Journal of Economics & Management Strategy
25 YEARS ANNIVERSARY

Learning through Simultaneous Play: Evidence from Penny Auctions

Ricardo Gonçalves
Católica Porto Business School and CEGE
Universidade Católica Portuguesa
Rua Diogo Botelho, 1327
4169-005 Porto, Portugal
rgoncalves@porto.ucp.pt

Miguel A. Fonseca
University of Exeter Business School
Streatham Court
Rennes Drive
Exeter EX4 4PU, United Kingdom
m.a.fonseca@exeter.ac.uk

This paper contributes to the emerging empirical literature on penny auctions, a particular type of all-pay auctions. We focus on the potential learning effects that bidders may experience over time but also (and particularly) across auctions as a result of their auction participation. Using detailed bid-level information, we find that, similarly to earlier literature, bidders suffer from a sunk cost fallacy, whereby their probability of dropping out of an auction is decreasing in the number of bids they have already placed in that auction. Although we do find that learning through repeated participation alleviates the sunk cost fallacy, participation in simultaneous penny auctions emerges as a much more effective learning mechanism, ultimately contributing toward bidders earning higher individual surpluses.

1. Introduction

The penny auction was popularized by firms like Swoopo and is still used by online auction companies, as well as by traditional retailers across the world.[1] In a penny auction, bidding usually starts at zero and bidders must pay a bid cost to increase the sale price by a small amount—typically one penny, hence the name of the auction. [2] A key attraction of this type of auction is the possibility of paying substantially less than the retail price for an object. However, this does not necessarily mean the auctioneer has made a loss. For instance, if each bid in a penny auction costs $1 to place, an iPad with a retail price of $500 that is sold for $75 in a penny auction yields a revenue of $7,575 to the

We thank a telecommunications operator (which has requested its identity not to be disclosed) for providing us the bid-level database for the auctions analyzed in this paper. Financial support from Fundação para a Ciência e Tecnologia (through project PEst-OE/EGE/UI0731/2011) is gratefully acknowledged. We thank Ricardo Ribeiro and participants in a seminar at Universidade Católica Portuguesa, as well as in the 7th Meeting of the Portuguese Economic Journal (2013, Covilhã, Portugal) and in the 40th Annual Conference of the European Association for Research in Industrial Economics (EARIE) (2013, Évora, Portugal) for their helpful comments and suggestions. We also thank the editor, a co-editor, and a referee for the valuable suggestions made.

1. Swoopo filed for bankruptcy in 2011.
2. Although penny auctions bear some resemblance to the war-of-attrition game, a type of all-pay auction, they are not a special case of any type of all-pay auctions (Hinnosaar, 2014).

Journal of Economics & Management Strategy, Volume 25, Number 4, Winter 2016, 1040–1059

auctioneer ($7,500 from bid costs plus $75 from the actual sale price) and a substantial profit margin. Indeed, the recent interest in penny auctions has been driven both by its popularity as an e-commerce mechanism, as well as by the empirically observed high-profit margins—a clear violation of auction theory (expected revenue should equal the good's value) and yet another real-world example of overbidding in auctions.

The standard game theoretical analysis assumes that players arrive at a Nash equilibrium through an introspective process, in which they form beliefs about their opponents' actions, and beliefs about their opponents' beliefs about their own actions, and so on. While real people may be able to engage in this type of mental process in simple games with few players and with a unique Nash equilibrium, it is less reasonable to expect this to be true in more complex games with many players and multiple equilibria, as is the case of penny auctions. Instead, as already suggested in the penny auction literature, which we review below, it may be that repeated experience (over time, within a given auction or in subsequent auctions) in this type of game allows players to learn what the optimal strategy is (or at least allows players to identify and use payoff-enhancing strategies).

However, a dimension that has thus far been ignored is the potential contribution that participation in simultaneous auctions has in this learning process. Indeed, "experience" may be obtained "vertically," over time through bid submission or auction participation, but also "horizontally," within a time window but through bid submission or simultaneous participation in more than one auction. Theoretically (Mertens, 1992), one would expect such parallel auctions to be independent and thus not affect bidding behavior; in reality, it may actually speed up the bidders' learning process. To the best of our knowledge, this is a little-explored subject in the economics literature.[3] In penny auctions, bidders are learning a complex object. We conjecture that by virtue of participating in concurrent auctions with different types of bidders, or even by bidding in auctions at different stages, subjects can learn (i) how to play a particular optimal strategy, or learn that (ii) there are potentially different optimal strategies faster than if they only take part in one auction.

There is a small (experimental) literature on behavioral spillovers that is somewhat related to these learning effects. Cason et al. (2012) consider a behavioral spillover to exist when observed behavior in a game is different depending on whether that game is played together with other games or in isolation and acknowledge that learning effects can be a source of such spillovers.[4]

3. In neuroscience, Sigman and Dehaene (2008) show the coexistence of serial and parallel brain processes during the performance of a cognitive task. Gombrich (2011) puts forward a tentative definition of series learning (equivalent to our "vertical" learning) as opposed to parallel learning (equivalent to our "horizontal" learning), borrowing from the working of electric circuits: through series learning, one learns one thing after another to arrive at a total amount of knowledge, while through parallel learning, one learns several things at the same time to arrive at the same amount of total knowledge. In the machine learning (and artificial neural networks) literature, Caruana (1995) proposes and tests several mechanisms through which neural nets learning through multiple related tasks can outperform sequential learning, as it enables a more generalizable representation of a particular feature. Wason (1960) pioneered the paradigm of rule discovery, which studies how humans develop hypotheses from observing data from unknown data generation processes. This is also illustrated well by Baxter (1995), who points out that engaging in multiple tasks enables learning more general representations of concepts.

4. In particular, Cason et al. (2012) look at (two different) coordination games and find strong spillovers when the games are played sequentially, but not when they are played simultaneously. In the same vein, Falk et al. (2013) analyze two identical and completely independent (coordination or public good) games, played simultaneously, and also find no evidence of behavioral spillovers. By contrast, Bednar et al. (2012) do find spillovers when (two different) games are played simultaneously, but these games are different from Cason et al. (2012).

With the objective of analyzing such learning effects, we study data from 403 penny auctions conducted by a telecommunications operator between June and December 2011.[5] Our database offers distinct advantages to study the role of learning through participation in simultaneous auctions. The telecommunications operator auctions off essentially five types of items: mobile phones, tablet PCs, laptops, mobile services (e.g., games, etc.), and USB pens for mobile internet access. While certainly differentiated within each type, these items belong to a relatively homogeneous category of products: devices and/or services that may enhance users' benefits from their current mobile subscriptions. In that sense, these auctions may attract a relatively homogeneous set of potential bidders and this appears to be adequate to study learning effects. By comparison, Swoopo or other major penny auction web sites sell a wider variety of items to certainly more heterogeneous bidders.

We find that the telecommunications operator has an average profit margin of 219%: it obtains an average revenue of €1,224 per auctioned item, whose average retail price is €370. There is significant dispersion in profit margins (including large negative profit margins), and its distribution exhibits positive skewness and is thus right-tailed. Our detailed bid-level dataset allows us to investigate drivers of bidding behavior. First, we analyze the decision to participate in a penny auction and find that the probability of participation by an individual bidder is decreasing with the number of simultaneous auctions and bidder experience—evidence of both vertical and horizontal learning effects. Second, we investigate what determines the likelihood of an active bidder dropping out of the auction. We find evidence of Augenblick's (2016) sunk cost fallacy in our data: the probability of an individual bidder leaving an auction decreases with the number of bids she has already placed in that auction. However, experience contributes significantly toward alleviating it (as in Augenblick, 2016), as does—to a more noticeable extent—the simultaneous participation in more than one penny auction, suggesting that horizontal learning may be more effective than vertical learning. This empirical evidence appears to support the existence of behavioral spillovers, both when auctions are conducted sequentially, as well as simultaneously, and the latter is in contrast with the (experimental) evidence of Falk et al. (2013).[6]

The remainder of the paper is organized as follows. Section 2 contextualizes our paper in the literature. Section 3 describes the data and Section 4 contains the main analysis. Section 5 concludes. Appendix A contains additional material.

2. Literature Review

The literature on penny auctions is relatively recent. Augenblick (2016) proposes a tractable penny auction model, for which he obtains a Markov perfect equilibrium. He uses auction-level and bid-level data from Swoopo to show that the sunk cost fallacy (as bidders spend more money on bids, they become more reluctant to leave the auction although the bidding costs are sunk) explains the persistence of auction participation above and beyond the prediction of the normative model. In turn, this provides a rationale for the observed high and positive (51%) average profit margins in Swoopo auctions. Bidders apparently learn how to use more effective bidding strategies

5. The operator requested its identity not to be disclosed.

6. Particular care must be taken in interpreting (and comparing) this result because unlike the experimental setup of Falk et al. (2013), the environment that characterizes simultaneous penny auctions does not control for all possible differences between the auctions.

(including aggressive bidding strategies), but such learning takes place at a very slow rate and requires a large number of submitted bids (and consequently high losses).

Hinnosaar (2014) provides an alternative analysis of the penny auction game and shows that there are multiple symmetric, stationary, subgame-perfect Nash equilibria. Platt et al. (2013) analyze a penny auction model that explicitly allows bidders to have risk-loving preferences (rather than the traditional assumption of risk neutrality), in an attempt to explain excess profits. Using data from Swoopo, they find that bidding patterns are consistent with some degree of risk loving, thus rationalizing the observed high average profit margins. However, they do not analyze individual-level behavior because they rely solely on auction-level data.

Byers et al. (2010) also try to explain excess profits by considering informational asymmetries, whereby some bidders are better informed regarding the number of active bidders at any given point in time.[7] Unlike Augenblick (2016), Byers et al. (2010) find that although aggressive bidding strategies are associated with higher winning probabilities, they also lead to lower bidder surpluses. Wang and Xu (2012) use data from BigDeal.com to study the role of bidder sophistication. In particular, they find that auctioneers benefit from penny auctions only if participating bidders are inexperienced. Experienced and strategically sophisticated bidders, who over time learn how to time their bids so as to maximize their surplus, earn a small and positive surplus.

Goodman (2012) also uses data from Swoopo and explores effective strategies (bidding frequently, bidding immediately after the previous bid, and using automated bidding services) that bidders use to increase their surplus, typically associated with aggressive and reputation-building behavior. Experimentally, Caldara (2012) finds that overbidding decreases over time, risk-averse individuals submit fewer bids than risk-neutral individuals, strategic sophistication is an important driver of outcomes (confirming Wang and Xu's, 2012, results), signaling strategies are used, but with little success, and the learning process also involves learning not to bid in further auctions.

3. Data Description

The data used in this paper consist of penny auctions conducted by a telecommunications operator roughly on a monthly basis for a period of five working days which we define as "auction week."[8] On average, there are approximately 10 penny auctions in each day of the auction week. Auction weeks were advertised in advance on the operator's web site as well as through targeted e-mails, blogs, social networks, and other web sites. Bidding was restricted to that operator's subscribers and could be done through SMS or through the operator's auction web site.

The operator's web site displayed all of the day's penny auctions, as well as each auction's starting time. Typically, each auction started 30 minutes after the previous one began. Mobile phones were the most commonly auctioned item, but laptops, tablet PCs, USB pens for broadband wireless access and services (e.g., one year of free SMS, 6 months of free online gaming, etc.) were also auctioned. For each auction, a detailed description of the item was provided, including its main characteristics, as well as its retail price.

7. Byers et al. (2010) also exploit the role of asymmetric valuations of the good for sale, as well as asymmetries in bidding costs.

8. As we will later see, the only exception to this rule occurred in December 2011, where the auction lasted for 10 working days (over a two-week period).

All auctions shared the following features: auctions were conducted in euros (€), the starting price was €0, and the auction finishing time was set for one hour after its start; each bidder placed bids, which raised the current high bid by a fixed €0.01 increment; each bid placed had a nonrefundable bid cost of €0.50 (charged directly to the bidder's phone bill). All bids placed within the first hour of the auction did not change its finishing time; however, all bids placed after the initial one hour period extended the auction finishing time by one minute. The winning bidder was whoever held the current high bid once the auction reached its finishing time. The winning bidder's total auction cost for the good was the number of bids placed times the bid cost plus the final auction price. All other bidders forfeited their bidding costs.

After each auction finished, the operator posted on the web site the auction duration, final price, the difference between final price and retail price, and the bid, time of bid, and identity of the bidder for the last five bids (including the auction winner). This information was publicly available until the start of the following auction week (roughly for a one-month period) and we manually collected it for 414 auctions pertaining to seven auction weeks between June and December 2011.

In addition to this, we contacted the operator in order to obtain more detailed bid-level data, which includes, for each auction, the time of the bid and the (numeric) identity of the bidder. In total, we received bid-level information for all 414 auctions in the auction database, totaling 956,742 bids placed by 38,733 unique bidders. A closer inspection of the two databases revealed a "coverage" problem, which has led us to exclude 11 auctions from the analysis and focus on 403 auctions, for which we have 940,968 bids placed by 38,000 unique bidders in the bid-level database.[9]

Table A1 (Appendix A) presents the descriptive statistics of the auction-level database. On average, the final auction price was €24.01 and bid costs yielded €1,200.50 in revenue, which means that using as a reference the average retail price of €369.97, each auction yielded a net profit of €854.77—a profit margin of 219%.[10] Figure A1

9. Note that the auction-level database, by registering the final bid, allows us to infer the total number of bids placed in each auction because each bid increment is fixed and equal to €0.01. Comparing the total number of bids in each auction both in the auction-level as well as in the bid-level database, we noted that the latter typically contained fewer bids than suggested by the former. This problem was particularly acute for 11 auctions in the original sample. In 10 out of these 11 auctions, the coverage was below 30%. We suspect that this may have been due to an error in the bid-registering software (the operator could not explain this), as all these 11 excluded auctions occurred on a given day (14th December 2011) and the missing bid-level observations related to bids placed after 7 pm. The average coverage of these 403 auctions was 97.2%. The coverage problem was widespread and affected 76% of these 403 auctions—that is, the two databases only coincided fully for 24% of all auctions. But in most cases, the differences were relatively small—basically a few missing bids from the bid-level database for each auction (the median is 8 missing bids per auction), which could be explained by small errors in the operator's registering procedure: 44.5% of auctions were missing less than 1% of bids; 57% were missing less than 2%; and 66.5% were missing less than 5%. Therefore, only 9.5% of auctions had missing bids in excess of 5%.

10. The average profit margin in our dataset (219%) is clearly higher than in Augenblick (2016) (51%). This could be associated with the type of good most frequently auctioned in our dataset: mobile phones account for 85% of total auctions. Augenblick (2016), Platt et al. (2013), Byers et al. (2010), and Goodman (2012) all focus on Swoopo auctions—typically auctions of consumer goods, such as home electronics, computer accessories, videogames-related items, and other consumer goods. Platt et al. (2013, figure 2) present evidence of significantly higher deviations (reported in standard deviation units) between observed and theoretical average revenue for videogames-related items: the distribution is clearly right-tailed, with 30% (approximately 50%) of items having average revenue greater than one (one-half) standard deviation above the model's prediction. Byers et al. (2010, figure 3) present evidence of an average profit margin of 86%, with the item that is more frequently auctioned obtaining an average profit margin of 365% (Wii Play with Wii Remote). Goodman (2012, table 3) provides a breakdown of average profit margins by bid increment and by good value: although the average (overall) profit margin is 67%, for some bid increments (2 cents and 6 cents), it is higher than that—126% and 132%, respectively—, and, for those same increments and for goods whose value is in

(Appendix A) shows that some auctions present negative (and large, in absolute value) profit margins, but others yield extremely high profit margins. There is also some variability across auction weeks and between item types. On average, there are 6.88 auctions active at the same time (simultaneous auctions) and in each auction, 2,401 bids were placed by 439 unique bidders. Table A2 (Appendix A) presents the descriptive statistics of the bid-level database, which contains 176,908 unique bidder-auction combinations.

4. Analysis and Results

4.1. Auction-Level Analysis

Using data from the auction-level database, we wish to understand the main drivers of final prices in penny auctions. The auction price is, in effect, a duration-type variable: higher auction prices indicate, according to the bidding rules, longer durations and, because we have normalized auction prices by dividing them by the retail price of the auctioned item (as suggested by Augenblick, 2016), we can compare the durations of auctions of different items.[11] On this basis, we have estimated a model assuming the Gompertz and the Weibull distributions for the hazard rates, as well as Cox's partial likelihood model (which places no restrictions on the shape of the baseline hazard rate).[12, 13]

As potential explanatory variables we have included: (i) the average number of simultaneous auctions; (ii) the number of normalized (divided by the retail price) first hour bids;[14] (iii) the auction week in question; (iv) the type and (v) brand of the auctioned item; (vi) the number of unique bidders; and (vii) a Herfindahl–Hirschman concentration index (HHI) of unique bidders' bids. Table A4 displays the results of our estimations.

First, note that all models produce very similar estimates, which is somewhat reassuring. Second, the exponentiated (significant) coefficient of the average number of simultaneous auctions indicates that all else constant, an additional simultaneous auction increases the baseline hazard rate by 21–27%. Third, the number of normalized first hour bids significantly decreases the hazard rate (by 81–88%), that is, first hour

the \$25–\$50 range, it is even higher (477% and 240%, respectively). Therefore, although average profit margins in our dataset's penny auctions are higher than in Swoopo auctions, there are several instances in which the latter yield comparable or even higher profit margins.

11. In order to establish a parallel with survival analysis, we refer to the normalized auction prices as "normalized time." Figure A2 (in Appendix A) presents the estimated (smoothed) hazard rate using a kernel with a bandwidth of 2, where we can see that the hazard rate appears to be (almost always) increasing with normalized time; in Figure A2, we have also plotted the logarithm of the cumulative hazard against that of normalized time. The fact that we obtain a relatively straight line suggests that the underlying hazard rate in the data could be coming from a Weibull distribution.

12. We have estimated the hazard rates for each parametric distribution and for each model and obtained their respective Akaike's information criterion (AIC). This method takes into account a model's log-likelihood, but also the number of covariates and the number of model-specific parameters (see Akaike, 1974, or Cleves et al., 2010, p. 281, on its application for survival analysis). We have thus identified the two parametric distributions with lower AIC scores in Table A3 (Appendix A). The Weibull and the Gompertz distributions emerge as sensible candidates.

13. As Jenkins (2005) notes, models differ not only in terms of the shape assumed for the hazard rate but also in terms of their specification and interpretation. They can be divided into proportional hazard (PH) models or accelerated failure time (AFT) models. Of the estimated models, Cox and Gompertz are PH, although the Weibull model can be written in both specifications. In PH models, (exponentiated) coefficients are hazard ratios: a coefficient of 1.05 for an explanatory variable indicates that a unit increase in that explanatory variable increases the baseline hazard rate by a factor of 1.05.

14. This variable measures the extent to which bidders may be using signaling strategies by bidding in the first hour (a weakly dominated strategy), as documented by Goodman (2012).

bidding activity is an important explanatory factor of auction duration. And finally, all else constant, an additional unique bidder decreases the hazard rate by 0.5%. Relating these results to our main research question, we find that at the auction level:

Observation 1: *Auction duration varies inversely with the number of simultaneous auctions.*

This could be interpreted as a supply-induced effect (more simultaneous auctions could induce a fixed number of constrained bidders to bid less) or as a first hint for the existence of (aggregate) horizontal learning effects (if the supply-induced effect is small).

4.2. Bid-Level Analysis

4.2.1. Decision to Participate in an Auction

We have used the bid-level database to construct a balanced panel with 15,314,000 observations, which includes all unique bidders (38,000) in all auctions (403). In that panel, the dummy variable "participate" takes on the value of 1 if bidder i participated in auction j and 0 otherwise.

Our purpose is to look deeper into the motivation of each bidder to participate in an auction. In addition to the average number of simultaneous auctions, the auction week in question, the type, and brand of the auctioned item, we also consider three additional variables: the retail price, the number of bids previously submitted by each bidder (in other auctions), and the number of auctions previously won by each bidder. The number of previously submitted bids is of particular interest for two reasons: on the one hand, because it entails a cost, bidders who have already submitted many bids prior to auction j may be (relatively more) financially constrained and hence decide not to participate in that auction; on the other hand, it is clearly a proxy for bidder experience (as in Augenblick, 2016, or Goodman, 2012). Moreover, we have adapted the variables related to the number of unique bidders and the HHI of unique bidders' bids to this setting. [15]

Table I contains the results of a random-effects probit specification, with "participate" as the dependent variable. The results are presented as marginal effects evaluated at the mean. Almost all variables are statistically significant and the coefficients' signs appear plausible. [16] For instance, a higher retail price increases the probability of bidder participation, whereas a higher number of simultaneous auctions decreases it. Interestingly, the probability of a bidder participating in an auction is decreasing with the number of previously submitted bids. This suggests that as bidders become more experienced over time, they learn to participate less often in penny auctions.[17]

15. Given that the participation decision occurs at a given moment in time (when the bidder submits her first bid in a given auction), we have computed the number of unique bidders and the HHI of unique bidders' bids in that auction until that point in time - effectively measuring potential 'perceived' competition up to the point where the bidder decides to participate. If the bidder decides not to participate, we use the number of unique bidders and HHI of unique bidders' bids at the end of the auction.

16. The marginal effects appear to be very small, but one must bear in mind that 98.8% of the observations for the dependent variable ('participate') take on the value of 0, that is, we only observe 176,908 bidder-auction participations (see Table A2) in a total of 15,314,000 possible combinations. Therefore, although the marginal effects (evaluated at the mean) are small, they entail rather more noticeable *percentage* increases in the participation probability.

17. In addition, the number of auctions previously won has a positive coefficient, which implies that previous auction wins have a positive effect on the probability of participating in subsequent auctions. Finally, perceived competition (measured through the number of unique bidders until the participation decision) has a negative effect on the participation probability. However, unique bidders' bid concentration has the opposite

TABLE I.
RANDOM-EFFECTS PROBIT RESULTS ON DECISION TO PARTICIPATE

Dependent Variable: Participate	RE Probit Coef. (Std. Error)
Dummies for auction week	yes
Dummies for type of good	yes
Dummies for brand of good	yes
Retail price	0.00002 (0)***
Average number of simultaneous auctions	−0.00167 (0.00001)***
HHI of unique bidders' bids until participation decision	0.000002 (0)***
Number of unique bidders until participation decision	−0.00002 (0)***
Number of previously submitted bids	−0.00002 (0)***
Number of previous auction wins	0.00785 (0.00019)***
N	15,314,000
Log-likelihood	−837,004
LR (chi-square)/Wald (chi-square)	169,935

*** Significant at the 1% level; ** significant at the 5% level; * Significant at the 10% level.

OBSERVATION 2: *More experienced bidders are less likely to enter in new auctions.*

OBSERVATION 3: *The higher is the number of simultaneous auctions, the lower is the probability of entering an auction.*

4.2.2. INDIVIDUAL PSEUDO-HAZARD RATES

We have also adopted Augenblick's (2016) methodology to estimate individual-level pseudo-hazard rates, which are defined as the probability that an individual exits an auction given the number of bids she has already submitted up to that point. Therefore, we have created a dummy variable "leave" that is equal to 1 if a bidder's bid in an auction is her last bid in that auction and 0 otherwise. We have then estimated a random effects probit model for two different specifications. Under a first specification, we have included as explanatory variables: (i) the average number of simultaneous auctions; (ii) a first hour bid dummy variable; (iii) the number of simultaneous auctions in which the bidder is actively participating; (iv) the number of unique bidders until the moment a bid is submitted; (v) the Herfindahl–Hirschman concentration index of unique bidders' bids until the moment a bid is submitted; (vi) the number of bids submitted by the bidder in that auction; (vii) the number of bids submitted by the bidder in previous auctions (a proxy for experience);[18] (viii) the number of auctions previously won by the bidder; (ix) the bid amount; (x) the auction week dummies; (xi) the type; and (xii) brand of the auctioned item.

Under a second specification, we have followed Augenblick (2016) by adding an interaction variable between experience (number of previously submitted bids) and the number of bids placed in the auction. We have also interacted simultaneous auction participation with the number of bids placed in the auction. Finally, we have also interacted the first hour dummy variable with experience and the number of simultaneous auctions in which the bidder is participating. Table II presents the results as marginal effects evaluated at the mean.

effect: that is, a higher concentration of bids (which could be interpreted as a "seriousness of competition" indicator) actually increases the probability of participation.

18. Rather than using the absolute number of bids (in the auction and in previous auctions), we have chosen to use their logs because of convergence problems in model estimation.

Table II.
Pseudo-Hazard Rate Estimation

Dependent Variable: Leave	RE Probit Coef. (Std. Error)	RE Probit Coef. (Std. Error)
Dummies for auction week	yes	yes
Dummies for type of good	yes	yes
Dummies for brand of good	yes	yes
Average number of simultaneous auctions	0.02955 (0.00055)***	0.02958 (0.00055)***
First hour bid (dummy)	−0.00657 (0.00366)*	−0.01419 (0.00589)**
First hour bid (dummy) × ln(number of previously submitted bids)		−0.00392 (0.00262)
First hour bid (dummy) × ln(number of simultaneous auctions in which bidder is participating)		0.01911 (0.00584)***
Number of unique bidders until bid submitted	0.00011 (0.00001)***	0.00011 (0.00001)***
HHI of unique bidders' bids until bid is submitted	0.00000 (0)	0.00000 (0)
ln(number of bids submitted in auction)	−0.11852 (0.00071)***	−0.17833 (0.00184)***
ln(number of previously submitted bids)	0.05555 (0.00099)***	0.04739 (0.00102)***
ln(number of simultaneous auctions in which bidder is participating)	0.05263 (0.00132)***	0.04533 (0.00178)***
ln(number of previously submitted bids) × ln(number of bids submitted in auction)		0.01428 (0.00043)***
ln(number of simultaneous auctions in which bidder is participating) × ln(number of bids submitted in auction)		0.00597 (0.00089)***
Number of previous auction wins	0.01592 (0.00459)***	−0.00215 (0.00461)
Bid amount	0.00011 (0.00011)	−0.00001 (0.00011)
N	940,968	940,968
Log-likelihood	−385,991	−385,315
LR (chi-square)/Wald (chi-square)	40,800	42,429

*** Significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

First, the coefficient on the average number of simultaneous auctions is positive and statistically significant. Second, first hour bids decrease the pseudo-hazard rate, that is, a bidder who submits a bid in the first hour of the auction is less likely to drop out.[19] Third, the sunk cost fallacy is present: the probability of leaving the auction decreases with the number of bids already placed by the bidder in that auction. By contrast, participation in simultaneous auctions increases the drop out probability, as does experience. Fourth, we find that bidder experience does tend to reduce the effect of an additional bid on the probability of leaving the auction (interaction term between experience and the number of bids submitted in the second column), that is, it alleviates the sunk cost fallacy, a result that is similar to that obtained by Augenblick (2016). Interestingly, participation in multiple simultaneous auctions also appears to have an important role toward alleviating the sunk cost fallacy, as the interaction variable between simultaneous auction participation and the number of bids submitted has a positive (and significant) coefficient (second column).

Observation 4: *There is evidence of behavioral spillovers, some of which consist of vertical and horizontal learning effects. Both types of learning contribute to alleviating the sunk cost fallacy.*

19. This seems to suggest that first hour bids may indeed be used as signaling or reputation-building strategies, as suggested by Goodman (2012).

4.3. Individual Level Analysis

In order to gain a better understanding of bidding behavior, we now turn our attention to individual bidders and to the possible existence of a positive relationship between bidder surplus and experience.[20] We explore the issue in more detail by calculating total auction surplus for each bidder-auction combination and then regress it on several variables of interest using Ordinary Least Squares (OLS). Table III reports the results. The first two columns refer to a generic specification (with and without bidder-specific fixed effects) where a bidder's surplus in an auction may be explained by a quadratic function of experience (measured through the number of previously submitted bids), a quadratic function of the number of simultaneous auctions in which the bidder is participating, a dummy variable that takes on the value of 1 if a particular bidder has won that specific auction and two measures of bidding competition (the number of unique bidders in that auction until the bidder in question decides to quit and an HHI index of those unique bidders' bids).

In both specifications, experience exhibits a broadly negative relationship with surpluses, whereas participation in simultaneous auctions has a broadly positive impact on surplus. Evaluated at the mean, the marginal effect of experience on surplus under both specifications is negative; in addition, it is negative for almost all experience levels under column 1's specification and only becomes positive for very high experience levels (99^{th} percentile). Also, evaluated at the mean, the marginal effect of simultaneous auction participation on surplus is positive; however, it becomes negative for high levels of simultaneous auction participation (4^{th} quartile).

These results are largely consistent with earlier literature and show that experience and simultaneous auction participation induce different bidder responses in the bidding process. Some factors could help explain why more experienced bidders exhibit a different bidding behavior and we look at them under the specifications of the third and fourth columns (with and without bidder-specific fixed effects) of Table III. For example, both Augenblick (2016) and Goodman (2012) identify bidding runs as a potential signaling mechanism through which a bidder signals her commitment and interest in winning the auction. In our data, once bidder-specific fixed effects are considered, such bidding streak episodes do not appear to affect surpluses in a significant way.[21] In addition, we follow Goodman (2012) and introduce a quadratic function of own bid proportion in total auction bids—possibly a signaling device used by more experienced bidders. We find that for 99.7% of bidder-auction observations, own bid proportion has a negative effect on auction surplus.[22] First hour bids may be particularly effective signaling devices, as they allow for reputation effects to be established early in the auction (Goodman, 2012). Therefore, we have also included as potential explanatory variables the number of first hour bids and the proportion of a bidder's first hour bids in total first hour bids and find that similarly to Goodman (2012), first hour bids appear to be a useful and effective signaling mechanism.[23]

20. Augenblick (2016), Goodman (2012) and Wang and Xu (2012) find such a positive relationship.

21. Without fixed effects, bidding streak episodes have a negative impact on surplus, a result that is opposite to that obtained by Augenblick (2016) and Goodman (2012).

22. This result contrasts with that of Goodman (2012), who finds that own bid proportion has a positive impact on surplus for a much lower threshold, that is, bidding frequently and (possibly) in a well-timed manner raises surplus.

23. The number of first hour bids attempts to capture *absolute* reputation signals, whereas the proportion of a bidder's first hour bids attempts to capture *relative* reputation signals. We find that the marginal effect of an additional first hour bid is negative (and larger, in absolute value, than €0.50, the bid cost) but the marginal

TABLE III.
INDIVIDUAL-LEVEL AUCTION SURPLUS REGRESSIONS

Dep. Variable: Individual Auction Surplus	OLS Coef. (Std. Error)	OLS with fixed effects Coef. (Std. Error)	OLS Coef. (Std. Error)	OLS with fixed effects Coef. (Std. Error)
Dummies for auction week	yes	yes	yes	yes
Dummies for type of good	yes	yes	yes	yes
Dummies for brand of good	yes	yes	yes	yes
Individual-specific dummies (fixed effects)	no	yes	no	yes
Number of previously submitted bids	−0.0363 (0.0005)***	−0.0529 (0.001)***	−0.0198 (0.0005)***	−0.0347 (0.001)***
(Number of previously submitted bids)^2	0.000006 (0)***	0.000014 (0)***	0.000002 (0)***	0.000008 (0)***
Number of simultaneous auctions in which bidder is participating	0.085 (0.05)*	0.143 (0.07)**	0.108 (0.04)**	0.165 (0.06)***
(Number of simultaneous auctions in which bidder is participating)^2	−0.013 (0.006)**	−0.023 (0.009)***	−0.002 (0.005)	−0.005 (0.008)
Proportion of own bids in total bids			−761.71 (5.42)***	−732.78 (6.55)***
(Proportion of own bids in total bids)^2			1,200.89 (28.35)***	1,030.99 (35.88)***
Number of bidder first hour bids			−0.678 (0.02)***	−0.668 (0.03)***
Bidder proportion of first hour bids			67.73 (2.32)***	64.13 (2.96)***
Number of bidding streak episodes			−0.002 (0.0007)***	−0.001 (0.0009)
Auction winner dummy	333.99 (0.53)***	335.31 (0.58)***	345.60 (0.5)***	346.17 (0.55)***
HHI of unique bidders' bids until bidder quits	−0.00001 (0.0001)	0.00016 (0.0001)**	0.00027 (0.0001)***	0.00031 (0.0001)***
Number of unique bidders until bidder quits	−0.0021 (0.0001)***	−0.0029 (0.0001)***	−0.0027 (0.0001)***	−0.0033 (0.0002)***
Constant	−1.026 (0.21)***	−2.570 (0.3)***	0.305 (0.2)	−0.710 (0.28)**
N	176,908	176,908	176,908	176,908
R^2	0.70	0.71	0.74	0.74
F-test	14,406	12,086	14,971	12,147

*** Significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Importantly, with the exception of first hour bids, these signaling strategies are rather ineffective. When they are taken into account (third and fourth column of Table III), the marginal effect of experience on surplus is not as high (in absolute value) as in columns 1 and 2, but remains negative for virtually all experience levels. At the same time, the positive marginal effect on surplus of simultaneous auction participation is higher than in columns 1 and 2, further reinforcing the previous result that this is a much more effective learning tool.

OBSERVATION 5: *Simultaneous participation in multiple auctions (horizontal learning) is a more effective learning mechanism than bidding over time within and across auctions.*

4.4. Robustness of Results

4.4.1. Budget Constraints

The participation of bidders in simultaneous auctions raises the possibility that they may be allocating a fixed budget across auctions, in the vein of Colonel Blotto games, a constant-sum game in which two players simultaneously distribute forces across n battlefields and, within each battlefield, the player who has allocated a higher force wins (see Roberson, 2006, or Hart, 2008).[24] There are three notable differences between Colonel Blotto games and penny auctions: (i) in Colonel Blotto games, the chosen strategy has an opportunity cost (the potential benefit of using a force in a different battlefield), whereas in penny auctions, there is a direct bidding cost; (ii) the Colonel Blotto game is a static constant-sum game, whereas penny auctions are dynamic nonconstant sum games; (iii) and in Colonel Blotto games, the final (nonnegative) payoff is the proportion on wins across battlefields (regardless of the force levels allocated to each battlefield), whereas in penny auctions, the final payoff crucially depends on the overall bid costs across all auctions.

Nevertheless, a relevant similarity is that bidders in penny auctions may be budget constrained when choosing their strategies and it is important to test whether that is the case. The data do not allow for a direct test and we have thus followed an indirect approach. First, within each auction week, we assume that budget constraints may affect bidding behavior on a daily basis, that is, the behavior of bidder i in auction day t is affected by her decisions (and total financial expenditure) in day $t-1$. Second, we assume that bidders are not financially constrained across auction weeks. [25] With this framework in mind, we use the bid-level database and calculate, for each bidder and for each auction day in which she was active, bidder-specific and auction day-specific variables (in total, 80,349 active bidder-auction day observations).[26]

effect of a bidder's first hour bid proportion is positive. When the latter is evaluated at the mean, the marginal effect of an additional bid is very close to the bid cost (€0.52). This suggests that provided the proportion of a bidder's first hour bids is relatively high (certainly higher than the mean), there may be an expectation of a net positive marginal effect on surplus of first hour bids (insofar as the expected surplus gain through a first hour bid is higher than the bid cost), that is, first hour bids appear to be a useful and effective signaling mechanism.

24. We would like to thank an associate editor and an anonymous referee for suggesting this line of analysis.

25. This strikes us as sensible, because auctions are typically conducted on a monthly basis; that is, roughly a month goes by between auction weeks.

26. The bidder-specific variables are (i) the total number of bids placed, (ii) the total number of auctions in which she has participated, (iii) the total financial expenditure incurred (which includes bid costs as well as the final auction price if she has won an auction), and (iv) the number of previous auction wins. We have also created lagged variables for (ii) and (iii). The auction day-specific variables are: the total number of auctions, the average retail price of all items to be auctioned during the day, and the number of auctions for each item brand and for each item type.

Our general approach was the following: does a bidder's choice of (i) how many bids to place or (ii) the total number of auctions in which she participates depend on other bidder-specific or auction day-specific variables, including the previous day's total financial expenditure (a proxy for the possible existence of a budget constraint)? To answer this question, we have implemented two specifications. In the first specification, we have expanded our dataset to include all "inactive" auction days that immediately follow an "active" auction day because, if a bidder is financially constrained, it may be that exhausting the available funds in day t makes her not bid in day $t+1$. This yields a total number of 131,089 bidder-auction day observations. Because of the large number of zeros for the dependent variable ((i) or (ii)), we conducted a tobit regression with clustered (by bidder) standard errors. The results are presented as "tobit (i)" and "tobit (ii)" in Table IV. In a second specification, we have only resorted to the active bidder-auction day data and performed a fixed-effects (FE) regression. This is presented as "FE (i)" and "FE (ii)" in Table IV.

When the dependent variable is the total number of auctions in which the bidder participates, the financial expenditure in the previous day yields a positive and significant (albeit small) or insignificant effect, in the tobit (ii) and FE (ii) models, respectively. That is, bidders do not appear to have daily financial constraints in what concerns their choice of how many auctions to participate in. When the dependent variable is the total number of bids placed, the previous day's financial expenditure appears with a positive or negative (significant) coefficient in the tobit (i) and FE (i) models, respectively.[27] The latter is the only model for which the existence of a daily financial constraint could not be rejected, although the marginal impact is rather small.[28]

Observation 6: *Bidder behavior does not seem to be affected by financial constraints.*

4.4.2. Possible differences in bidder strategies

Another possible concern with our results is whether the effect of participating in multiple simultaneous auctions is truly a learning effect. [29] Indeed, different types of bidders may be using different strategies. In order to analyze this possibility, we have classified bidders in one of three possible categories: "simultaneous bidders" are bidders who have participated in more than one simultaneous auction throughout the period we analyze; "seasoned bidders" are bidders who have never participated in more than one auction at a time, but who have either submitted multiple bids in a single auction or a single bid in multiple (nonsimultaneous) auctions; and "single bid bidders" are bidders who have only submitted one bid in one single auction. Table V summarizes some interesting descriptive statistics (averages for each bidder category).

First, most bidders are simultaneous bidders (68%) and account for a very large proportion of all submitted bids (92%). Single bid bidders are relatively few (12%) and seasoned bidders (21%) account for a very small proportion of all submitted bids

27. In a way, the FE specification (indirectly) tests the possible existence of financial constraints *conditional on a bidder being active*, whereas the tobit specification is more general and takes into account the possibility that financial constraints may lead to bidder inactivity (zero values for the dependent variables). If the latter is the true underlying model, then FE would yield biased estimates (as it considers only the uncensored observations) and this can explain the different signs across models for the "total expenditure in the previous day" variable (as well as for other explanatory variables).

28. The average total expenditure in the previous day is €2.25, which means that the respective marginal effect (evaluated at the mean) is $-0.27 \times 2.25 = -0.61$ (the mean number of submitted bids is 11.71). Therefore, evaluated at the mean, a €1 increase in the previous day's total expenditure (a large increase in relative terms) reduces the number of submitted bids by 0.61.

29. We would like to thank a referee for suggesting this line of analysis.

TABLE IV.
INDIRECT TEST FOR THE EXISTENCE OF BUDGET CONSTRAINTS

Dep. Variable: (i)= Total Number of Bids Placed by Each Bidder during an Auction Day; (ii)=Total Number of Auctions in which Bidder is Active during an Auction Day	Tobit (i) Coef. (Std. Error)	Tobit (ii) Coef. (Std. Error)	FE (i) Coef. (Std. Error)	FE (ii) Coef. (Std. Error)
Number of daily auctions for each brand	yes	yes	yes	yes
Number of daily auctions for each type of good	yes	yes	yes	yes
Number of previous auction wins	15.54 (2.17)***	1.11 (0.07)***	9.49 (0.92)***	−0.19 (0.07)***
Total expenditure in previous day	0.46 (0.03)***	0.01 (0.002)***	−0.27 (0.01)***	0.00 (0.001)
Number of auctions where bidder was active in previous day	−3.82 (0.22)***	−0.33 (0.01)***	0.80 (0.07)***	−0.09 (0.005)***
Average retail price of daily auctions	−0.07 (0.003)***	−0.01 (0.0002)***	0.00 (0.003)	0.00 (0.0002)***
Number of daily auctions	−1.92 (0.18)***	−0.25 (0.01)***	2.05 (0.82)**	0.33 (0.06)***
Constant	28.61 (1.45)***	4.95 (0.12)***	9.78 (1.66)***	2.81 (0.12)***
N	131,089	131,089	80,349	80,349
Pseudo-R^2	0.012	0.030	0.005	0.005
F-test	35.5	554.9	46.9	145.9

*** Significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table V.
Descriptive Statistics of Bidder Categories

Variable	Simultaneous Bidders (% of Total)	Seasoned Bidders (% of Total)	Single Bid Bidders (% of Total)	Total
Total number of bidders	25,652 (68%)	7,864 (21%)	4,484 (12%)	38,000
Total number of submitted bids	865,970 (92.0%)	70,514 (7.5%)	4,484 (0.5%)	940,968
Average number of submitted bids in all auctions	33.8	9.0	1.0	24.8
Average number of auctions in which they participate	6.3	1.4	1.0	4.7
Average number of bids per auction	5.1	6.9	1.0	5.0
Average daily expenditure (€)	7.5	3.4	0.5	5.8
Average experience (number of submitted bids) when submitting first bid	11.54	0.97	0.0	7.99
Average percentage of wins in auctions in which they participate	0.12%	0.09%	0.04%	0.10%
Average retail price in auctions in which they participate	456.7	501.4	489.1	469.8
Average number of unique bidders until they decide to enter	345.9	394.5	378.1	359.8
Average percentage of first hour bids	5.7%	7.4%	10.3%	6.6%

(7.5%). Second, simultaneous bidders participate in more auctions than seasoned bidders (6.3 vs. 1.4) but submit relatively fewer bids (5.1 vs. 6.9). Therefore, overall, simultaneous bidders submit a larger number of bids across all auctions than seasoned bidders (34 vs. 9) mainly because they participate in more auctions. Third, simultaneous bidders are much more experienced at the time when they submit their first auction bid and win slightly more often, despite participating in auctions for items with relatively lower retail prices. Fourth, simultaneous bidders submit their first bid at a time when fewer unique bidders have already submitted bids and submit relatively fewer first hour bids.

Leaving aside the case of single-bid bidders, this allows us to conjecture that although there are some differences between simultaneous and seasoned bidders, the main difference appears to be related to the number of auctions in which they participate and their experience when they make the decision to enter. That is, their being active in simultaneous auctions emerges as the most striking difference between bidder types. This suggests that our results may indeed be picking up a learning effect. In order to confirm this conjecture, we ran the pseudo-hazard rates regression presented in Section 4.2.2 for simultaneous bidders only. The results are presented in Table A5 (Appendix A) and indicate that for simultaneous bidders, the sunk cost fallacy is present, but learning effects—both through experience as well as through participation in simultaneous auctions—are rather helpful in mitigating its effects, thus broadly confirming the results presented in Section 4.2.2.

5. Discussion and Conclusion

This paper contributes to the emerging empirical evidence on penny auctions, an auction format that typically raises significantly higher revenue than the underlying market value of the auctioned item. Our particular interest is the effect of learning over time and across auctions as bidders become more experienced. In the dataset we have used, possibly because it is related to a relatively more homogeneous set of products than that used by earlier literature, profit margins are very high: on average, 219%. We find that using auction-level data, the standard supply and demand variables have the expected effects on final auction prices: increased supply (through a higher number of penny auctions occurring simultaneously) reduces final auction prices, whereas increased demand (through a higher number of unique bidders) increases them. In addition, first hour bidding activity contributes positively toward higher auction prices, to the benefit of the auctioneer.

Using bid-level data, we find that more experienced bidders learn to participate less often in penny auctions and, as one would expect, the participation decision is negatively influenced by perceived competition. We also find evidence of the sunk cost fallacy, identified by Augenblick (2016): bidders are more reluctant to drop out the higher is the number of bids they have already submitted. Experience acquired through bids submitted in one or more auctions over time does alleviate the sunk cost fallacy, but participation in multiple simultaneous auctions emerges as a significantly more effective learning mechanism. Therefore, we find evidence of behavioral spillovers, some of which appear to be learning effects. Looking at individual auction surplus confirms these results: experience exhibits a negative relationship with surplus, but participation in simultaneous auctions is associated with higher surpluses. Signaling and reputation-building strategies do appear to be used, but, with the exception of first hour bids, with little success in terms of outcomes.

Although our penny auction data have some limitations—for instance, it does not allow us to use bidder sophistication indicators such as those used by Wang and Xu (2012)—and is certainly smaller than that used by previous authors, it is sufficiently homogeneous to present a coherent picture of bidding behavior in penny auctions. Our most striking finding is that "horizontal learning" (within a time window but through bid submission or participation in more than one simultaneous auction) is a much more effective learning mechanism than "vertical learning" (through bid submission or auction participation over time). This, to the best of our knowledge, is a novel result in the penny auction literature and bears some relationship with the behavioral spillovers literature (Bednar et al., 2012; Cason et al., 2012; Falk et al., 2013). More generally, this result raises relevant questions on behavioral spillovers/learning effects that may occur in other sequential games (e.g., bargaining games). In particular, it raises the possibility that in some games, players may arrive more quickly at equilibrium strategies when playing two identical or similar games simultaneously rather than sequentially, although preliminary evidence by Falk et al. (2013) does not appear to corroborate this claim. In addition, this paper is eminently empirical and our finding of evidence of both "vertical" and "horizontal" learning effects clearly suggests that a theoretical extension to Augenblick's (2016) model that incorporates these two features is warranted. Such a model would yield equilibrium predictions for the role played by each type of learning effect, which could then be reconciled with our empirical results. These are likely to be the next steps in our research.

Appendix A: Auxiliary Tables and Figures

Table A1. Description and Summary Statistics of Main Variables in Auction Level Database

Auction Level Data					
Variable	Number of Observations	Mean	Standard Deviation	Minimum	Maximum
Final auction price	403	24.01	25.25	0.10	167.89
Retail price	403	369.97	198.07	9.90	999.90
Profit (euros)	403	854.77	1191.21	−788.76	7962.49
Profit (percentage of retail price)	403	219.30	249.84	−95.92	1327.30
Average number of simultaneous auctions	403	6.88	1.67	1.80	12.97
Number of unique bidders	403	438.98	333.37	8.00	2019.00
HHI of unique bidders' bids	403	230.72	307.04	28.31	2861.23
Profit percentage of auctions over time					
Auction week 1	61	192.51	270.44		
Auction week 2	46	385.50	302.84		
Auction week 3	44	296.61	264.00		
Auction week 4	48	183.11	202.49		
Auction week 5	48	185.60	205.33		
Auction week 6	47	203.66	195.49		
Auction week 7	109	170.47	234.69		
TOTAL	403	219.30	249.84		
Profit percentage by type of good					
PC	24	292.92	241.76		
Pen	2	−3.63	44.63		
Services	22	84.21	235.47		
Tablet PC	12	350.11	244.51		
Mobile phone	343	219.54	248.32		
TOTAL	403	219.30	249.84		
Profit percentage of mobile phones by brand					
Apple	55	307.81	313.71		
BlackBerry	38	185.21	174.00		
Google	4	352.32	267.62		
HTC	49	110.89	182.18		
Huawei	3	346.12	139.76		
LG	39	99.04	140.06		
Nokia	43	278.79	268.49		
Optimus	44	277.99	242.01		
Samsung	40	240.12	264.60		
Sony	28	205.90	252.79		
TOTAL	343	219.54	248.32		

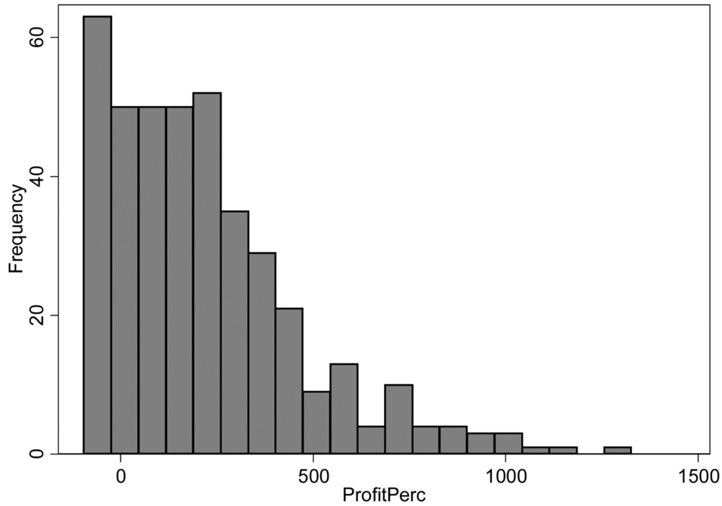

FIGURE A1. PROFIT MARGINS IN PENNY AUCTIONS

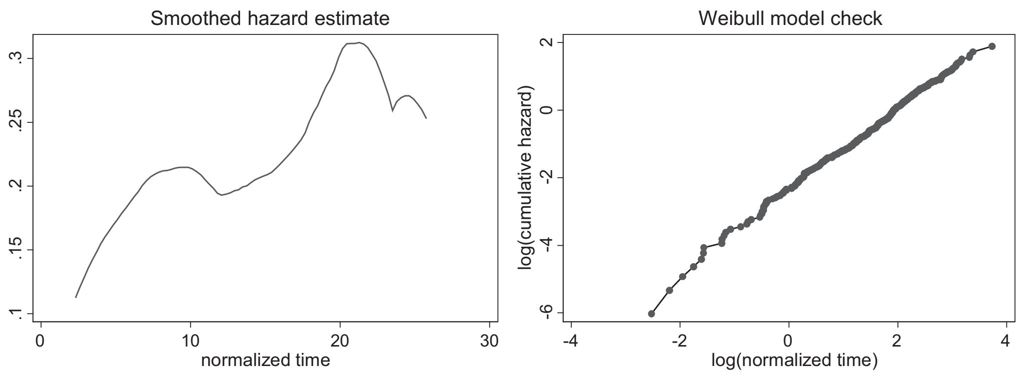

FIGURE A2. EMPIRICAL HAZARD RATE (LEFT) AND GRAPHICAL CHECK OF ADEQUACY OF WEIBULL DISTRIBUTION (RIGHT) FOR FINAL AUCTION PRICES

Table A2.
Description and Summary Statistics of Main Variables in Bid-Level Database

Bid-Level Data					
Variable	Number of Observations	Mean	Standard Deviation	Minimum	Maximum
Total number of bids	940,968				
Expected total number of bids*	967,781				
Coverage	97.2%				
Average number of bids per auction	403	2,334.9	2,405.8	10	16,709
Expected total number of bids per auction*	403	2,401.4	2,525.2	10	16,789
Average coverage	403	98.1%	4.3%	65.1%	100.6%
Total number of unique bidders	38,000				
Average number of individual bids per auction	176,908	5.32	14.12	1	1518

* using data from auction database.

Table A3.
AIC Scores for Final Auction Prices and Underlying Distributions

Distributions	AIC Score
Exponential	948.6
Weibull	596.7
Gompertz	590.2
Lognormal	640.8
Loglogistic	605.6
Generalized Gamma	598.7

Table A4.
Hazard Rate Estimation

	Cox(PH) Coef. (Std. Error)	Gompertz (PH) Coef. (Std. Error)	Weibull (PH) Coef. (Std. Error)
Dummies for auction week	yes	yes	yes
Dummies for type of good	yes	yes	yes
Dummies for brand of good	yes	yes	yes
Average number of simultaneous auctions	1.269 (0.057)***	1.211(0.052)***	1.268 (0.055)***
Number of (normalised) first hour bids	0.121 (0.038)***	0.165 (0.049)***	0.193 (0.054)***
Number of unique bidders	0.994 (0.0004)***	0.995(0.0003)***	0.995 (0.0004)***
HHI of unique bidders' bids	1.001 (0.0002)***	1.001 (0.0002)***	1.001 (0.0002)***
Constant		0.070(0.038)***	0.011 (0.006)***
N	403	403	403
Log-likelihood	−1,710	−268	−271
LR (chi-square)	616	553	541

*** Significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

Table A5.
Pseudo-Hazard Rate Estimation for "Simultaneous Bidders" Only

Dependent Variable: Leave	RE probit Coef. (Std. error)	RE probit Coef. (Std. error)
Dummies for auction week	yes	yes
Dummies for type of good	yes	yes
Dummies for brand of good	yes	yes
Average number of simultaneous auctions	0.02959 (0.00056)***	0.02981 (0.00056)***
First hour bid (dummy)	0.00434 (0.00391)	0.00179 (0.00691)
First hour bid (dummy) × ln(number of previously submitted bids)		−0.00369 (0.00274)
First hour bid (dummy) × ln(number of simultaneous auctions in which bidder is participating)		0.01063 (0.00609)*
Number of unique bidders until bid submitted	0.00012 (0.00001)***	0.00013 (0.00001)***
HHI of unique bidders' bids until bid is submitted	−0.000001 (0)	−0.000002 (0)
ln(number of bids submitted in auction)	−0.11932 (0.00073)***	−0.18247 (0.002)***
ln(number of previously submitted bids)	0.06402 (0.00102)***	0.05624 (0.00106)***
ln(number of simultaneous auctions in which bidder is participating)	0.06467 (0.00135)***	0.05868 (0.00181)***

Continued

Table A5.
Continued

Dependent Variable: Leave	RE probit Coef. (Std. error)	RE probit Coef. (Std. error)
ln(number of previously submitted bids) × ln(number of bids submitted in auction)		0.01455 (0.00045)***
ln(number of simultaneous auctions in which bidder is participating) × ln(number of bids submitted in auction)		0.00664 (0.00092)***
Number of previous auction wins	0.01225 (0.00457)***	−0.00441 (0.00463)
Bid amount	0.00004 (0.00012)	−0.00007 (0.00012)
N	865,970	865,970
Log-likelihood	−351,912	−351,287
Wald (chi-square)	41,042	42,219

*** Significant at the 1% level; ** significant at the 5% level; * significant at the 10% level.

References

Akaike, H., 1974, "A New Look at the Statistical Model Identification," *IEEE Transaction on Automatic Control*, 19, 716–723.

Augenblick, N., 2016, "The Sunk-Cost Fallacy in Penny Auctions," *Review of Economic Studies*, 83, 58–86.

Baxter, J., 1995, "Learning Internal Representations," *Proceedings of the Eighth Annual Conference on Computational Learning Theory*, 311–320.

Bednar, J., Y. Chen, T. Liu, and S. Page, 2012, "Behavioral Spillovers and Cognitive Load in Multiple Games: An Experimental Study," *Games and Economic Behavior*, 74, 12–31.

Byers, J., M. Mitzenmacher, and G. Zervas, 2010, "Information Asymmetries in Pay-Per-Bid Auctions," *Proceedings of the 11th ACM conference on Electronic Commerce (EC '10)*, 1–12.

Caldara, M., 2012, "Bidding Behavior in Pay-to-Bid Auctions: An Experimental Study," Mimeo: University of California, Irvine.

Caruana, R., 1995, "Learning Many Related Tasks at the Same Time with Backpropagation," *Advances in Neural Information Processing Systems*, 7, 656–664.

Cason, T., A. Savikhin, and R. Sheremeta, 2012, "Behavioral Spillovers in Coordination Games," *European Economic Review*, 56, 233–245.

Cleves, M., R. Gutierrez, W. Gould, and Y. Marchenko, 2010, *An Introduction to Survival Analysis Using Stata*, third edition. College Station, TX: Stata Press.

Falk, A., U. Fischbacher, and S. Gächter, 2013, "Living in Two Neighborhoods - Social Interaction Effects in the Laboratory," *Economic Inquiry*, 51, 563–578.

Gombrich, C., 2011, September 22, "Series Learning versus Parallel Learning [Web log comment]," Retrieved from http://www.carlgombrich.org/series-learning-versus-parallel-learning/.

Goodman, J., 2012, "Reputations in bidding fee auctions," Mimeo: Northwestern University.

Hart, S., 2008, "Discrete Colonel Blotto and General Lotto games," *International Journal of Game Theory*, 36, 441–460.

Hinnosaar, T., 2014, "Penny Auctions," Mimeo: Collegio Carlo Alberto.

Jenkins, S., 2005, "Survival Analysis," Unpublished manuscript.

Mertens, J., -F.,1992, "The Small Worlds Axiom for Stable Equilibria," *Games and Economic Behavior*, 4, 553–564.

Platt, B., J. Price, and H. Tappen, 2013, "The Role of Risk Preferences in Pay-to-Bid Auctions," *Management Science*, 59, 2117–2134.

Roberson, B., 2006, "The Colonel Blotto Game," *Economic Theory*, 29, 1–24.

Sigman, M., and S. Dehaene, 2008, "Brain Mechanisms of Serial and Parallel Processing during Dual-Task Performance," *Journal of Neuroscience*, 28, 7585–7598.

Wang, Z., and M. Xu, 2012, "Learning and Strategic Sophistication in Games: Evidence from Penny Auctions on the Internet," Mimeo: Northeastern University.

Wason, P.C., 1960, "On the Failure to Eliminate Hypotheses in a Conceptual Task," *Quarterly Journal of Experimental Psychology*, 12, 129–140.