

Individual bidder behaviour in repeated auctions

Olga Wojciechowska and Michael Waterson¹

University of Warwick

Draft- February 2024

Abstract

We examine bidders' behaviour in auction sales of the iPhone4 on eBay in the context of a significant shortage of the product at listed price, leading to achieved prices significantly above the posted price, on average. We test the direct prediction of the successive auctions model that bidders increase their bids over successive auctions and are influenced by the effects of information gained from previous auctions, finding that bidders indeed react both to their direct experience and to experience gained from studying previous auctions. In addition, the results are suggestive of bidders being reluctant to reveal their true valuation of the product initially but that they do so only over time. Our results are novel in being able to track individual bidders' behaviour rather than simply auction outcomes.

Keywords: Repeated auctions; eBay; consumer valuations.

JEL codes: L63, L81, D12, D44.

¹ Corresponding author Michael.waterson@warwick.ac.uk

1. Introduction

A retail buyer typically need not fully reveal their valuation of an item when purchasing it; all we can say is that they value it, in expectation, at more than the posted retail price. But in an auction, prospective buyers must form a clearer idea of their true valuation. There is then an interesting question of how such ideas are formed and whether they are consistent. Our real-world experiment is designed to answer this question in the context of successive eBay auctions for essentially identical items in short supply.

In such a successive auction context, how do bidders who want one such item behave? A potentially appropriate strategy is to discount their bid in an early auction by the option value of winning in the next period. Hence the prediction that bidders should increase their final bid over sequential auctions in which they participate. Milgrom and Weber (1999) demonstrate this in their General Symmetric Model. However, in terms of the winning bid, this might otherwise be expected to decrease over time, since those bidders with higher valuations will achieve the object first, and then leave. Given independent private values and risk neutrality, Weber (1983) shows that these two potential effects exactly offset each other in the outcome. However, with affiliated values, the winners curse decreases over time, so we should expect more aggressive bids over time. On the other hand, empirical research (Ashenfelter, 1989; Lambson and Thurston, 2006) has found decreasing *sale* prices over successive auctions. Various theoretical explanations have been offered for this (McAfee and Vincent, 1993; Bernhardt and Scoones, 1994; Ginsburgh, 1998 are examples). Clearly there is a need to tease out empirically the effects of *individual* bidders' strategies from the amalgam of aggregate effects in examining the impact on overall achieved sale price in an auction and its path across successive auctions. In pursuit of this, we aim to capture individuals' behaviours in successive auctions for essentially the same object.² To our knowledge,

² There are two main variants, 16GB and 32GB memory, which difference we account for in our analysis.

ours is the first paper to examine individual bidder behaviour for homogeneous products systematically to achieve this goal.³

We use a commonly employed data source to do this, but our sample is unusually appropriate for its purpose. The iPhone 4 was seen as a significant advance on previous models and, at the time, Apple's approach was to engender hype for its new phone model when opening it to a new country market by releasing relatively small numbers of handsets onto that market. In this case, we study the UK market through bidding for the phone on eBay in the immediate aftermath of launch in June 2010; obviously those consumers who did manage to buy the new phone at list price found themselves with an asset worth more at the time than they paid for it. It transpired that there was a very active market on the UK eBay, and in the six weeks of sales we study, covering a significant period of the excess demand, transactions amounting to £1.5m went through the site with up to 200 sales per day. Many participants bid in significant numbers of auctions before winning. With a high degree of confidence, we are able to recover the final bids of individuals over different auctions, so as to track their behaviour- indeed our approach to doing this may be of interest more generally.⁴ It focuses on the make-up of the demand curve rather than the aggregate demand curve itself, revealing behavioural traits that are not apparent in the latter.

Our findings, to a significant extent, match empirical behaviour to theoretical predictions and thus break what might be seen as a mismatch between theory and empirics. We test first whether bidders increase their own bids in subsequent auctions and second, explore whether past achieved prices influence bidding. To do this, we take carefully selected subsets of our sample.

³ An interesting recent paper (Anudsen et al., 2022) analyses bids in housing auctions, including bidder behaviour across successive auctions after they have been unsuccessful, and finds bidders increase their bids over time. However, houses vary over many dimensions. There are of course also laboratory experiments on repeated auctions- see for example Kagel and Levin, (1987)

⁴ A note on terminology here: Many people bid more than once in an auction. We define their final bid as the last bid they placed in that auction, subscripted t . This should be distinguished from their final observed bid, the last bid of theirs that we observe in the dataset, which we designate by the subscript T .

The data includes all iPhone 4 sales on eBay between 17th June 2010 and 7th August 2010. It is a unique source for the analysis of bidder behaviour in auctions in many ways. Most importantly it contains data on over 2000 auctions and over 3000 bidders, for essentially the same product over this period. Secondly, the site contains not only the winning bid but also all other bids in each auction, which means the analysis does not have to be limited to final auction prices. Thirdly, the data was collected at the time of the shortage of the iPhone 4, immediately after its first introduction to the UK market. There is significant interest in these organic auctions, which differ from the relatively limited participation in “auctions” purposefully set up for experimental research papers.

Though the dataset comes directly from Waterson and Doyle (2012), the question asked here is quite distinct. In their paper, the aim was to explain the final price paid by the winner by reference to the product’s characteristics. By contrast, we examine the behaviour of individual bidders across auctions to unearth their apparent strategies.

Because there was uncertainty about availability of the smartphone in physical retailers, the prices achieved were commonly significantly higher than the Apple list price that pertained after the supply shortage was resolved, reflecting high valuations for the object among bidders. The temporary supply shortage is also one of the reasons for the exceptionally high interest in the online auctions for this device at the time (Waterson and Doyle, 2012).

On the eBay site, usernames of bidders are partially encoded, so the full usernames cannot be unambiguously identified. We address this issue in order to identify individual bidders taking part in more than one auction. The fact that usernames were matched across auctions for the first time makes the generated dataset unique. Our approach to this is discussed in Section 2 and our methods discussed in Section 3. We then test the two main questions, on bidders’ own bids across successive auctions and on their reaction to achieved winning bids in previous auctions, in the two subsequent sections, before concluding briefly in Section 5.

2. Data collection and analysis

In practice, because of the intensive activity on the eBay site for the iPhone 4, two separate assistants were engaged to collect data. They employed different methods of data collection, with duplications and discrepancies resolved subsequently in cleaning the data. One researcher used manual collection (facilitated by the “watch” tool), the other developed a web crawler for data capture. The data was originally used in Waterson and Doyle (2012) which gives more detail. It constitutes a large and unique dataset. To give an indication of its size, we collected over 27,000 bids across over 2500 auctions within our 6-week period. After cleaning to remove problems with some of the observations, such as potential scam attempts either by buyers or sellers (identified as users “no longer registered”) we made use of data from 2393 completed auctions in total. However, we restrict the sample in different ways in order to engage in testing for individuals’ behaviours.

There are several issues concerning data collection and generation to be explained. The first and most obvious is to identify individual bidders accurately. Buyers’ usernames on eBay are partially anonymized for privacy reasons. The format adopted on eBay is a string containing the first character of the username followed by number of stars (*), which cover the middle part, and followed by the last character, then the number of current wins in brackets. As an example, user entry “a***s (19)” means that the first letter of the username was “a”, the last “s” and that the person had made a total of 19 purchases on eBay by that date. In the case that the user did not buy any new product over the duration of data collection, this encoding would give in fact almost 100% certainty that each distinct entry related to a different person. The data collection took place over 44 days, so it is possible that additional purchases were made over that period. Moreover, some buyers can win more than one product over the dataset duration or continue bidding in other auctions after winning a phone, therefore reasonable increases in the number of won auctions are possible, but not decreases. This information is used in the algorithm to identify unique users. The choice of usernames on eBay allows using any letter, lower or upper case, numbers, as well as special characters, which include full stops, asterisks, underscores, and dashes. Usernames must contain at least 6 characters. A username is a

unique identifier of a person and is automatically assigned by eBay unless the user changes it.

No matter what the total length of the username, just by knowing the first and last characters gives the number of unique permutations as $66^2 = 4356$ (66 is the total number of possible characters used), which means that the probability of randomly picking two identical pairs of characters is $1/66^2 = 2.296 * 10^{-4}$. This assumes that usernames are chosen randomly, that is that people do not modify their username from that automatically generated by eBay. We also have the number of total wins on eBay. This gives an additional means by which to distinguish users, in cases of more than one username with the same first and last character. Two extreme approaches are: 1) treating the users as the same whenever the first and last letters are the same, or 2) only when both first and last letter *and* the number of wins is the same.

We sort the users first by the total number of wins, and then by time. The variable *user* is created by requiring that for each username, if the number of wins decreases but clock time has increased, then the following entry is presumed to be a new user.⁵ This approach yields 3994 unique users out of over 6000 possible unique eBay users.

A second issue is the need to capture what a bidder sees in placing a bid on eBay. During the auction, when the bidder places their bid, the bid itself is not visible to other users, only the “live bid”, which is the second highest bid plus the required increment. The magnitude of the increment, which changes depending on bid amount, is publicly known and available on the auctioneer’s website. It is, therefore, possible to retrieve live bids from bids data, and it has been done for example in Jank and Shmueli (2010). The R script used is available on the book’s website, and it has been utilized here in order to retrieve the live bids (with necessary adjustments to the dataset). Auction bids and live bids can be quite different. Most importantly, while it is possible that a lower bid is placed after a higher one by another user, live bids are always monotone increasing.

⁵ We also engaged in some robustness checking around this point, but the set identified as individual users differs little. For more detail, see Wojciechowska (2018).

Live bids show, in fact, what the highest bidder would pay if the auction ended at the given time, but do not reveal the current highest bid. It has to be also noted that while the code provided by Jank and Shmueli (2010) recreates the live bids, information available to the bidder at the time of bidding is the current live price - which is the live bid just before the new bid is placed - precisely the live bid of previous bid.

In eBay terminology the ability to place high bids in advance is called “proxy bids”, and live bids are related to as “bids” - even though these “bids” are automatically made by the system on behalf of the bidder. Interestingly at the time the bidder places their bid they are not able to know what their live bid will be or whether they will even become the highest bidder. If we use the terminology conventionally used in auction theory, the proxy bids should be called bids, while eBay is a type of second price auction (but with a hard close), where the highest bidder pays the second highest bid with an added fixed increment specified by the rules. Over the auction duration eBay keeps track of the second highest bid (plus the increment) and gives this information to intending auction participants and observers. This information may be partially revealing about the valuation of some of the bidders in the auction, and it can influence the bidding strategy. Due to the possibility of multiple bidding, as well as the fact the auctions in this case last one or three days in practice, the actual information about the valuations revealed by these live bids is very limited. It does, though, have a significant impact on the bids placed. Jank and Shmueli (2010) show that the information on live prices and time alone can be used for prediction of final price in an auction, since there are patterns of how these prices evolve over the course of the auction. The mechanism under which the auction evolves is discussed by means of a brief example in a short appendix to the paper.

3. Testing Method

Essentially, there are two different approaches to testing our hypotheses. One involves tracking successful individuals (in that they obtained an iPhone4) across successive auctions and comparing *their bids* in these auctions, the question being whether they raise their bids across these and if so, whether it is a continuing effect. In other words, we test the hypothesis coming most directly from Milgrom and Weber (1999). This is

a clean but relatively limited approach to examine the personal effect. The second is to take a larger sample and investigate the role of past *achieved* prices in previous auctions on a bidder's behaviour, examining an additional question- do people modify their bids based on seeing what *other bidders* do as well as their own past behaviour?

Specific cuts of our overall sample frame are needed to carry out these analyses. For the first approach, we use the final bids of bidders until their first win, with the proviso that they must have participated in auctions *strictly* sequential in time, i.e. not including individuals (who may be professionals) who have several simultaneous bids in play. We then sort into bidder sets who participated in two, three and four auctions before having a win in the final period for the same item.

For the second approach, we model the final bid of a bidder in an auction as a function of the winning prices in past auctions in which they have participated or, in a variant, for the same object. We use controls to correct for anything in the detailed characteristics changing between auctions. The sample is limited to participants in at least three sequential auctions in the case where we consider their own bids as an influence. The reasoning is as follows: In the first recorded auction for the phone participated in by this individual,⁶ the individual learns the auction's final achieved price, sent to them by eBay. It is the individual's final bid in the penultimate auction which forms our dependent variable. Different individuals will have different valuations, of course, so we also include in the regressions a *proxy* for the individual's valuation (and bid), namely the price they paid for their successful purchase. In the alternative where we examine the effect of previous auction prices on the bidder, they need not have participated in these, rendering the sample significantly larger.

4. Do successful bidders raise bids over successive auctions?

Hypothesis 1a: Bids are increasing with auction number in sequential auctions.

⁶ Given that our sample collection starts on the release date of the iPhone4 in the UK, this is almost certainly the first auction for the phone that they participated in.

Hypothesis 1b: Bidders with a higher valuation discount their bids by more than bidders with lower valuation.

Table 1 shows tests of these hypotheses. The bidder's final bid (or successful payment) at time t is written Pb_{it} . In order to allow for the limited degree to which the products may not be homogeneous between successive auctions, we include controls for whether the item is unlocked to all networks, whether it is locked on the O2 network or another network, and whether it has the same capacity in Gigabytes or not.⁷ Bidders who entered bids twice (i.e. won with their final bid on the second occasion) who had a higher valuation clearly discounted their first bid by more than did bidders with a lower valuation. Key distinguishing characteristics did not appear to affect this discounting on average, save that the gap between bids was smaller, other things equal, when both items bid on were unlocked. However, whilst in absolute terms they discounted their bid by more, they did not discount their bid proportionately more than bidders with a lower valuation (not shown in the table). For those we observe bidding three times, the difference between their successful bid and their second bid was significantly related to the size of their final bid, but no item characteristics appear to be important. The difference between their second bid and their first bid was also significantly related to the size of their second bid, but nothing else. In cases where we observe four bids, the final bid being successful, a similar pattern emerges, although the relationship between the difference between the second and first bid on the second bid is significant only at the 10% level.

Nevertheless, it is important to note that in the case of the final, successful, auction the price paid does not represent the amount actually bid but is non-strictly less than that, because the winner pays the second highest bid in the auction plus the increment. Therefore, the test is conservative as regards the penultimate to final auction, since on unsuccessful occasions we observe the actual amount bid.

⁷ The previous version of the phone, the iPhone3, featured a period of exclusivity to operator O2, so many buyers seeking to upgrade would be as interested in a phone locked to O2 as a phone which was unlocked. We do not include a control for whether the phone was listed as "New", since in the few cases where they are not described as such, this may be because they have simply been taken out of the box to photograph, and all must be at least nearly new.

The general message we take from this is that hypothesis 1 is confirmed, and that bidders are more successful if they are more willing to raise their bids after failure to obtain the item, the more so if they raise them significantly. Bid discounting increases with the final payment, and therefore with valuation - as predicted by the view that economic agents are forward-looking in sequential games. But it does raise the question of how bidders determine their valuation for the phone, or at least how willing or reluctant they are to approach their valuation.

5. Learning from past achieved prices

The dataset contains final bids in each auction by bidders. Bidders can learn from outcomes of previous auctions either through personal experience, where there are common bidders across auctions in which they participate, or through observing a number of past achieved prices for similar objects, without necessarily participating. Both these pieces of information are readily available from the eBay site. Accordingly, we consider both the achieved price in the last auction the bidder participated in and alternatively the average achieved price over the past N auctions for the object, where $N = 5, 10, 20$.

In the regressions where we include bidders' own previous experience, we only include bidders who have participated in at least three auctions, as explained in Section 3 above. As a result, we are left with 2515 observations on bidders who bid in three or more auctions, 1753 of which auctions are one-day. Within this set, there are significant numbers bidding in up to five auctions. It is natural to view these from the perspective of the final bid, back to the earlier bids, rather than in historical timing.

The slice of the dataset for this test contains 12063 final bids, coming from 2536 auctions

with the unit of observation being up to 3867 final bids.⁸ Here we take the approach of restricting the sample to render it more homogeneous and then controlling for remaining key characteristics. We remove sellers with a history which is insufficient to qualify for a “star”, also bidders who have since left, offers to sell which do not have photographs and auction durations apart from the most common one or three days. We also drop cases where, rather than an auction, the phone is offered under “buy it now”. We control for phone type (16 or 32 GB memory), whether it is unlocked to any network or confined to the O2 network.⁹

It is a maintained assumption that there are some common bidders across auctions, in the sense that if the set of bidders was to change completely, there would be no guide as to the distribution of valuations of the other bidders. Therefore, we check whether bidders carry across to some extent. There are two ways of doing this. One is to consider the set of recent previous auctions for the object, the other to consider the set of previous auctions in which this particular player bid. The proportions are small, but using either method, we find some overlap between successive auctions.

Hypothesis 2a: Previous auction outcomes in auctions in which the bidder participated affect the bids a bidder makes.

Hypothesis 2b: The outcome prices of recent auctions affect the bids a bidder makes.

The proposed estimation equation is given in (1) below, where the dependent variable is the individual’s final bid in the auction, while the treatment variable is the price paid by the winner in period $t - 1$. The hypothesis posed is that the amount of the bid is not independent from the results of previous auctions, either those in which the bidder participated (2a) or relevant recent auctions (2b). Written succinctly, the model is as follows:

$$Pb_{it} = P_{t-1} + \log(Pb_{iT}) + C_t + C_a + f(t) + E_{iat} \quad (1)$$

⁸ As we have seen, bidders may bid in more than two auctions, so this number can (and does) exceed the number of separate bidders. In this test we do not require strictly sequential bids by bidders.

⁹ There were few auctions lasting more than three days. Also, at the time, “buy it now” was uncommon for phone auctions.

where the dependent variable, Pb_{it} , is the final bid of bidder i in the auction in question, at time t . The treatment variable is the achieved winning price of a previous auction in $t-1$, denoted P_{t-1} in the above equation. In practice, we adopt two alternative definitions of this, first and most directly (Hypothesis 2a) the achieved price in the penultimate auction in which the bidder participated (Pb_{t-1}), and second (Hypothesis 2b) the (mean) price in recent auction(s) for the same object- the most recent, the most recent five and the most recent ten. The latter hypothesis allows us to make use of a considerably larger sample.¹⁰ There are numerically more overlaps across participating bidders as we add more auctions. We also have Pb_{iT} , our proxy for the bidder's valuation, which is the final bid we observe for this bidder amongst those we observe for the bidder in subsequent auctions. Finally, we include control variables relating to the particular auction in which the bid was placed, C_a , - for example the seller rating, model etc., also controls relating to the time at which the final bid has been placed, C_t , namely the current number of bids or bidders in the auction, and the time in the auction - whether the bid is close to the start or end of the auction, and time. E_{iat} is the normal error term.

So far as functional form is concerned, we adopt a pragmatic approach in which the bid the bidder makes and the previous outcome bid are used in linear form as having a direct linear impact but those variables that have more nearly lognormal distributions are used in logarithmic form, focusing on parsimony of representation.¹¹ Auctions in which alternative functional forms including interactions between variables produce very similar results and are available from the authors on request.

In respect of testing Hypothesis 2a, where we use the achieved price in the penultimate auction in which the bidder participated as the value for P_{t-1} , the data used contains only

¹⁰ There are many bidders who joined more than one auction and therefore there is more than one final bid by them in an auction represented in the data used to test hypothesis 2b.

¹¹ Note that when the bidder is successful (auction at time T) the price paid does not, unlike their bid in previous unsuccessful attempts, represent their valuation but instead a proxy of their valuation, the bid they beat plus the increment. The relation between the two is likely to be nonlinear, so there is no oddity in using the log form price for their period T.

the bidders who took part in at least three auctions. If we use price(s) from recent auction(s) then it is not crucial that the bidder participated in three auctions. There is also the question of whether we should focus narrowly on a particular length of auction, in addition to controlling for the key differences in item across auctions. In presenting results, Table 3 shows three representative regressions, one where a one-day duration definition is adopted, another where both one-day and three-day auctions are used and the last where the five most recent auctions are used instead of the immediately previous auction the bidder participated in.¹² The penalty for adopting the narrowest definition is a smaller sample size, of course.

Examining Table 2, we observe that the bidder's valuation proxy (Pb_T), has a substantial positive impact on the final bid the bidder makes in the observed auction ~~bid~~ later, implying that a higher valuation leads to a higher final bid in the auction under examination. Using either the achieved price in the penultimate auction in which the bidder participated, or the mean achieved price across the five most recent auctions for the object, shows that these influence the final bid the bidder makes in this auction, suggesting that indeed the bidder does learn from observations of previous auctions. These observations may come either from the bidder's own experience in the previous auction (regressions 1 and 2) or from examining results from similar previous auctions (regression 3) and drawing on this experience, in line with hypotheses 2a and 2b.

In terms of key controls, when the bidder has participated in more auctions this has an ambiguous effect on the final bid, which is surprising. However, the remainder of the signs make sense: the greater the competition the lower the bid made, the higher the prevailing price at time of bid, the higher the bid made, whereas the more bids the bidder makes in the course of the auction (where their own bidding behaviour is relevant), the lower the bid made. These results suggest that bidders are attracted by a relatively low price to participate and then drop out later in the face of competition or are nervous about raising their bid. Surprisingly, if anything the later the time of the final bid within the auction, the lower the bid all else equal, although this variable is not precisely

¹²¹² We also tried last 10 and last 20 auctions, with similar results.

estimated in the first two examples.¹³ Unsurprisingly, the 32GB model commands a premium, as (not shown) does an unlocked model- these latter results mirror those in Waterson and Doyle (2012). In addition, results not shown here suggest that bidders for this more powerful model behave in a more sophisticated manner than bidders for the 16GB model, perhaps because they are more experienced and more intensive phone users. We achieve an impressive overall explanation of their bidding behaviour in terms of R-squared.

6. Conclusions

We introduced the paper by pointing to a need to tease out empirically the effects on individual bidders' strategies where there are many auctions for the same or a very similar object. This is important because simply observing auction outcomes mixes bidder behaviour, actions of other bidders, and seller behaviour. By identifying the behaviour of individual bidders in a context where there are many alternative/sequential auctions taking place for the same object, we find that bidders do indeed modify their strategy in response to observing previous auctions, bidding more aggressively in later auctions.

One clear message from tracking individual successful bidders is that they do raise their bids over successive auctions, as would be predicted in the case of a shortfall between the number of bidders and the number of objects (Milgrom and Weber, 1999). Indeed, so much so that they are seemingly reluctant to express their true valuations in their early bids- the large increase in bids over subsequent auctions would not otherwise occur. Individuals who are successful in obtaining a phone generally start with a much lower bid than that which yields success. This is in the context where in the aggregate, Waterson and Doyle (2012) found, using the same set of data, a mild (and declining) increase in achieved prices over the period studied. Therefore, individual behaviour in bidding cannot be directly inferred from aggregate outcomes, which points to a clear

¹³ We should recall here a feature of eBay auctions. Because of the hard close, there is often a flurry of bids in the final minutes of the auction. This is suggestive of the lognormality in the data observed for this and other variables.

gain in understanding through tracking individual consumer behaviour.

These insights persist in the broader samples represented in testing hypothesis 2. Again, bidders, whether successful or not, are positively influenced by their valuation, as represented by the proxy. A second point is that there appears from table 2 to be some evidence of affiliated values or learning in bidders' behaviour- bids are increased as a result of learning about achieved prices in auctions either in which they participated or by study of previous auctions. Additionally, there is evidence of rationality in that higher bids are made for the more powerful model, for example.

Overall, these results suggest a rather more powerful conclusion- contrary to a common assumption that bidders will bid their true valuation in a second-price auction which is one of a sequence, they underbid in early auctions and then raise their bids in order to improve success, as predicted by sequential auction theory, but additionally guided by their experience. We believe this is the first study to demonstrate this for an essentially homogeneous good.

Tables

Table 1: Tracking successful bidders over time.

Dependent variable, bid increment $Pb_{it} - Pb_{i,t-1}$ in each case.

Bid numbers	Bid	Coefficients; standard error						R sq
		Bid, Pb_{it}	Increased GB (16 to32)	Decreased GB	Both unlocked	Same firm lock	O2	
2 (244 obs)	Final	0.941	46.25	-53.1	-121.6	-54.48	-67.84	0.22
	(se)	<i>0.148</i>	<i>41.92</i>	<i>38.35</i>	<i>51.63</i>	<i>48.51</i>	<i>41.71</i>	
3 (71 obs)	Final	1.036	25.91	-36.01	-46.05	-81.11	69.33	0.23
	(se)	<i>0.306</i>	<i>86.81</i>	<i>83.80</i>	<i>144.5</i>	<i>129.3</i>	<i>143.8</i>	
	Penultimate	0.700	40.67	-147.1	-69.41	39.49	-131.5	0.38
	(se)	<i>0.133</i>	<i>75.67</i>	<i>80.29</i>	<i>157.2</i>	<i>153.3</i>	<i>171.0</i>	
4 (23 obs)	Final	0.748	-73.44	-168.6	-190.9	114.5	-124.5	0.46
	(se)	<i>0.304</i>	<i>109.3</i>	<i>110.5</i>	<i>172.9</i>	<i>166.0</i>	<i>203.3</i>	
	Penultimate	1.413	-140.4	134.0	-366.6	249.5	162.2	0.53
	(se)	<i>0.385</i>	<i>148.3</i>	<i>129.2</i>	<i>196.6</i>	<i>176.0</i>	<i>282.4</i>	
	Pre-penult	0.354	-29.66	71.19	13.48	-86.97	No case	0.38
	(se)	<i>0.182</i>	<i>138.9</i>	<i>111.1</i>	<i>112.6</i>	<i>101.8</i>		

Table 2: Explaining the bidder's final bid in a particular auction.

Regressions to explain bidder's final bid in an auction, Pb_{it}

Days duration	1+3	1	1+3
Price in the immediately previous auction bidder participated in (P_{t-1})	0.04	0.06	
Standard error	<i>0.016</i>	<i>0.019</i>	
Mean of last five auction prices			0.08
			<i>0.026</i>
Proxy for valuation (log) (bidder's final observed final bid, Pb_{iT})	17.44	23.37	14.15
	<i>1.71</i>	<i>1.80</i>	<i>1.36</i>
Count of auctions joined by the user (log)	-15.82	14.23	
	<i>42.93</i>	<i>2.24</i>	
Percent of time passed in auction (log)	-11.42	-4.97	-22.15
	<i>6.62</i>	<i>2.77</i>	<i>5.58</i>
Bid number in the auction sequence (log)	16.17	44.55	13.85
	<i>6.30</i>	<i>3.62</i>	<i>5.45</i>
Number of other bidders in auction	-11.32	-10.31	-11.48
	<i>2.02</i>	<i>0.80</i>	<i>1.68</i>
Prevailing price at time of bid	0.66	0.65	0.71
	<i>0.01</i>	<i>0.01</i>	<i>0.01</i>
Count of user's bids in the auction (log)	-43.76	-46.40	-41.03
	<i>3.54</i>	<i>4.32</i>	<i>2.78</i>
Control for 32GB model	37.59	38.35	33.83
	<i>3.80</i>	<i>4.50</i>	<i>3.09</i>
Other fixed effects include Network unlocked (+sig) and interaction effects			
R squared	0.871	0.851	0.874
Number of observations	2515	1753	3867

Appendix: Explaining the sequencing in the course of an eBay auction

To take an example at random, in 2020 an iPhone 6s (not in short supply) was in an auction, with the price, after 17 bids, set at £96. On the website, if I wished to place a bid, I needed to make it at least £98, the site tells me. If I chose £100, the current bid would be set to £100, if I was the highest bidder. However, another participant may have entered a “proxy bid” of £150, in which case when I enter £100, I will be told that I am outbid at the price of £100 and if I decline to do so, the price on the site moves to £100 in favour of the proxy bidder and the only effect of my bid is to raise the current revealed price. The value of the proxy bid is not revealed unless my bid happens to be above it. In fact, earlier in the auction, a relatively inexperienced bidder had placed bids in succession of £48, £51, £53 and £59, only to be trumped by a proxy bid placed two days earlier at £60. The inexperienced bidder then clearly, from the site, placed a proxy bid at £80, baulking another inexperienced bidder who later placed three successive bids below that in vain, before dropping out. The phone sold for £118, but this does not necessarily represent the maximum amount bid of course.

References

- Andre K. Anundsen, Andreas Erling, Røed Larsen and Mikkel Walbækken (2022). Bidding behaviour in housing auctions. Oslo Metropolitan University Working Paper 22-1.
- Orley Ashenfelter (1989). How Auctions Work for Wine and Art. *The Journal of Economic Perspectives*, 3(3):23–36.
- Dan Bernhardt and David Scoones (1994). A Note on Sequential Auctions. *American Economic Review*, 84(3):653–657.
- Victor Ginsburgh (1998). Absentee Bidders and the Declining Price Anomaly in Wine Auctions. *Journal of Political Economy*, 106, 1302–1319.
- Wolfgang Jank and Galit Shmueli (2010). *Modeling online auctions*. Wiley, Oxford.
- Jon H. Kagel, Ronald M. Harstad and Dan Levin (1987). Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study. *Econometrica* 55(6), 1275-1304.
- Paul Klemperer (1999). Auction Theory: A Guide to the Literature. *Journal of Economic Surveys*, 13(3):227–286
- Val E Lambson (2006). Sequential auctions: theory and evidence from the Seattle Fur Exchange. *RAND Journal of Economics*, 37(1), 70-80.
- R Preston McAfee and Daniel Vincent (1993). The Declining Price Anomaly. *Journal of Economic Theory*, 60(1):191–212.
- Paul R Milgrom and Robert J Weber (1999). A theory of auctions and competitive bidding, II. *The Economic Theory of Auctions*, ed. by P. Klemperer, Cheltenham, U.K.: Edward Elgar.
- M Waterson and C Doyle (2012). Your Call: eBay and Demand for the iPhone 4. *International Journal of the Economics of Business*, 19(1):141–152.
- Weber, Robert J. (1983). Multiple-Object Auctions, in R. Engelbrecht-Wiggans, M. Shubik, and R. M. Stark, eds., *Auctions, bidding, and contracting: Uses and theory*. New York: New York University Press.
- Olga Wojciechowska (2018) *Online Auctions- Examination of Bidders' Strategies: Theory and Evidence*. Unpublished PhD Thesis, Coventry: University of Warwick.