided by "Average Adjusted Value."

Table 3: Auction Hazard Rate and Aggregate Sunk Costs

| Dependent Var: $\mathbf{1}[\text{End}]$ | (1) | (2) | (3) | (4) |
|---|----------------------------|----------------------------|---------------------------|---------------------------|
| Ln[Aggregate Sunk Costs] | -0.000120*** (0.000032) | -0.000331*** (0.000038) | -0.000661*** (0.00011) | -0.000946*** (0.00014) |
| Ln[Net Value of Good] | -0.00190*** (0.00011) | -0.00452*** (0.00036) | -0.00862*** (0.00099) | -0.00857*** (0.00064) |
| Constant | 0.0145*** (0.00096) | - - | - - | - - |
| Feb 2009-May 2009 Only | - | - | X | - |
| Nailbiter Only | - | - | - | X |
| Auction Characteristics FEs | - | X | X | X |
| Observations | 94,065,963 | 94,065,963 | 13,382,471 | 3,382,471 |
| Adjusted R^2 | 0.0029 | 0.0048 | 0.0088 | 0.0076 |

Notes: Standard errors in parentheses (clustered on items in all regressions). Linear regressions of the binary outcome of an auction ending at a given point on the log of total amount of sunk costs that have been spent in that auction and the log of the net value of the good. Columns (2)-(4) include auction characteristic fixed effects. Column (3) excludes auctions not included in the bid-level dataset. Column (4) excludes auctions which allow an automated system. Constant not reported for regressions with fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Individual Behavior and Individual Sunk Costs

| Dependent Var: $\mathbf{1}[\text{Leave}]$ | (1) | (2) | (3) | (4) | (5) |
|---|------------------------|------------------------|------------------------|------------------------|-------------------------|
| Ln[Individual Sunk Costs] | -0.0434*** (0.0016) | -0.0445*** (0.0017) | -0.0271*** (0.0011) | -0.0188*** (0.0017) | -0.0803*** (0.0025) |
| Ln[Net Value of Good] | 0.00118 (0.0019) | -0.0310*** (0.0028) | -0.0285*** (0.0033) | -0.0378*** (0.0047) | -0.0269*** (0.0031) |
| Ln[Experience]*Ln[Sunk Costs] | - - | - - | - - | - - | 0.00702*** (0.00012) |
| Constant | 0.209*** (0.013) | - - | - - | - - | - - |
| Feb 2009-May 2009 Only | X | X | X | X | X |
| Nailbiter Only | - | - | - | X | - |
| Auction Characteristics FEs | - | X | X | X | X |
| User FE | - | - | X | X | X |
| Observations | 13,178,970 | 13,178,970 | 13,178,970 | 1,249,038 | 13,178,970 |
| Adjusted R^2 | 0.113 | 0.116 | 0.205 | 0.264 | 0.205 |

Notes: Standard errors in parentheses (double clustered on users and items in all regressions). Linear regressions of the binary outcome of an individual player leaving an auction on the log of amount of sunk costs the individual has spent in that auction and the log of the net value of the good. Columns (2)-(5) include auction characteristic fixed effects. Columns (3)-(5) include individual user fixed effects. Column (4) excludes auctions which allow an automated system. Column (5) adds an experience-sunk cost interaction effect. Ln[Experience] is also included in this regression (coefficient = .0069). Constant not reported for regressions with fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 5: Descriptive Profit Statistics of Competition

| Company | Active Since | Auctions/Day | Profit/Day | Profit Perc |
|----------|--------------|--------------|-------------|-------------|
| Swoopo | 10/2005 | 271.77 | \$63,322.53 | 62.74% |
| BidStick | 10/2008 | 38.22 | \$3,656.38 | 51.76% |
| GoBid | 02/2009 | 9.12 | -\$110.74 | 7.0% |
| Zoozle | 02/2009 | 6.64 | \$164.27 | 3.31% |
| RockyBid | 03/2009 | 9.98 | -\$628.72 | -11.9% |
| BidRay | 04/2009 | 1.75 | \$127.31 | 62.31% |

Notes: Auction and profit statistics from five major competitors as of mid-2009. Statistics calculated from October 2008 to June 2009. Companies ordered by entry date.

A Appendix

A.1 Proof of Main Results

The results for hazard rates hold for non-Markovian strategies (in which players condition on the leader history) with leader history H_t replacing (t, l_t) and some notational changes.

Proposition 1

Proof: Assume that an equilibrium exists in which $h(t^*, l_t^*) < 1$ for some history (t^*, l_t^*) where $t^* > \frac{v-c}{k} - 1$. Then, there must be some player $i \neq l_t^*$ with $p_t^i > 0$. Given some (t, l_t) , define the probability that player i has a bid accepted at (t, l_t) as $a^i(t, l_t) \in [0, 1]$ and the probability that the game ends at (t, l_t) as $q(t, l_t) \in [0, 1]$. Note that as $p_t^i > 0$, it must be that $a^i(t^*, l_t^*) > 0$. Player i 's continuation payoff in the proper subgame starting at (t^*, l_t^*) is then: $E[\sum_{t=t^*}^{\infty} a^i(t, l_t)(-c + q_{t+1}(t+1, i)(v - (t+1)k))] < E[\sum_{t=t^*}^{\infty} a^i(t, l_t)(-c + q_{t+1}(t, i)(v - (\frac{v-c}{k} + 1)k))] < E[\sum_{t=t^*}^{\infty} a^i(t, l_t)(-c + q_{t+1}(t+1, i)(c - k))] < 0$. But, player i could deviate to setting $p_t^i = 0$ and receive a payoff of 0. Therefore, this can not be an equilibrium.

Proposition 2

Proof: Note that the hazard function associated with the strategies matches those in the Proposition: for $t = 0$, $h(t, l_t) = 0$; for $0 < t \leq F$, $h(t, l_t) = (1 - (1 - \sqrt[n-1]{\frac{c}{v-tk}}))^{n-1} = \frac{c}{v-tk}$; for $t > F$, $h(t, l_t) = 1$.

Claim: this set of strategies is a Markov Perfect Equilibrium.

First, consider if $k > 0$. Note that the game is non-stationary. I will show that, for any (t, l_t) , the following statement (referred to as statement 1) is true: there is no strictly profitable deviation from the listed strategies at (t, l_t) and the continuation payoff from entering (t, l_t) as a non-leader is 0. For the subgames starting at (t, l_t) with $t > F$, refer to the proof of Proposition 1 for a proof of the statement. For the subgames starting at (t, l_t) with $t \leq F$, the proof continues using (backward) induction with the statement already proved for any (t, l_t) with $t > F$. At (t, l_t) , non-leader player i will receive an expected continuation payoff of 0 from not betting (she will receive 0 at (t, l_t) and will enter some $(t+1, l_{t+1})$ as a non-leader, which has a continuation payoff of 0 by induction). By betting, there is some positive probability her bid is accepted. If this is the case, she receives $-c$ at (t, l_t) , and will enter $(t+1, i)$ as the leader. The probability that she wins the auction at $(t+1, i)$ is $h(t+1, i) = \frac{c}{v-(t+1)k}$, in which case she will receive $v - (t+1)k$. The probability that she loses the auction at $(t+1, i)$ is $1 - \frac{c}{v-tk}$, in which case she will enter $(t+2, l_{t+2})$ as a non-leader, which must have a continuation payoff of 0 by induction. This leads to a

total continuation payoff from her bid being accepted of $-c + \frac{c}{v-(t+1)k}(v - (t+1)k) = 0$. Alternatively, if the bid is not accepted, she enters $(t+1, l_{t+1})$ as a non-leader and receives a continuation payoff of 0 by induction. Therefore, the continuation payoff from betting must be 0. Therefore, statement 1 is true for all periods and this is a Markov Perfect Equilibrium.

Second, consider if $k = 0$.⁴⁴ First, note that, given the strategies listed in the Proposition, the continuation payoff for every player of entering a period t as the leader is $\pi_L(t) = c$ and a non-leader is $\pi_{NL}(t) = 0$, which we refer to as π_L and π_{NL} as they are independent of time. Consider the (potentially non-equilibrium) strategy that leads to the highest payoff for player i . Suppose this provides a higher expected payoff for player i than the equilibrium strategy. Then, the strategy must involve bidding at some point as never bidding leads to a payoff of 0. Consider a period t in which player i bids with positive probability and define the continuation payoff for player i given this strategy as $\pi_L^*(t)$ if the player is the leader and $\pi_{NL}^*(t)$ if the player is not the leader. Define the probability of bidding at period t for this strategy as $p > 0$. Define the probability of player i 's bid counting at period t given she makes a bid and given other players' equilibrium strategies as q . Then, $\pi_{NL}^*(t)$ can be written as $p(q(\pi_L^*(t+1) - c) + (1-q)\pi_{NL}^*(t+1)) + (1-p)\pi_{NL}^*(t+1)$. We will now show that (1) $\pi_{NL}^*(t+1) = \pi_{NL}^*(t)$ and (2) $\pi_L^*(t+1) = c + (1 - \frac{c}{v})\pi_{NL}^*(t)$. Consider (1). Note that all players other than i play the same strategy regardless of the history. Therefore, (i) $\pi_{NL}^*(t+1) \not\leq \pi_{NL}^*(t)$ because player i could just change her strategy from period $t+1$ and onward to match that of period t and onward (achieve a higher payoff at time t) and (ii) $\pi_{NL}^*(t+1) \not\geq \pi_{NL}^*(t)$ because player i could just change her strategy from period t and onward to match that of period $t+1$ and onward, (contradicting the supposition that her the strategy at time t led to the highest payoff). Therefore, (1) is true. Now, consider $\pi_L^*(t+1)$. Given the other players' strategies, with probability $\frac{c}{v}$, no other player will bid, so that player i will win the auction and receive v , while with probability $1 - \frac{c}{v}$, some player will bid and player i will enter period $t+2$ and have $\pi_{NL}^*(t+2)$, which is equal to $\pi_{NL}^*(t+1)$ and therefore $\pi_{NL}^*(t)$, by the same logic as (1). Therefore, (2) is true. Given (1) and (2), $\pi_{NL}^*(t) = p(q(\pi_L^*(t+1) - c) + (1-q)\pi_{NL}^*(t+1)) + (1-p)\pi_{NL}^*(t+1)$ can be rewritten as $\pi_{NL}^*(t) = p(q(c + (1 - \frac{c}{v})\pi_{NL}^*(t) - c) + (1-q)\pi_{NL}^*(t)) + (1-p)\pi_{NL}^*(t) = \pi_{NL}^*(t) - \frac{c}{v}pq\pi_{NL}^*(t)$. As $c > 0$, $v > 0$, and $p > 0$, this can only be true if $\pi_{NL}^*(t) = 0$, which is a contradiction.

Proposition 3

Proof:

⁴⁴The author wishes to thank an anonymous referee who pointed out a hole in the previous proof for the $k = 0$ case. The previous proof checked for one-step deviations, but this is not sufficient as the game does not satisfy continuity at infinity. The current proof for the $k = 0$ is a modified version of Lemma B.3 in Hinosaar (2013).

First, consider if $k > 0$.

Consider statement (1). I will show that, for each period t , (A) for any (t, l_t) that is reached in equilibrium, $h(t, l_t)$ must match those in Proposition 2 if $t > 1$ and (B) the continuation payoff from any player $i \neq l_{t-1}$ entering any $(t-1, l_{t-1})$ that is reached in equilibrium, $\pi_i(t-1, l_{t-1})$, must be zero. By Proposition 1, the statement (A) is true for all (t, l_t) where $t > \frac{v-c}{k} - 1 = F$. Now, consider statement (B). As $h(t, l_t) = 1$ for every period $t > F$, it must be that $p_i = 0$ for each player $i \neq l_t$ for every period $t > F$. Then, it must be that $\pi_i(t, l_t) = 0$ if $t > F$ for all players $i \neq l_t$ as no player bids for any $t > F$. Finally, consider $\pi_i(F, l_F)$ for any player $i \neq l_F$. There are three possible outcomes for player $i \neq l_F$ at (F, l_F) , all of which lead to a continuation payoff of 0. First, the game ends. Second, another player enters period $F+1$ as the leader, where player i 's continuation payoff is $\pi_i(F+1, l_{F+1})$ where $i \neq l_{F+1}$, which must be 0 by the above proof. Third, player i enters period $F+1$ as the leader, in which case her payoff must be $-c + h(F+1, i)(v - Fk) + (1 - h(F+1, i))\pi_i(F+2, l_{F+2}) = -c + v - (\frac{v-c}{k})k = 0$ as $h(t, l_t) = 1$ for any (t, l_t) if $t > F$ by Proposition 1. Therefore, $\pi_i(F, l_F) = 0$ for $i \neq l_F$ and statement (B) is proven if $t > F$.

For $1 < t \leq F$, the proof continues using (backward) induction with the statement already proved for all periods t with $t > F$. First consider statement (A). Taking the other players' strategies as fixed, define the probability of each player $i \in \{1, 2, \dots, n\}$ being chosen as the leader in $t+1$ at (t, l_t) as $q_i^{j=B}(t, l_t)$ if player j bids and $q_i^{j=NB}(t, l_t)$ if player j does not bid. Note that $q_i^{i=B}(t, l_t)$ must be strictly positive. For part (1) of the statement, consider some (t^*, l_t^*) which is reached in equilibrium in which $h(t^*, l_t^*) \neq \frac{c}{v-tk}$. Consider any $(t^* - 1, l_{t-1}^*)$ that proceeds (t^*, l_t^*) and any $(t^* - 2, l_{t-2}^*)$ that proceeds $(t^* - 1, l_{t-1}^*)$. Note that $l_t^* \neq l_{t-1}^*$ and $l_{t-1}^* \neq l_{t-2}^*$. The expected difference in continuation payoff from player l_t^* in period $t-1$ for history $(t^* - 1, l_{t-1}^*)$ from bidding and not bidding is:

$$q_{l_t^*}^{l_t^*=B}(t^* - 1, l_{t-1}^*)(-c + h(t^*, l_t^*)(v - tk)) + (1 - q_{l_t^*}^{l_t^*=B}(t^* - 1, l_{t-1}^*)) \sum_{j \neq l_t^*} q_j^{l_t^*=B}(t^* - 1, l_{t-1}^*) * \pi_{l_t^*}(t, j) - \sum_{j \neq l_t^*} q_j^{l_t^*=NB}(t^* - 1, l_{t-1}^*) * \pi_{l_t^*}(t, j).$$

By induction, $\pi_{l_t^*}(t, j) = 0$ for any $j \neq l_t^*$. Therefore, the above equation simplifies to $q_{l_t^*}^{l_t^*=B}(t^* - 1, l_{t-1}^*)(-c + h(t^*, l_t^*)(v - tk))$. Now, consider the situation in which $h(t^*, l_t^*) < \frac{c}{v-tk}$. In this case, the difference in continuation payoff is negative, and therefore player l_t^* must strictly prefer to not bid at any $(t^* - 1, l_{t-1}^*)$ that proceeds (t^*, l_t^*) . But then (t^*, l_t^*) will not be reached in equilibrium and we have a contradiction. Next, consider the situation in which $h(t^*, l_t^*) > \frac{c}{v-tk}$. In this case, the difference in continuation payoff is positive, and therefore player l_t^* must strictly prefer to bid in period $t-1$ at any $(t^* - 1, l_{t-1}^*)$ that proceeds (t^*, l_t^*) . This implies that $h(t^*, l_t^*) = 0$ in equilibrium. However, now consider l_{t-1}^* in period $t-2$ in any $(t^* - 2, l_{t-2}^*)$ that proceeds $(t^* - 1, l_{t-1}^*)$. Claim: in each potential state of the world at $(t^* - 2, l_{t-2}^*)$ (the other players' bids and the

auctioneer's choice of leader are unknown), player l_{t-1}^* weakly prefers to not bid and, in at least one state of the world, l_{t-1}^* strictly prefers to not bid. First, consider the states of the world in which no other player is bidding. Here, a bid from player l_{t-1}^* leads to an expected continuation payoff of $-c + h(t^* - 1, l_{t-1}^*)(v - (t - 1)k) + (1 - h(t^* - 1, l_{t-1}^*))\pi_{l_{t-1}^*}(t, l_t) = -c$ as $h(t^* - 1, l_{t-1}^*) = 0$ in equilibrium and $\pi_{l_{t-1}^*}(t, l_t) = 0$ by induction. The expected continuation payoff from not bidding in these states of the world is 0, as the game ends. Therefore, in these states, player l_{t-1}^* strictly prefers to not bid. Second, consider the states of world in which another player bids and player l_{t-1}^* 's bid will be accepted. Here, the expected continuation payoff from bidding is $-c$ (as above) and the expected continuation payoff from not bidding is $\pi_{l_{t-1}^*}(t - 1, l_{t-1})$ for some $l_{t-1} \neq l_{t-1}^*$. $\pi_{l_{t-1}^*}(t - 1, l_{t-1})$ must be weakly greater than 0, as a player could guarantee an expected payoff of 0 from never bidding. Therefore, in these states, player l_{t-1}^* strictly prefers to not bid. Note that one state from these first two categories of states must occur, so player l_{t-1}^* strictly prefers to not bid in at least one state. Finally, consider the states of the world in which another player bids and player l_{t-1}^* 's bid will not be accepted. Here, (t, l_{t-1}) is constant if player l_{t-1}^* bids or not, and therefore player l_{t-2}^* weakly prefers to not bid. Therefore, in equilibrium, player l_{t-1}^* must not bid at any $(t^* - 2, l_{t-2}^*)$ that proceeds any $(t^* - 1, l_{t-1}^*)$ that proceeds any (t^*, l_t^*) . But, then we have a contradiction as (t^*, l_t^*) cannot occur in equilibrium.

Next, I will prove statement (B) for period t . Consider $\pi_i(t - 1, l_{t-1})$ for any player $i \neq l_{t-1}$ in any $(t - 1, l_{t-1})$ that is reached in equilibrium. There are three possible outcomes for player i at $(t - 1, l_{t-1})$, all of which lead to a continuation payoff of 0. First, the game ends. Second, another player enters period t as the leader, in which case player i 's continuation payoff is $\pi_i(t, l_t)$ for some $l_t \neq i$, which must be 0 by induction. Third, player i enters period t as the leader, in which case her payoff must be $-c + h(t, i)(v - tk) + (1 - h(t, i))\pi_i(t + 1, l_{t+1}) = -c + \frac{c}{v - tk}(v - tk) + (1 - \frac{c}{v - tk})\pi_i(t + 1, l_{t+1}) = 0$ as $h(t, i) = \frac{c}{v - tk}$ for period t by above and $l_{t+1} \neq i$, so $\pi_i(t + 1, l_{t+1}) = 0$ by induction. Therefore, it must be that $\pi_i(t - 1, l_{t-1})$ for any player $i \neq l_{t-1}$ in any $(t - 1, l_{t-1})$ that is reached in equilibrium and the statement is proved.

Consider Statement (2):

Assume there is an equilibrium in which player i uses Markov Strategies and $p_0^i > 0$, $p_1^i > 0$ for all i .

For each period $t > 1$, I will prove that player i must follow the strategies listed in the Proposition 2. First, note that by Proposition 1, $h(t, l_t) = 1$ where $t > \frac{v - c}{k} - 1 = F$, so it must be that $p_t^i = 0$ for each player i for every period $t > F$. For periods $1 < t \leq F$, the proof is by induction with period 2 as the initial period. Period 2: As $p_0^i > 0$, $p_1^i > 0$ for all i , it must be true that for each player i , $(2, l_2 = i)$ occurs on the equilibrium path. Suppose that $p_t^i \neq p_t^j$ for

some players i and j for $t = 2$. Then, $h(t, l_t = i) = \frac{\prod_{k=1}^n (1-p_t^k)}{(1-p_t^i)} \neq \frac{\prod_{k=1}^n (1-p_t^k)}{(1-p_t^j)} = h(t, l_t = j)$ for $t = 2$. But, by Statement (1) of Proposition **3**, it must be that $h(t, l_t = i) = \frac{c}{v-2k} = h(t, l_t = j)$ for $t = 2$ so we have a contradiction. Therefore, $p_t^i = p_t^j$ for all i and j and therefore $p_t^i = \sqrt[n]{1 - \frac{c}{v-2k}}$ for all i when $t = 2$. Period t : Suppose the statement is true for periods prior to t . Then, it must be true that for each player i , $(t, l_t = i)$ occurs on the equilibrium path. Now, follow the rest of the proof for $t = 2$ for any $t \leq F$ to show that the statement holds for any period $1 < t \leq F$. Therefore, in any Markov Perfect Equilibrium in which play continues past period 1, strategies must match these after period 1.

Second, consider if $k = 0$.

Consider Statement (1):

Assume that players use symmetric strategies: $p_t^i = p_t^j = p_t$. Note that this implies that $h(t, l_t = i) = h(t, l_t = j) = h(t)$. Define the continuation payoff for every player of entering period t as the leader as $\pi_L(t)$ and a non-leader as $\pi_{NL}(t)$. Claim: $\pi_{NL}(t) = \pi_L(t) - c$ for any period $t > 1$ that appears on the equilibrium path. First, suppose that there exists t on the equilibrium path such that $\pi_{NL}(t) > \pi_L(t) - c$. Using notation from the Proof to Proposition **3**, the difference in the expected payoff from bidding and not bidding for player i at period $t - 1$ is $(1 - q_i^{i=B}(t-1))\pi_{NL}(t) + q_i^{i=B}(t-1)(-c + \pi_L(t)) - \pi_{NL}(t) < 0$. Therefore, all non-leading bidders must strictly prefer to not bid in period $t - 1$. However, this implies that t cannot be reached on the equilibrium path, a contradiction. Second, suppose that there exists t on the equilibrium path such that $\pi_{NL}(t) < \pi_L(t) - c$. The difference in the expected payoff from bidding and not bidding for player i at period $t - 1$ is $(1 - q_i^{i=B}(t-1))\pi_{NL}(t) + q_i^{i=B}(t-1)(-c + \pi_L(t)) - \pi_{NL}(t) > 0$. Therefore, all non-leading bidders must strictly prefer to bid in period $t - 1$. This implies that $\pi_L(t - 1) = \pi_{NL}(t)$ as a leader in period $t - 1$ will necessarily become a non-leader in period t . It also implies that $\pi_{NL}(t - 1) \geq \pi_{NL}(t)$ as a non-leader in period $t - 1$ could not bid and guarantee $\pi_{NL}(t)$. Therefore, the difference in the expected payoff from bidding and not bidding for player i at period $t - 2$ is

$$\begin{aligned} & (1 - q_i^{i=B}(t-2))\pi_{NL}(t-1) + q_i^{i=B}(t-1)(-c + \pi_L(t-1)) - \pi_{NL}(t-1) = \\ & (1 - q_i^{i=B}(t-2))\pi_{NL}(t-1) + q_i^{i=B}(t-1)(-c + \pi_{NL}(t)) - \pi_{NL}(t-1) < \\ & (1 - q_i^{i=B}(t-2))\pi_{NL}(t) + q_i^{i=B}(t-1)(-c + \pi_{NL}(t)) - \pi_{NL}(t) = \\ & q_i^{i=B}(t-1)(-c) < 0 \end{aligned}$$

Therefore, players in $t - 2$ must strictly prefer to not bid. This implies that period t is not on the equilibrium path, a contradiction. Therefore, it must be that $\pi_{NL}(t) = \pi_L(t) - c$

for any period $t > 1$ that appears on the equilibrium path. Now, note in equilibrium $\pi_L(t) = H(t)v + (1 - H(t))\pi_{NL}(t + 1)$ and $\pi_{NL}(t) = H(t)(0) + (1 - H(t))(\frac{1}{n}(-c + \pi_L(t + 1)) + \frac{n-2}{n-1}\pi_{NL}(t + 1))$. Suppose $t > 1$ and t is on the equilibrium path. Note that $t + 1$ must also be on the equilibrium path: If $H(t) = 1$, then $\pi_L(t) = v > \pi_{NL}(t) + c$, a contradiction. Therefore, it must be that $\pi_L(t) = \pi_{NL}(t) + c$ and $\pi_L(t + 1) = \pi_{NL}(t + 1) + c$. Imposing this on the equations for $\pi_{NL}(t)$ and $\pi_L(t)$ yields the unique solution: $H(t) = \frac{c}{v}$. Therefore, the Proposition is true.

Consider Statement (2):

Assume there is an equilibrium in which players uses Symmetric Markov Strategies and $p_0 > 0, p_1 > 0$.

For each period $t > 1$, I will prove that players must follow the strategies listed in the Proposition 2. The proof is by induction with period 2 as the initial period. Period 2: As $p_0 > 0, p_1 > 0$, it must be true that period t occurs on the equilibrium path. By Statement (1) of Proposition 3, it must be that $h(t, l_t) = \frac{c}{v}$ and therefore $p_t = \sqrt[n-1]{1 - \frac{c}{v-tk}}$ when $t = 2$. Period t : Suppose the statement is true for periods prior to t . Then, it must be true that period t occurs on the equilibrium path. Now, follow the rest of the proof for $t = 2$ for any t to show that the statement holds for any period $1 < t$. Therefore, in any Symmetric Markov Perfect Equilibrium in which play continues past period 1, strategies must match these after period 1.

Proposition 4

Proof:

The proof is a transparent corollary of Proposition 2, by just considering the value of the good to be equal to $\tilde{v} = v + \theta s_i c$. By the naivety assumption, each player perceives that they are playing the equilibrium in Proposition 2 with value \tilde{v} , leading to the hazard rates, individual strategies, and profit statistics listed.