Loss aversion and learning to bid

Dennis A.V. Dittrich Werner Güth
Martin G. Kocher Paul Pezanis-Christou

August 2011

Bidding challenges learning theories; even with the same bid, experiences vary stochastically: the same choice can result in either a gain or a loss. In such an environment the question arises how the nearly universally documented phenomenon of loss aversion affects the adaptive dynamics. We analyse the impact of loss aversion in a simple auction using the experienced-weighted attraction model of learning. Our experimental results suggest that individual learning dynamics are highly heterogeneous and affected by loss aversion to different degrees. Apart from that, the experiment shows that loss aversion is not specific to rare decision making.

JEL Classification C91, D44, D83.

Keywords loss aversion, bidding, auction, experiment, EWA learning

There are very rare situations where bidding games – usually referred to as auctions – are dominance solvable, like in the random price mechanism (see Becker, de Groot, and Marshak, 1964) or in second-price auctions (see Gandenberger, 1961, and Vickrey, 1961). Apart from such special setups, bidding poses quite a challenge for learning. According to the usual first-price rule (the winner is the highest bidder, and she pays her bid to buy the object), one has to

---

D.A.V. Dittrich: Touro College Berlin, Am Rupenhorn 5, 14055 Berlin, Germany; davd@economicscience.net
W. Güth: Max Planck Institute of Economics, Jena; gueth@econ.mpg.de
M.G. Kocher: University of Munich; martin.kocher@lrz.uni-muenchen.de
P. Pezanis-Christou: School of Economics, University of New South Wales, Sydney 2052, Australia, p.pezanis-christou@unsw.edu.au
Published as Dittrich, Dennis A. V.; Werner Güth; Martin G. Kocher; and Paul Pezanis-Christou (2012); *Loss Aversion and Learning to Bid*. Economica 79, 226-257.

underbid one's own value (bid shading) to guarantee a positive profit in case of buying. Such underbidding may cause feelings of regret or loss when the object has not been bought, although its price was lower than one's own value. Similarly, even upon winning the auction, feelings of loss may be evoked since one could have won more by a lower bid. When the values of the interacting bidders are stochastic, whether one experiences a loss or a gain is highly stochastic.

In this paper we investigate two important issues. First, we are going to reassess the issue of learning in auctions. Although bidding experiments are numerous in economics (see Kagel, 1995, for an early survey), hardly any of them provides as much opportunity for learning as we do (e.g., Garvin and Kagel, 1994; Selten and Buchta, 1998; Armantier, 2002, and Güth et al., 2003). We believe that the number of repetitions in most studies has been too low to fully account for possible learning dynamics. In stochastic environments, learning is typically slow and the number of repetitions or past bidding experiences are important dimensions. Our one-bidder contest (Bazerman and Samuelson, 1983; Ball et al., 1991) allows for a considerable number of repetitions. To directly identify the path dependence of bidding behaviour (i.e., why and how bidding behaviour adjusts in the light of past results) without any confounding factors, we deliberately choose a bidding task that lacks strategic interaction.

We will analyse the learning dynamics by estimating the parameters of several variants of Ho et al.'s (2007) Experienced Weighted Attraction (EWA) learning model. The EWA learning model is a hybrid of reinforcement and belief models that uses information about forgone payoffs as well as past choice behavior. Either information would be ignored by pure reinforcement or pure belief learning. Estimating best-fitting parameter values for the EWA learning model thus effectively allows to compare a large number of different learning models in one go. We are interested in the individual behaviour which we expect to be heterogeneous with respect to the underlying parameters of the EWA model. Consequently, we are going to estimate the parameters for each participant separately and discuss the distribution of estimated parameters.

Second, we experimentally test for the impact of loss aversion (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) on the learning dynamics in this bidding task by incorporating a loss parameter in the EWA learning model. Loss aversion is referred to as the behavioural tendency of individuals to weigh losses more heavily than gains; there is ample evidence for loss aversion in both risky and risk-free environments (see Starmer, 2000, for an overview). In our setup, we avoid the possible ambiguity of the reference point implied by loss aversion: losses are monetary losses compared to no-trading. Obviously, in a first-price auction all possible choices can prove to be a success (yield a gain) or a failure (imply a loss), at least in retrospect. Therefore, the usual and robust finding of loss aversion should be reflected in the adaptation dynamics.

More specifically, we distinguish losses from bidding and hypothetical, i.e., non-realized, retrospective gains from not bidding as both shape the future attraction of bidding. In doing so we can reasonably exclude idiosyncratic risk attitudes since due to cumulative payments (over 500 rounds) the variance of total earnings should be close to zero. Our main hypothesis on the impact of loss aversion claims that, in absolute terms, an actual loss will change bidding dispositions more than an equally large gain. To the best of our knowledge, loss aversion has so far received little attention (one example, considering only hypothetical losses, is Ockenfels and Selten, 2005) when specifying learning dynamics in situations comparable to ours. Note that – since risk aversion might disappear when playing many rounds due to diversification – one could expect that the so far documented confirmation of loss aversion might also be specific to rare decision making.

In the basic bidding task (originally studied experimentally by Samuelson and Bazerman, 1985, and more recently by Selten et al., 2002, and Charness and Levin, 2009) to which subjects are exposed in our experiment, the only bidder and the seller have perfectly correlated values. Put differently, the seller's valuation is always the same constant and proper share of the bidder's value. Whereas the seller knows her value, the bidder only knows how it is randomly generated.

The bid of the potential buyer determines the price if the asset is sold. In our experiment (like in earlier studies) the seller is captured by a robot strategy accepting only bids exceeding her evaluation. Thus, the bidder has to anticipate that, whenever her bid is accepted, it must have exceeded the seller's valuation. Neglecting this fact may cause the winner's curse as in standard common-value auctions involving more than one bidder.

If the seller's share or quota in the bidder's value is high, i.e., when the positive surplus is relatively low, such a situation turns out to be a social dilemma: According to the solution under risk neutrality[1], the bidder abstains from bidding in spite of the positive surplus for all possible values. However, if the quota is low, the surplus is fully exploited: the optimal bid exceeds the highest valuation of the seller and thereby guarantees trade. Earlier studies have only focused on the former possibility. In our experiment each participant, as bidder, repeatedly experiences both low and high quotas. It is straightforward that the bidding tasks participants are exposed to constitute an ideal environment to study learning because other, possibly confounding behavioural determinants such as the effects of social preferences are nonexistent.

When bids can vary continuously, it is rather tricky to explore how bidding behaviour is adjusted in the light of past results. Traditional learning models like reinforcement learning, also referred to as stimulus-response dynamics or law of effect (see Bush and Mosteller, 1955), cannot readily be applied. Therefore, we offered participants only a binary choice, namely to abstain from or to engage in bidding. In order to observe directly the inclination to abstain or engage, respectively, we allowed participants to explicitly randomise their choice among

1 Playing as many rounds as in our experiment renders risk neutrality a natural assumption when deriving a benchmark solution.

the two possible strategies (for earlier experiments offering explicitly mixed strategies, see, e.g., Ochs, 1995, or Anderhub et al., 2002).

In retrospective analysis, both situations where bidders should abstain from or engage in trade will render both choices as good and as bad with positive probability. If restricted to binary choices a bidder's (un)willingness to engage in trade could be detected by comparing behaviour in different rounds of the same game. This, however, would interfere with learning which we expect to be relevant. Allowing for mixing is a way to elicit (un)willingness to trade without confounding it with learning. Will loss aversion be equally strong for a bidder who has chosen probability 1 and one who came to the same choice by implementing proper mixing? If not, our experiment would serve as a particularly challenging setup for observing loss aversion, and, thus, it would even corroborate any confirmation of loss aversion effects. Another justification for the use of explicit mixed strategies is that in a binary choice protocol we could never be sure whether and, if so, participants randomly determine their choice behaviour. Offering explicit random devices therefore implies a better control of individual decision making.

More specifically, each participant played 250 rounds each of both the same high- and same low-quota game where a low quota suggests an efficient benchmark solution and a high quota an inefficient one. As mentioned earlier, we consider such extensive experience necessary since learning in stochastic settings is typically slow. We also wanted sufficiently many loss and gain experiences for each participant when trying to explore the path dependence of bidding behaviour and how adaptation is influenced by loss aversion.

As our analysis of the data shows, learning in the experiment is indeed slow; half of the subjects seem not to learn using the payoff maximising strategies within the given number of rounds. According to the EWA learning model, observed learning is characterized by a long memory, participants do not discount past experiences. This is reasonable given that even though the bidding task is highly stochastic, it remains stable throughout the experiment. Retrospective gains and losses have a slightly more heterogeneous impact on the learning dynamics. Even so, the majority of participants attach considerable importance to retrospective gains and losses. With respect to loss aversion we observe a substantial degree of heterogeneity in our data. Although the majority of participants puts a higher weight on losses than on gains, there are a few participants who discount losses. Nonetheless, allowing for loss aversion significantly improves the goodness-of-fit of the EWA learning model.

The paper proceeds as follows: The next section introduces the basic model in its continuous and its binary forms. In Section II we discuss the details of the experimental design and the laboratory protocol. Section III presents the experimental data. In Section IV we discuss experience-weighted attraction learning and loss aversion; we estimate several learning functions for our experiment participants and assess the importance of losses in their learning behaviour. Finally, Section V concludes.

1. The model

Before introducing the binary bidding task, let us briefly discuss its continuous version. Let v denote the value of an asset to be sold by a seller. This value is randomly drawn from a uniform distribution defined on $[0, 100]$ and is unknown to the bidder; the latter only knows how it has been generated. The seller knows v but has an evaluation that is equal to qv, with $q \in (0, 1]$ and $q \neq 0.5$. She agrees to sell the asset only if the bidder's bid satisfies $b > qv$. The value of q is known to both the seller and the bidder.

Since a risk neutral buyer earns $v - b$ when $b > qv$ (or $v < b/q$) and 0 otherwise, her expected payoff takes the following expression:

$$E(v-b) = \begin{cases} \int_0^{b/q} (v-b)\,\frac{1}{100}\mathrm{d}v = \frac{b^2}{100q}\left(\frac{1}{2q} - 1\right) & \text{for } 0 \le b \le 100q \\ \int_0^{100} (v-b)\,\frac{1}{100}\mathrm{d}v = 50 - b & \text{for } 100q < b < 100. \end{cases} \tag{1}$$

The first equation stands for the bidder's expected payoff when submitting a bid smaller than $100q$, the seller's highest possible evaluation of the asset. The second equation stands for the expected payoff when submitting a bid greater than this maximum possible evaluation.[2]

It thus follows from equation (1) that any positive bid $b \le 100q$ implies a positive expected profit if $q < 0.5$. Thus, the optimal bid b^* depends on q via

$$b^* = b^*(q) = \begin{cases} 100q & \text{for } 0 < q < 0.5 \\ 0 & \text{for } 0.5 < q < 1 \end{cases} \tag{2}$$

so that $E(v - b^*) = 50 - 100q$ for $0 < q < 0.5$ and 0 for $0.5 < q < 1$. In case of $q > 0.5$, the positive expected surplus of

$$\int_0^{100} (1-q)\frac{1}{100} v\,\mathrm{d}v = 50(1-q) \tag{3}$$

is lost, whereas in case of $q < 0.5$ it is fully exploited.

The binary version of this game assumes only two possible bids, 0 and B, instead of all bids in $[0, 100]$. In our experiments, all participants will repeatedly encounter two bidding environments: one with $0 < q < 0.5$ for which it is optimal to bid B, and one with $0.5 < q < 1$ for which it is optimal to bid 0. In both environments, participants will have to choose between bidding 0 and bidding B (by assigning a weighing probability, cf. next section). We further assume that $0 < B < 100$ so that bidding B does not warrant trade and thus efficiency.

Let us finally investigate the stochastic payoff effects for bidding 0 or B. By bidding $b = 0$, the bidder never buys. However, as the asset's value v is revealed

2 Notice that the assumption of constant relative risk averse (CRRA) preferences in this common value context would lead to an intractable model as the first order condition for expected utility maximisation would lead to a term $(-b)^a$ whenever the CRRA parameter a is smaller than 1.

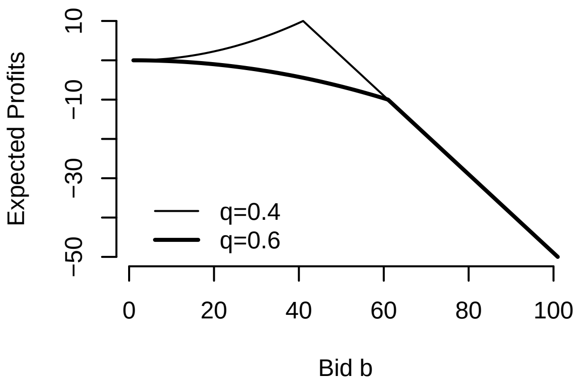

Figure 1: Expected profits for each bid b under the experimental conditions

after the sale, she should understand that compared to $b = B$, her previous bid $b = 0$ is retrospectively justified if $B > v$, which happens with probability $P_0^+ = B/100$. On the other hand, if $v > B > qv$, which happens with probability $P_0^- = B/(100q) - B/100$, having bid $b = 0$ counts as a loss of $v - B$ in retrospect.

For $b = B$, if no trade takes place (because $qv > B$, occurring with probability $P_0^n = 1 - (B/(100q))$), then there is no retrospective feeling of a loss or a gain. Similarly, $b = B$ implies a gain of $v - B$ if $v > B > qv$, which happens with probability $P_B^+ = B/(100q) - B/100$, and a loss of $v - B$ if $B > v$, which occurs with probability $P_B^- = B/100$. There is no trade whenever $qv > B$, which happens with probability $P_B^n = 1 - (B/(100q))$. The assumption $0 < B < \underline{q}(1 + \underline{q})^{-1}$ guarantees that all probabilities $P_0^+ = P_B^-$, $P_0^- = P_B^+$ and $P_0^n = P_B^n = P^n$ are positive. To summarise, regardless whether $q = \underline{q}$ or $q = \overline{q}$ applies, the gains and losses associated to bidding $b = B$ are actual, whereas those associated to bidding $b = 0$ are retrospective.

2. Laboratory protocol

Participants were instructed that there are two roles and they would act as buyers with the computer acting as the seller by following a fixed robot strategy. The instructions informed participants about parameter values, the distribution of the values, and the rule followed by the robot seller.[3]

We chose $B = 25$, $\underline{q} = 0.4$ and $\overline{q} = 0.6$. For such parameters and assuming risk neutrality, it is optimal to bid $b = B$ when $q = 0.4$ and to bid $b = 0$ when $q = 0.6$. We report in Figure 1 the bidder's expected profits for each value of q.

To observe the bidders' propensities P_t for bidding B, we asked them to assign a probability of submitting a bid B on an 11-point scale ranging from 0 to 100 percent with a step size of 10 percentage points. Therefore, the pure strategies of bidding 0 or B are equivalent to a choice of 0 or 100 percent on

3 A transcript of the German instructions can be found in the Appendix.

the 11-point probability scale; of course, we did not impose a default option for the mixed strategy choice in order not to influence subjects in any way.[4] One of the advantages of such a method is that it avoids the experimentally uncontrolled randomness of a binary choice between bidding $b_t = 0$ and $b_t = B$ when checking, for each individual participant, how well a postulated adaptive process captures her idiosyncratic learning behaviour.

While the feeling of loss aversion is probably less important when submitting bids by chance, we conjecture that the variety of situations encountered over a long sequence of play should ensure its salience in subjects' behaviour.[5] Note further that no subject in our experiment consistently determined his or her bid completely by chance and that pure strategy choices were rather frequent.

In spite of the 500 rounds of play, the whole experiment took less than two hours to complete, including the time needed to read the instructions. As already mentioned, we allowed for a sufficiently large number of repetitions to be able to test for learning effects since the assessment of loss aversion in a stochastic environment requires many repetitions. At the end of each round t, participants were informed about the value v_t, whether trade occurred or not, and about earnings for both possible choices, i.e. for $b_t = B$ and $b_t = 0$. The values $q = 0.4$ or $q = 0.6$ were displayed on each decision screen and alternated every ten rounds.

The conversion rate from experimental points to euros (1:200) was announced in the instructions. Participants received an initial endowment of 1500 experimental points (€ 7.50) to avoid bankruptcy during the experiment. As a matter of fact, such a case never occurred. Average earnings were € 10.56 (s.d.=1.3). Overall, 42 subjects participated in the two sessions that we conducted. The experiment was computerised using z-tree (Fischbacher, 2007) and was conducted with students of the University of Jena.

3. Descriptive results

We start by presenting some descriptive results on subjects' behaviour before assessing the extent of loss aversion within a learning framework.

4 The decision screen that we used in the experiment is displayed in the Appendix.

5 We are not aware of any study that specifically analyses this question. One referee pointed out that there could be a demand for mixing strategies when subjects are given the option to do so. This could indeed be the case in the first rounds of play as subjects may then find it difficult to figure out how to bid (i.e., mixing strategies can represent a way of spreading the risks associated to each pure strategy.) However, we expect such demand effect to vanish as the experiment proceeds since subjects should eventually figure out the sub-optimality of mixing strategies in either q-regime (as Observation 2 below suggests).

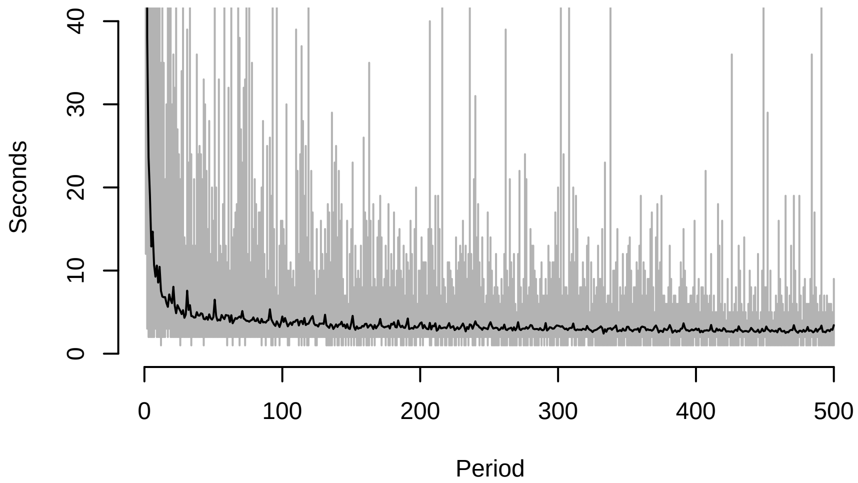

Note: The shaded area covers the actual decision time in each period for all participants.

Figure 2: Average decision time

3.1. Decision time

Since an experiment with 500 rounds of play is rather long, our first concern is naturally whether participants remained cognitively active during the whole experiment. In particular, boredom could induce subjects to consider no longer the available information and to adjust no longer their behaviour accordingly. Although we cannot exclude that some subjects got bored by the experiment, there are several indications that this was not at all predominant: Figure 2 shows the average decision time (in seconds) over the 500 rounds of play. The average decision time is the longest in the first period with about 490 seconds, and it declines as the experiment proceeds. By round 80 the average decision time falls below 5 seconds and it reaches 3.5 seconds in the last period.

We also observe periodic spikes in the average decision time until the very end. As most spikes occur at a q-regime switch, their presence suggest that participants considered the new information displayed on their screens. The average increase in decision time after a regime switch is about 30%, and the difference is highly significant ($p < 0.001$, according to a Wald-test using the results of a model that controls for repeated measurements). After the last regime switch (at round 490) the increase in decision time is about 67% (one-sided t-test, $p = 0.030$).

Eventually, we take the subjects' positive comments and reactions at informal end-of-session debriefings as evidence that they did not get bored by the experiment. Consequently, we are confident that participants took the experiment seriously and remained cognitively active till the very end.

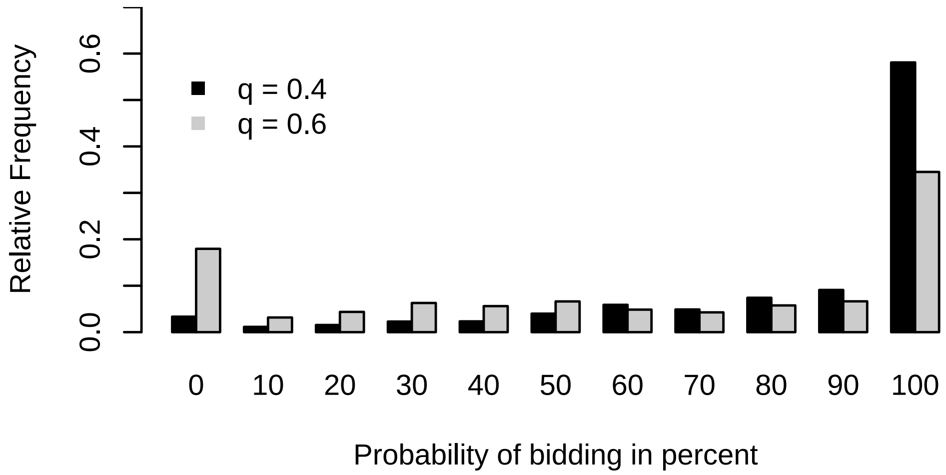

Figure 3: Relative frequency of probability choices

3.2. Strategy choices and learning

Observation 1 *Only 4 (out of 42) subjects always played the optimal pure strategy of bidding B when $q = 0.4$, and only 2 subjects always played the optimal pure strategy of bidding 0 when $q = 0.6$. Subjects performed much better when $q = 0.4$ than when $q = 0.6$.*

Table 1 displays the 500 weighing decisions of each participant along with the 25 individual average deviations from the optimal bids when $q = 0.4$ and when $q = 0.6$.[6] As deviations from the normative prediction are represented in grey, one can immediately see how subjects fared in comparison to the theoretical prediction and to what extent these deviations decline over time, to what extent learning occurs. Only a few participants play almost always optimally when $q = 0.6$ (i.e., subjects 2, 19, 26, and 39), and a few more do so when $q = 0.4$ (i.e., subjects 2, 9, 19, 26, 27, 28, and 39). For most subjects, the average individual deviations are significantly larger when $q = 0.6$ than when $q = 0.4$, which suggests a tendency to bid even if it is not optimal to do so. Such a bias towards bidding resembles the well-documented phenomenon of the winner's curse in standard first-price auctions. Another obvious aspect of the bidding data – that will have to be taken into account in the estimation of learning models – is a considerable heterogeneity across subjects.

Figure 3 shows for each value of q, the relative frequency of the eleven possible probability weights to be put on the event $b = B$. When $q = 0.4$, the optimal pure strategy of bidding B has been chosen 58% of times whereas the zero bid option has been chosen 3% of times. When $q = 0.6$, the optimal strategy of bidding 0 has been chosen 18% of times whereas bidding B for sure has been chosen 34% of times. The latter suggests that quite a few participants did not correctly take into account the expected revenue implications of such a choice

6 Subjects played 25 blocks of 10 periods of each q-regime.

Table 1: Choices and Deviations from the Normative Prediction

Subject	Choices	Mean deviation	
ID	1 Round 500	$\underline{q}$	$\overline{q}$
1			
2			
3			
4			
5			
6			
7			
8			
9			
10			
11			
12			
13			
14			
15			
16			
17			
18			
19			
20			
21			
22			
23			
24			
25			
26			
27			
28			
29			
30			

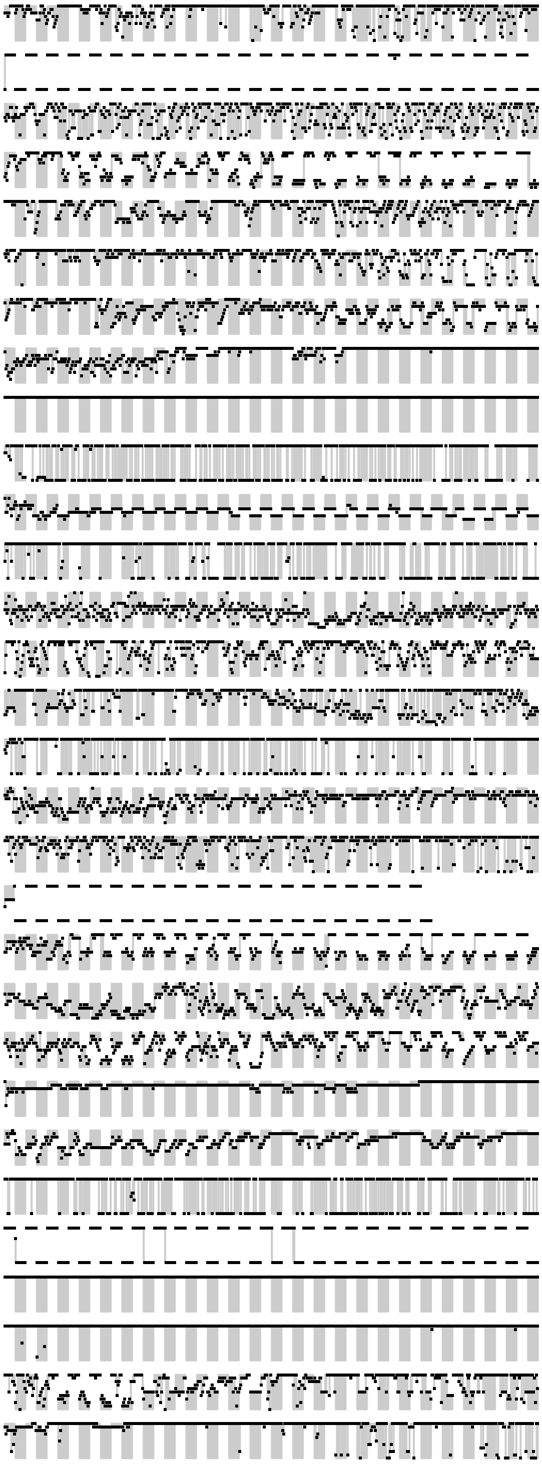

Table 1: Choices and Deviations from the Normative Prediction

Subject ID	Choices: Round 1 … 500	Mean deviation $\underline{q}$	Mean deviation $\overline{q}$
31			
32			
33			
34			
35			
36			
37			
38			
39			
40			
41			
42			

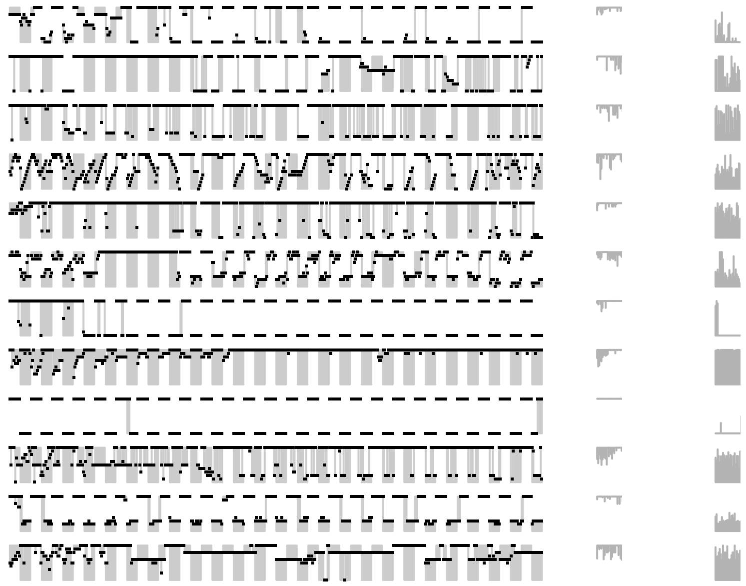

Note: Black dots represent actual choices during the 500 rounds, grey bars are deviations from the normative solution. Mean deviations from the normative solution in columns three and four are computed for each block of ten rounds under the respective q-regime. While a point at the bottom of a line indicates bidding with probability 0, a point at the top of a line indicates bidding with probability 1. The subject is closer to the normative solution the more white space the graph shows.

Subject 19 exceeded the time limit and had to stop after only 400 rounds.

when $q = 0.6$. In view of the many rounds played, arguing that subjects bid B for sure because they were extremely risk seeking, seems dubious. Note also that efficiency concerns cannot explain obsessive bidding since there has been no human seller.

To measure the extent of optimal play in individual behaviour, we look at the deviation of choices from the normative prediction, i.e. the relative size of the grey areas in Table 1. This measure of individual performance equals 0 if the participant always chose the expected payoff *maximising* pure strategy. It equals 1 if the participant always chose the expected payoff *minimising* pure strategy. Values between 0 and 1 result from mixing.

Figure 4 reports the distributions of deviations from the normative prediction for each q-regime and the individual difference in average deviations between treatments. The average deviation is equal to 0.207 when $q = 0.4$ and 0.593 when $q = 0.6$. The difference is significant at $p < 0.001$ according to a Wilcoxon rank sum test for paired data. It confirms that individual decisions are closer to the theoretical prediction when $q = 0.4$ than when $q = 0.6$.

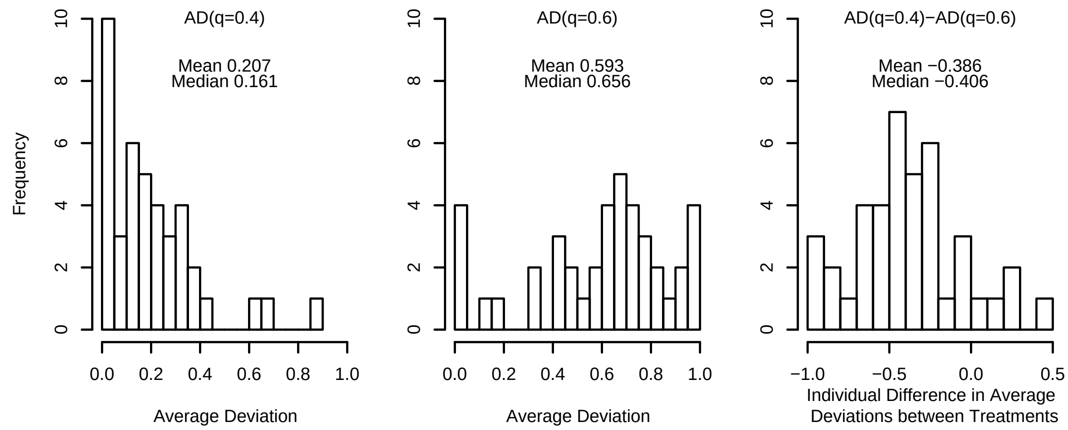

Figure 4: Distributions of deviations from payoff maximising bids.

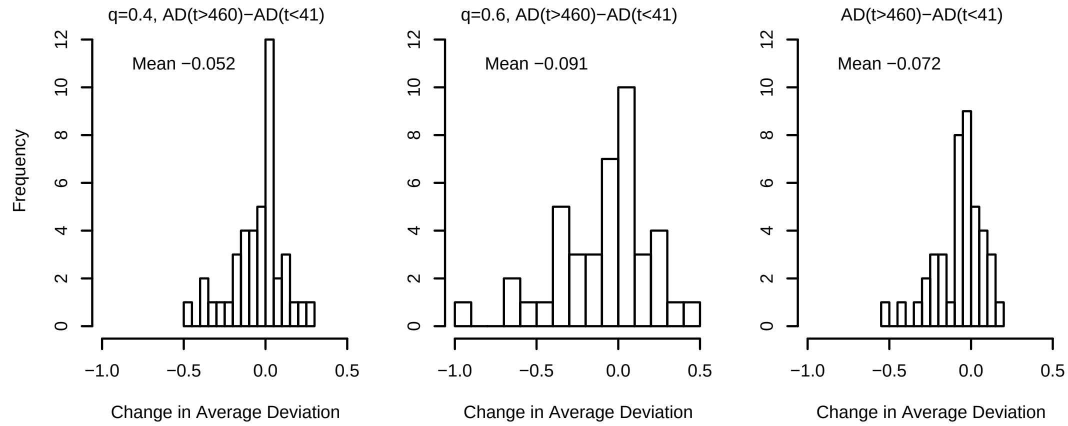

Figure 5: Difference in average deviations from the normative prediction between first and last 40 periods for both treatments

Observation 2 *On average, subjects performed significantly better in the last than in the first 40 rounds of the experiment. This holds for both q-regimes.*

Figure 5 plots the distributions of individual differences in the average deviations from the normative prediction between the first and the last 40 rounds of each q-regime. A negative difference indicates a reduction in deviations from the normative prediction over time and suggests that learning occurred over the course of the experiment. According to this measure, 22 participants moved *towards* the normative prediction when $q = 0.4$ and 23 when $q = 0.6$. Similarly, 14 participants moved *away* from the normative prediction when $q = 0.4$ and 15 when $q = 0.6$. On average, the participants' deviations decreased by 5 percentage points between the first and the last 40 rounds when $q = 0.4$ (t-test, $p = 0.035$) and by 9 percentage points when $q = 0.6$ ($p = 0.043$). Overall, 29 participants 'learnt' the normative solution and deviations decreased by 7 percentage points ($p = 0.003$).

In the last 100 rounds of play, 24% of participants (10 bidders) always played the optimal strategy when $q = 0.4$, and 10% of the participants (4 bidders) did so when $q = 0.6$. The latter outcome is in line with the 7% reported by Ball et al. (1991) for a game that corresponds to an environment where $q = 0.6$, which spanned over 20 rounds and for which no significant change in behaviour was found. Nevertheless, as only four bidders use the optimal strategy with $q = 0.4$ and $q = 0.6$, the well-documented pattern of the winner's curse is still very present. This result corroborates the findings of Selten et al. (2005) and Charness and Levin (2009), who also found this pattern to be persistent after 100 or 60 rounds of play, even in treatments designed to simplify subjects' bidding task.

Observation 3 *Revenues of bidders are higher under $q = 0.4$ than under $q = 0.6$. Many bidders fall prey to the winner's curse under $q = 0.6$.*

As expected, average earnings are also different across q-regimes: they are equal to 825.76 points when $q = 0.4$ (s.d.: 223.96; Min: 269; Max: 1274) and to -213.66 points when $q = 0.6$ (s.d.: 127.99; Min: -547; Max: 20).

4. EWA learning with(out) loss aversion

The situation our participants face in the experiment poses quite a challenge for learning theories. Obviously, whatever the bidder chooses, she can retrospectively experience both an encouragement and a regret with positive probability. As already mentioned, we are interested in two specific questions: (i) Will regret be more decisive than positive encouragement as suggested by (myopic) loss aversion (Benartzi and Thaler, 1995; Gneezy and Potters, 1997)? (ii) Will learning bring about an adjustment to bidding $b = B$ when $q = \underline{q}$ and $b = 0$ when $q = \overline{q}$?

In general, reinforcement learning relies on the cognitive assumption that what has been good in the past, will also be good in the future. In such a context, one does not have to be aware of the decision environment except for an implicit stationarity assumption justifying that earlier experiences are a reasonable indicator of future success. In our setup, actual as well as retrospective gains and losses can be derived unambiguously since the bidder is informed about the true value v_t after each round t, regardless whether the object has been bought or not.

Ho et al. (2007) proposed a one-parameter adaptive experience-weighted attraction (EWA) learning model which takes into account both past experience and expectations about the future. It nests reinforcement and belief learning and can predict the path dependence of individual behaviour in any normal-form game. Camerer (2003) shows that EWA has a better (predictive) fit than pure reinforcement and pure belief learning (and quantal response as a no-learning benchmark) in a number of games. He notes, however, that identification of EWA parameters requires a substantial number of periods. Ho et al. (2007) claim that the EWA learning model can easily be extended to games with incomplete information.

In the following, we describe how we incorporated loss aversion in the EWA learning model for our game of incomplete information. Denote bidder's jth choice by s^j, the round t choice by s_t, and the payoff from choice s^j by π^j. Let $I(\cdot,\cdot)$ be an indicator function that is equal to 1 if the arguments are equal and 0 otherwise. Using the notation of Ho et al. (2007), a bidder's attraction A_t^j for choice j at time t can be defined as

$$A_t^j = \frac{\phi N_{t-1} A_{t-1}^j + [\delta + (1-\delta) I(s^j, s_t)]\pi_t^j}{N_{t-1}\phi(1-\kappa_t) + 1} \tag{4}$$

where $N_t = N_{t-1}\phi(1-\kappa_t) + 1$ and $N_0 = 1$. Here, the parameter ϕ reflects profit discounting or, equivalently, the decay of previous attractions owing either to forgetting or deliberate discounting of past experience when the learning environment is changing. This parameter will be freely estimated, as in the original EWA model of Camerer and Ho (1999). The parameter δ stands for the weight placed on retrospective gains and losses (foregone payoffs). A parameter value of $\delta = 0$ reduces the EWA learning model to a pure reinforcement model. On the other hand, being fully responsive to foregone payoffs, i.e. $\delta = 1$, means that the decision maker's behaviour is not only reinforced but rather driven by a thorough cognitive retrospective analysis of behaviour.

Ho et al. (2007) proposed to tie the decay of attraction to the weight on foregone payoffs: $\phi = \delta$. They argue that a subject whose behaviour is driven by a thorough cognitive retrospective analysis is also more likely not to discount past experiences in a stable environment. We will estimate models where we place different restrictions on δ to analyse the implications of the assumptions of pure reinforcement and pure belief learning, as well as to test whether the heuristic simplification of Ho et al. (2007) is justified within our stochastic, yet stable bidding framework.

Finally, κ controls the growth rate of attraction. A value of $\kappa = 0$ implies that attractions will be averaged, a value of $\kappa = 1$ inplies that attraction will be cumulated.

Ho et al. (2007) argue that if in the past a subject often changed her strategy, she will be more likely to change her strategy now given the same feedback. This can be reflected in the model by letting κ be equal to a normalised Gini coefficient on choice frequencies. In our estimations we will focus on average attraction under pure reinforcement and pure belief learning, i.e. we restrict $\kappa = 0$. Furthermore, we will estimate the EWA model as proposed by Ho et al. (2007) and compute the normalised Gini coefficient on choice frequencies for each q-regime and period t separately in order to use this variable in place of κ_t.

In our analysis, the attraction A_t^j will only be calculated for the binary choice of bidding B or not bidding. We need to compute only A_t^B; the attraction of not bidding is always zero because the payoff for not bidding is always zero. The mapping of the attraction A_t^B to the mixed strategies via a probabilistic choice function is defined in the next section.

To augment the standard model by loss aversion, we introduce a loss weight l, and we define by ψ_t the actual losses and non-realised, retrospective gains in round t, depending on whether bidding B generated a loss or bidding 0 left some money on the table:

$$\psi_t = \rho_t^0(v_t - B) + \rho_t^B(B - v_t) \tag{5}$$

where $\rho_t^0 = 1$ if $v_t > B > qv_t$ and $b = 0$, and $\rho_t^0 = 0$ otherwise, and similarly, $\rho_t^B = 1$ if $B > v_t$ and $b = B$, and $\rho_t^B = 0$ otherwise.

The bidder's attraction A_t^B for the binary choice of bidding $b = B$ in round t then becomes

$$A_t^B = \frac{\phi N_{t-1} A_{t-1}^B + [\delta + (1-\delta) I(s^B, s_t)][\pi_t^B + l\psi_t]}{N_{t-1}\phi(1-\kappa_t) + 1}. \tag{6}$$

With such a specification, a loss weight of $l = 0$ indicates that losses influence the attraction of choosing $b = B$ to the same extent as gains. While a loss weight of $l = -1$ indicates that losses have no influence on the attraction, a loss weight of $l = 1$ means that losses influence the attraction of choosing $b = B$ by twice the extent of gains.

4.1. Estimation procedure

Learning models such as those considered here are models of individual behaviour. Yet, most econometric tests of these models report aggregate parameters that have been estimated from individual behaviour in interactive games. While such estimates often suggest a behaviour that is consistent with some kind of reinforcement learning, they typically do not illustrate the extent of heterogeneity in subjects' behaviour. For our analysis of loss aversion it is thus important to note that, e.g., Johnson et al. (2006) provide empirical evidence that individual loss aversion is heterogeneous. Additionally, using Monte Carlo simulations, Wilcox (2006) has shown that there exists a strong estimation bias if players are heterogeneous but a homogeneous representative agent model is estimated. We therefore estimate our models separately for each participant and discuss the distributions of the estimated parameters.[7] For the reasons of comparability, we shall also report the estimation outcomes of a model that assumes a representative agent and includes the data from all participants, comparable to the model used in Ho et al. (2008) based on centipede game data by Nagel and Tang (1998).

7 One could have estimated a random coefficients model to reduce the estimation bias but, as this would have required specific assumptions on the distribution of each random coefficient, we chose not to proceed this way.

In the tradition of reinforcement learning models, the attraction A_t^B for the binary decision to bid B in period t is assumed to determine the probabilistic choice behaviour (see Bush and Mosteller, 1955; Roth and Erev, 1998; Camerer and Ho, 1999). Since we allowed explicit mixing of the two pure strategies bidding $B = 25$ and not bidding we need to map the attraction A_t^B to the eleven available mixed strategies. The mixed strategies follow a natural order given by their respective probability of submitting the bid. Consequently, the mapping of the attraction A_t^B to the mixed strategies can be achieved by an ordinal multinomial choice model (see Agresti, 2010). If the natural order of the mixed strategies were not obvious to our participants an alternative modelling approach would include computing separate attractions for each available mixed strategy and mapping them to choices via an unordered multinomial choice model.[8]

We choose the ordered logit model for the mapping of attraction to choices: Let s^* be a single latent variable with

$$s_t^* = \beta_1 I(q_t, \underline{q})\underline{A}_t^B + \beta_2 I(q_t, \overline{q})\overline{A}_t^B + u_t \tag{7}$$

where $I(\cdot,\cdot)$ is equal to 1 when the arguments are equal and 0 otherwise. The attraction A_t^B is evaluated separately for each q-regime and is denoted $\underline{A}_t^B$ for $q = \underline{q}$ and $\overline{A}_t^B$ for $q = \overline{q}$. The initial attractions $\underline{A}_1^B$ and $\overline{A}_1^B$ are set equal to the expected payoff of the corresponding q-regime. The error u is logistically distributed with $F(z) = \exp(z)/(1 + \exp(z))$. The probability P of choosing the mixing strategy $s = j$ with $j \in \{0, 0.1, \ldots, 1\}$ is then

$$\begin{aligned} P_t(s_t = j) =& P(\alpha_{j-0.1} < s_t^* \leq \alpha_j) \\ =& F\left(\alpha_j - \beta_1 I(q_t, \underline{q})\underline{A}_t^B + \beta_2 I(q_t, \overline{q})\overline{A}_t^B\right) \\ &- F\left(\alpha_{j-0.1} - \beta_1 I(q_t, \underline{q})\underline{A}_t^B + \beta_2 I(q_t, \overline{q})\overline{A}_t^B\right) \end{aligned} \tag{8}$$

where the threshold parameters $\alpha_0 = -\infty$ and $\alpha_1 = \infty$. Thus, the parameters to be estimated are β_i, ϕ, δ, l and $\alpha_{0.1}, \ldots, \alpha_{0.9}$ by maximising the log-likelihood[9]

$$\mathcal{L} = \sum_t \sum_j I(s_t, j) \ln P_t(s_t = j). \tag{9}$$

The decay parameter ϕ, the weight on foregone payoffs δ and the loss weight l are only identified if the sensitivities to the EWA learning model β_1 and β_2 are (significantly) positive and if the decay parameter ϕ lies in $[0, 1]$. In our estimation, we therefore impose $\beta_i = \exp(\tilde{\beta}_i)$ and enforce the restriction on ϕ by applying a logistic transformation $\phi = 1/(1 + \exp(\tilde{\phi}))$.

8 We thank the editor, Frank Cowell, for pointing this out.

9 The R procedures (see R Development Core Team, 2010) that were used for the estimation can be obtained from the authors upon request.

4.2. Estimation results

To analyse loss aversion within an experience-weighted attraction framework we estimate eight models with various restrictions on the parameters to assess their separate impact on learning.

Our null-model includes only one treatment dummy so that it assumes a constant bidding behaviour within each treatment, with a possibly different constant behaviour in each treatment (as indicated by the normative solution). A learning model must provide a better fit than the null-model before we consider capturing any learning dynamics.

In a first learning model, we estimate equation (8) with the following constraints: $l = 0$, $\delta = 0$ and $\kappa = 0$. Such restrictions reduce the EWA model to a reinforcement learning model with profit discounting and no weight on foregone payoffs. We will refer to this model as Model 1. In a second model (Model 2), we relax the restriction on the loss aversion parameter and estimate it freely.

In Model 3 and 4 we restrict the weight on foregone payoffs to $\delta = 1$ and keep $\kappa = 0$. Such restrictions reduce the EWA model to a reinforcement learning model with profit discounting and equal weights of actual and foregone payoffs. The loss aversion parameter l is restricted to $l = 0$ in Model 3 and is estimated freely in Model 4.

The fifth model has been proposed by Ho et al. (2007) and assumes the exploitation parameter κ to be a Gini coefficient on choice frequencies and, in addition, that $\delta = \phi$, which is reasonable given the stable strategic environment of our individual decision-making experiment. The loss aversion parameter l is restricted to $l = 0$ in Model 5 and is estimated freely in Model 6.

In Models 7 and 8, we assess whether the assumption $\delta = \phi$ in Models 5 and 6 is valid for our data by estimating both parameters. This also allows us to assess the importance of retrospective gains and losses for learning. The loss aversion parameter l is restricted to $l = 0$ in Model 7 and freely estimated in Model 8.

The estimation results for each of the 42 subjects involved and the representative agent are reported in Table 2 of the Appendix and lead to the following observation:

Observation 4 *A subject is said to learn according to some EWA model if this model outperforms the null-model of constant behaviour (at the 1% significant level). Models 5 and 6 (inspired from Ho et al. 2007) identify 21 subjects (out of 42) as 'learners' whereas Models 1 to 4 identify only 11 'learners'.*

Subjects for which the EWA models do not indicate a significant learning pattern correspond to those identified earlier with non-decreasing deviations from the normative solution over time (cf. Observation 2 and Figure 5). For these subjects the joint likelihood ratio tests do not unambiguously support that β_1 and β_2 of equation (7) are significantly different from 0. The EWA model parameters are therefore not identified and the corresponding subjects need to

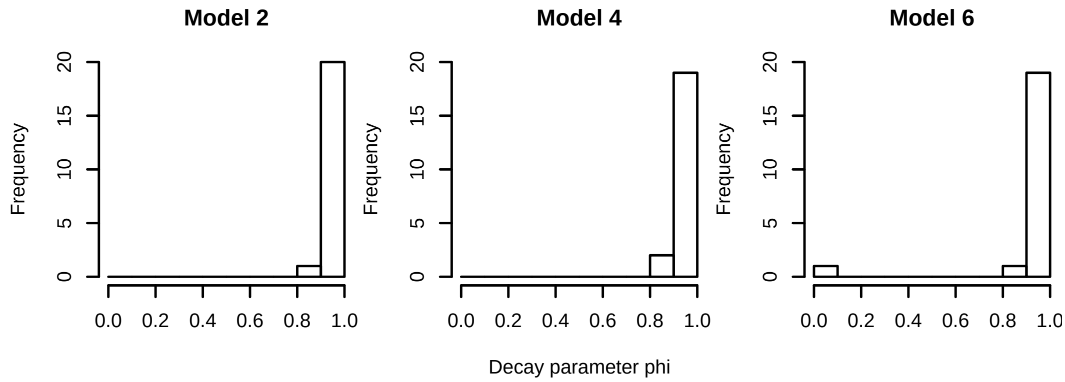

Figure 6: Distributions of estimated decay parameters ϕ in Models 2, 4, and 6.

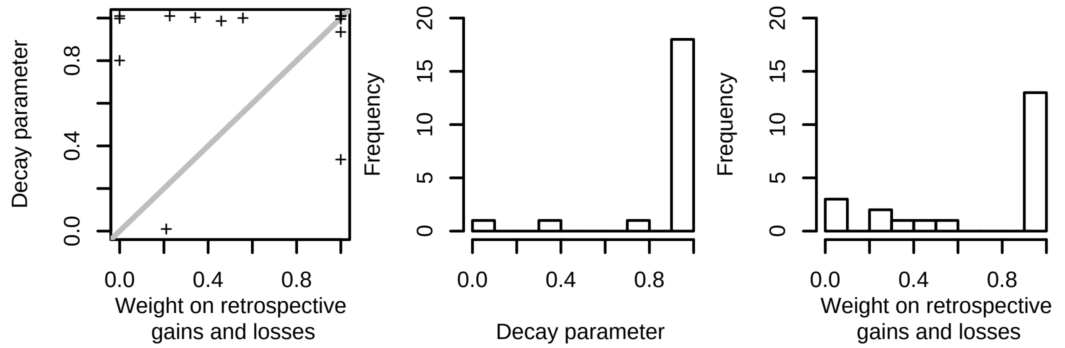

Figure 7: Distributions of estimated ϕ (decay) and δ (retrospective gains and losses) in Model 7.

be discarded from further analysis.[10]

Let us first discuss the parameter ϕ that reflects the decay of previous attractions. Notice that even though the bidding framework of our experiment is highly stochastic, it remains stable throughout the experiment. We therefore do not expect participants to discount past experiences, i.e. we expect ϕ to be close to 1. This is supported by the data as the mean values of ϕ are 0.988 for Model 2 and 0.977 for Model 4.

Including retrospective gains and losses slightly decreases the decay parameter. Figure 6 further reveals that the distributions are all unimodal. As expected, most participants who displayed a significant learning pattern did not discount past experience. This is also reflected in the estimates for the representative agent model that shows a ϕ estimate close to 0.99. The histograms for Models 6 and 7 (cf. Figures 6 and 7) reveal that tying the weight of retrospective gains and losses δ to the decay parameter ϕ (Model 6) has a small positive effect on the latter as the mean values of ϕ are 0.943 in Models 6 and 8 and 0.904 in Model 7.

10 These are the subjects with IDs 2, 3, 5, 7, 9, 10, 13, 15, 16, 18, 25–29, 33, 36, 38, 39, 41, 42 (cf. Table 1).

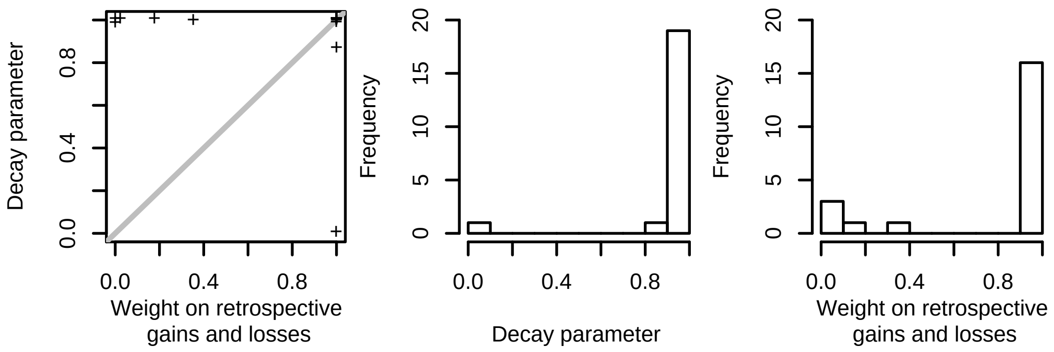

Figure 8: Distributions of estimated ϕ (decay) and δ (retrospective gains and losses) in Model 8 that includes a loss weight.

This leads us to the question whether tying δ to ϕ is a legitmate simplification for our data. The left panel of Figure 7 shows the (joint) distribution of the estimates for δ and ϕ in Model 7. Most observations lie in the upper right corner of the joint distribution graph, confirming the intuition that both parameters should be linked closely. There are, however, some exceptions. Appropriate likelihood ratio tests show that $\delta = \phi$ has to be rejected for only 6 (5) of the 21 individuals at the 5% (1%) level. The largest difference for which $\delta = \phi$ cannot be rejected is 0.672 whereas the smallest difference that is rejected is 0.012. The latter is certainly not economically relevant as two estimates that differ by such a small amount are, for all matters of interpretation, equivalent. The clustering of ϕ and δ at 1 is even more pronounced in Model 8 that includes the loss weight l (see Figure 8). For reasons of parsimony we consider the simplification $\delta = \phi$ for our experimental data as legitimate. The estimations for the representative agent, however, strongly contradicts this simplification because the estimated coefficients are then $\delta = 0$ and $\phi = 0.99$.

Notice that although Model 1 (Model 2) has the same number of freely estimated parameters as Model 3 (Model 4), it uses less information as it does not take account of the bidder's retrospective gains and losses ($\delta = 0$). Surprisingly, however, a pairwise comparison of the models' log-likelihoods reveals that the fit of Model 1 (Model 2) is not substantially worse than the one of Model 3 (Model 4) ($p > 0.148$, according to one-sided t-tests). The EWA models (Models 5–8) introduced by Ho et al. (2007), however, perform considerably better than all other versions of EWA learning in our experiment ($p < 0.001$).

Observation 5 *A standard reinforcement-learning model like Model 1 (or Model 2) explains the observed behaviour as effectively as the more sophisticated EWA models such as Model 3 (or Model 4) that account for foregone payoffs. Still, controlling for the growth rate of attractions further improves the fit of the model.*

Finally, we assess the impact of loss aversion on learning in our experiment. Figure 9 shows the distributions of the estimated loss weights in Models 2, 4, 6 and 8 (excluding for Model 2 one and for Models 6 and 8 two extreme values).

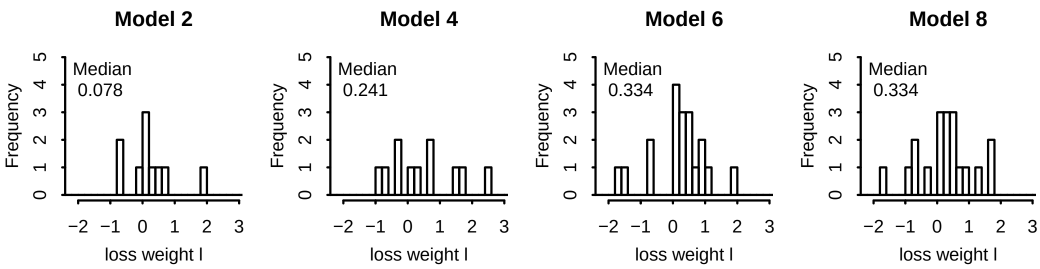

Figure 9: Distributions of estimated loss weight l in Models 2, 4, 6, and 8.

In Model 2, which neglects retrospective gains and losses, the median value of the loss aversion parameter l is 0.078; in Model 4, which weighs retrospective gains and losses as much as actual ones, the median value of l is 0.241. In both models the estimate is positive for 7 out of the 11 subjects for whom the model is a significant improvement over the null-model. The loss aversion parameter for the representative agent is -6.5 in Model 2 and 0.746 in Model 4. In Model 6, which links the weight of retrospective gains and losses δ to the decay parameter ϕ, and in Model 8, which allows both parameters to vary freely, the median loss aversion parameter is 0.334 and the estimate is positive for 16 (in Model 6) and 15 (in Model 8) of the 21 subjects for whom the model is a significant improvement over the null-model. Allowing for loss aversion improves the model fit for 14 out of 21 individuals in both models. The loss aversion parameter for the representative agent is -1.336 in Model 6 and 0.125 in Model 8. As reported in previous research (e.g., Wilcox, 2006, and Hichri and Kirman, 2007), the estimates for the representative agent in our experiments do not capture population heterogeneity and are thus seriously biased.

At last, the conduct of joint likelihood ratio tests indicates that including a loss weight significantly improves the explanatory power of the simple reinforcement model as well as of the more sophisticated EWA models (at the 5% level).

Observation 6 *Accounting for loss aversion in EWA learning models significantly improves their goodness-of-fit.*

On average, we observe a tendency for losses to have a higher impact on learning than gains. Yet, for some participants the opposite is true; they discount losses. This may be explained by a house-money effect. Thaler and Johnson (1990) find that under some circumstances a prior gain can increase subjects' willingness to accept gambles. This finding is labelled the house money effect, because gamblers often use the phrase "playing with the house money" to express the feeling of gambling while ahead. The essence of the idea is that until winnings are completely depleted, losses are coded as reductions of a gain, as if losing some of "their money" does not hurt as much as losing one's own cash. In our experiment subjects had a generous endowment that fully covered potential losses and may have induced such a feeling. Further experimental investigations would be needed to check explicitly how this endowment affects subjects′ learning dynamics.

4.3. Estimation bias and power analysis

As our analysis and conclusions rely on the outcomes of likelihood ratio tests for nested models, we proceed with assessing the power of this test regarding the loss aversion parameter l. Like Salmon (2001), we are also particularly interested in evaluating a potential estimation bias with a series of Monte-Carlo simulations. This is important as Cabrales and Garcia-Fontes (2000) observed that some EWA parameters are likely to be downward biased and often inaccurate if the number of periods is too small.

For the Monte-Carlo simulations of our Models 5 and 6, we independently draw the true parameter values of the sensitivities β_1 and β_2 and the decay parameter ϕ from uniform distributions that cover the range of the observed estimated parameters, i.e. $\beta \in (0, 0.1]$, $\phi \in [0.8, 1.0]$, and the weight on retrospective gains and losses $\delta = \phi$. The loss aversion parameter is set to $l \in \{-1.4, -1.0, -0.6, -0.2, 0.2, 0.6, 1.0, 1.4\}$. For each value of l and different lengths of play, i.e., series from 100 rounds up to 1000 rounds in steps of 100 rounds, we estimated our models for 200 independently simulated subjects. This allows us to determine the power of the likelihood ratio test for the loss parameter l and a possible estimation bias for different lengths of the experiment. Such an exercise is important given our belief that the number of repetitions in earlier studies has been too low to fully account for possible learning dynamics or to estimate the impact of loss aversion in such a stochastic environment where learning is typically very slow.

Figure 10 shows for each value of the loss parameter l, the mean, the median, the 25% quantile, and the 75% quantile of the estimated loss parameter and the power of the likelihood ratio test with $H0 : l = 0$ at the 1% level.[11]

The test's power increases with the number of periods and the increase is most rapid over the first 300 rounds of play. Interestingly, the power shows more extreme values the smaller the magnitude of the loss parameter. For $l \in \{-0.2, 0.2\}$ the power is lower at small numbers of rounds and higher at larger numbers of rounds than when the loss parameter is larger in absolute terms. Additionally, the precision of the estimates depends inversely on the true value of the loss parameter l. A larger value of the true loss parameter leads to a higher variance in its estimate as indicated by the interquartile range. The interquartile range decreases with the number of rounds, indicating that longer periods of play and therefore more data points lead to more precise estimates. However, after about 500 periods the interquartile range seems to be rather

11 A significant likelihood ratio test for one simulation adds only to the power if the sign of the estimated loss parameter is correct too.

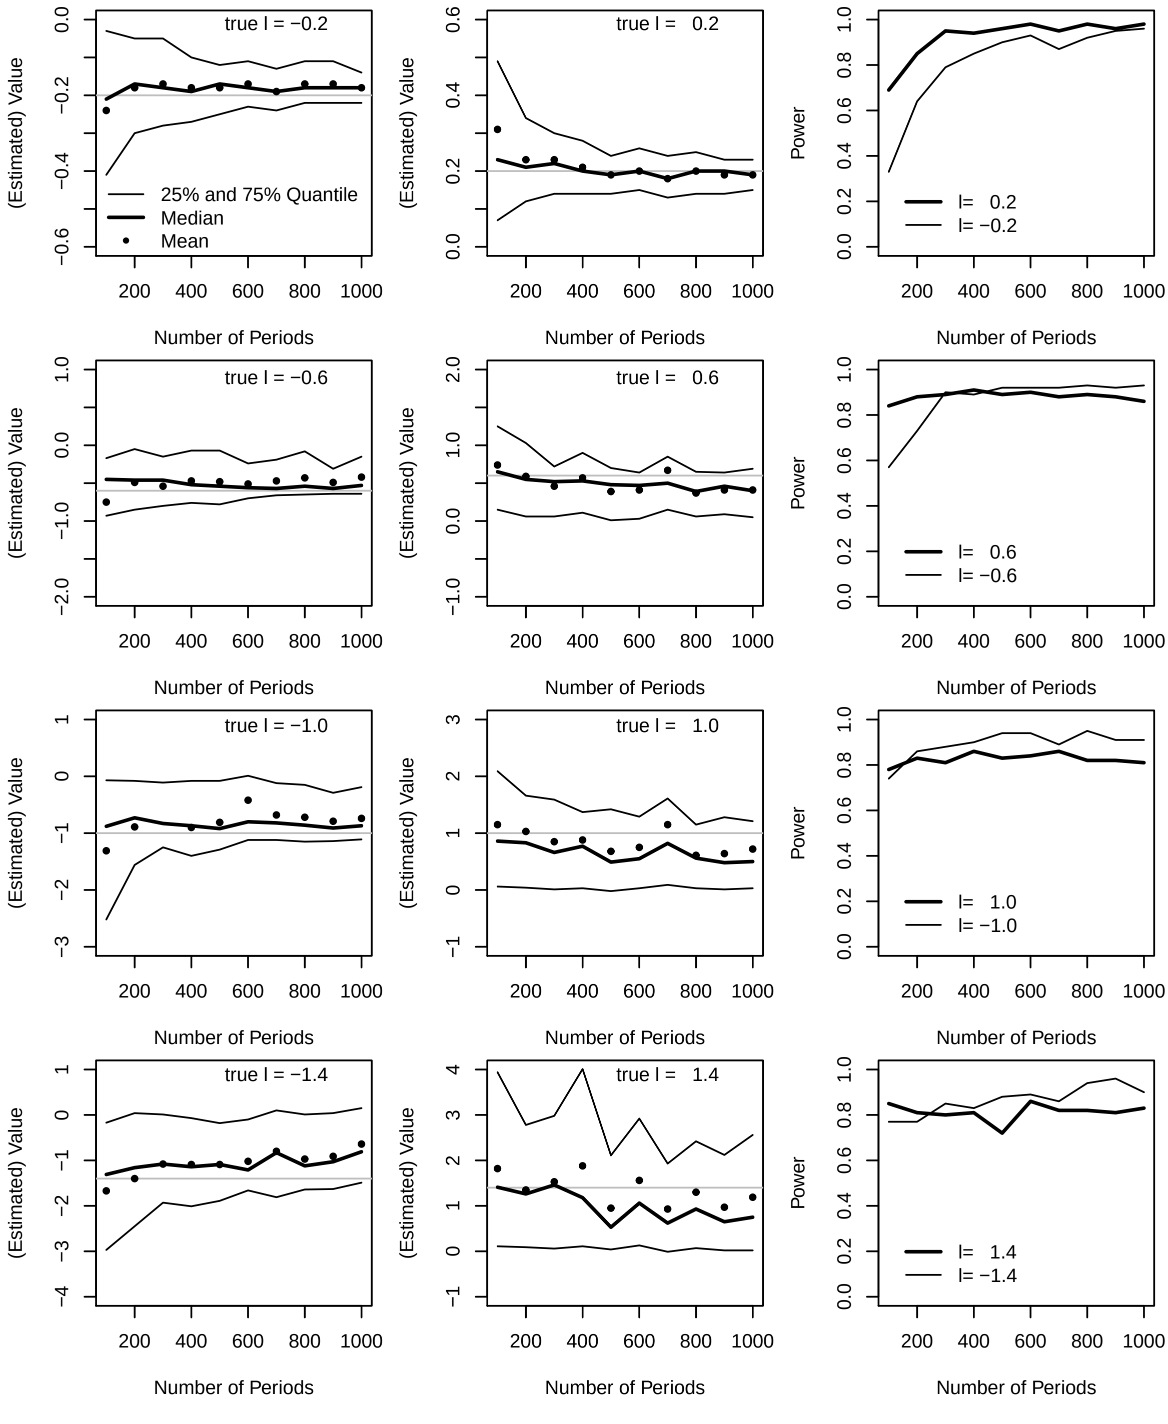

Figure 10: Estimation bias for loss parameter l and power analysis of test for loss aversion H0: $l = 0$ vs. H1: $l \neq 0$

stable. Therefore, our choice of conducting the experiment with 500 rounds is *ex post* justified by the results of the simulations. While shorter experiments would have yielded a much lower power and might have led to erroneous conclusions, i.e. not detecting the significance of loss aversion, we would not have gained much power or estimation precision with longer experiments.

The point estimates show a further noteworthy pattern. If the number of periods is small, absolute median and mean point estimates of the loss aversion parameter are slightly larger in magnitude than their true values. For longer periods of play we observe, however, a regression to zero, i.e. the true absolute effect size is under-estimated for most lengths of play.

In conclusion, even though our significance tests are quite powerful, point estimates of the loss aversion parameter are slightly downward biased in magnitude. This leads to a rather conservative assessment of the impact of loss aversion on learning dynamics in our paper.

5. Conclusion

One of the most robust findings of auction experiments is the winner's curse in common value auctions. However, all reasonable adaptation or learning dynamics suggest that 'winners' should learn to update their profit expectations more properly. In the bilateral trade situation at hand, the winner's curse means that the buyer neglects that the price of B will be accepted only if it exceeds the seller's evaluation of the asset (qv), i.e. only if the value v is smaller than B/q. Our experiment is designed in a way rendering it optimal for the buyer, in terms of expected revenues, to abstain from bidding B when the seller's quota (q) is high and to bid B with probability one when it is low.

Although we have provided ample opportunity for learning (250 rounds for each q-regime), bidding B with a positive probability when the seller's quota is high remains a strong pattern throughout the experiment and weakens only slightly over time. In our view, this ex post justifies the relatively large number of rounds of play. Our expectation that learning in stochastic environments can be rather slow is confirmed at least for the case when one should abstain from bidding. In contrast, learning to bid B with probability one when the seller's quota is low was much faster. This is also in line with the size of the financial incentives. The expected gain for bidding B when the seller's quota q is low is larger than the absolute expected loss for bidding B when the seller's quota is high. Bidding B is more profitable for the lower quota than abstaining is for the high one, and is therefore reinforced more strongly. The finding for low and high quotas together seem to show that whether learning is fast or slow does not depend on the general environment but whether one should learn omitting (to bid) or committing (to bid) which, in turn, depends here finally on a numerical parameter, namely the quota of lower evaluation by the seller.

In general, although we have reduced the bidding task to its simplest expression by removing confounding effects such as bidders' strategic interaction, it still seems to be quite a challenge for individuals to learn bidding optimally. We observe considerable individual heterogeneity of learning dynamics, and there is only a small fraction of people who almost instantaneously converge to optimal bidding. Moreover, learning to avoid the winner's curse seems to be rather weak. Either participants understand sooner or later that bidding is dangerous when q is high, or never recognise it clearly and are only slightly discouraged by (the more frequent) loss experiences.

Important conclusions concerning learning are that (1) loss aversion influences the individual learning dynamics but is highly heterogeneous, and (2) even though retrospective gains and losses have some influence on learning to bid their impact seems rather minor. Finally, we have shown that (3) loss aversion does not seem to be specific to rare decision making.

Acknowledgements

We are grateful for the assistance of Ben Greiner who helped in programming the experiment.

References

Agresti, A., 2010, Analysis of ordinal categorical data, Hoboken: John Wiley & Sons.

Anderhub V., D. Engelmann and W. Güth , 2002, An experimental study of the repeated trust game with incomplete information, Journal of Economic Behavior and Organization, 48, 197-216.

Armantier O., 2004, Does observation influence learning? Games and Economic Behavior, 46, 221-239.

Avrahami J., W. Güth and Y. Kareev, 2001, Predating predators: An experimental study, Homo Oeconomicus, 18, 337-352.

Becker G.M., M.H. de Groot and J. Marshak, 1964, Measuring utility by a single-response sequential method, Behavioral Science, 9, 226-232.

Ball S.B., M.H. Bazerman and J.S. Carroll , 1991, An evaluation of learning in the bilateral winner's curse, Organizational Behavior and Human Decision Processes, 48, 1-22.

Bazerman M.H. and W.F. Samuelson, 1983, I won the auction but don't want the prize, Journal of Conflict Resolution, 27, 618-634.

Benartzi S. and R.H. Thaler, 1995, Myopic loss aversion and the equity premium puzzle, Quarterly Journal of Economics, 110, 73-92.

Bush R. and F. Mosteller, 1955, Stochastic models for learning, New York: Wiley.

Cabrales, A. and W. Garcia-Fontes, 2000, Estimating learning models from experimental data, Working Paper, Universitat Pompeu Fabra 501.

Camerer, C., 2003, Behavioral Game Theory, Princeton: Princeton University Press.

Camerer, C. and T.-K. Ho, 1999, Experience-weighted attraction (EWA) learning in normal-form games, Econometrica, 67, 827-873.

Charness, G. and D. Levin, 2009, The origin of the winner's curse: A laboratory study, American Economic Journal: Microeconomics, 1: 207-36.

Fischbacher U., 2007, z-Tree - Zurich toolbox for readymade economic experiments - experimenter's manual. Experimental Economics, 10, 171-178.

Gandenberger O., 1961, Die Ausschreibung, Heidelberg: Quelle und Meyer.

Garvin S. and J. Kagel, 1994, Learning in common value auctions: Some initial observations, Journal of Economic Organization and Behavior, 25, 351-372.

Gneezy U. and J. Potters, 1997, An experiment on risk taking and evaluation periods, Quarterly Journal of Economics, 112, 631-645.

Güth W., R. Ivanova-Stenzel, M. Königstein and M. Strobel, 2003, Learning to bid – An experimental study of bid function adjustments in auctions and fair division games, Economic Journal, 103, 477-494.

Hichri W. and A. Kirman, 2007, The emergence of coordination in public good games, European Physical Journal B, 55, 2, 149-159.

Ho T.-K., C. Camerer and J.-K. Chong, 2007, Self-tuning experience weighted attraction learning in games, Journal of Economic Theory, 133, 177-198.

Ho T.-K., X. Wang and C. Camerer, 2008, Individual differences in EWA learning with partial payoff information, Economic Journal, 118, 37-59.

E. Johnson, S. Gächter and B. Herrmann, 2006, Exploring the nature of loss aversion, CeDEx Discussion Paper No. 2006-02, University of Nottingham.

Kagel J., 1995, Auctions: A survey of experimental research; in J. Kagel and A. Roth (eds.), The Handbook of Experimental Economics, Cambridge, MA: MIT Press.

Kahneman D. and A. Tversky, 1979, Prospect theory: An analysis of decision under risk, Econometrica, 47, 263-91.

Langer T. and M. Weber, 2005, Myopic prospect theory vs. myopic loss aversion: how general is the phenomenon?, Journal of Economic Behavior and Organization, 56, 25-38.

Nagel R. and F. Tang, 1998, Experimental results on the centipede game in normal form: an investigation on learning, Journal of Mathematical Psychology, 42, 356-84.

Ochs J., 1995, Games with unique, mixed strategy equilibria: An experimental study, Games and Economic Behavior, 10, 202-217.

Ockenfels, A. and R. Selten, 2005, Impulse Balance Equilibrium and Feedback in First Price Auctions, Games and Economic Behavior, 51, 155-170.

R Development Core Team, 2010, R: A language and environment for statistical computing, Vienna, Austria: R Foundation for Statistical Computing.

Roth A.E. and I. Erev, 1998, Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria, American Economic Review, 88, 848-880.

Salmon, T., 2001, An evaluation of econometric models of adaptive learning, Econometrica, 69, 1597-628.

Samuelson W.F. and M.H. Bazerman, 1985, The winner's curse in bilateral negotiations; in V.L. Smith (ed.): Research in Experimental Economics, Greenwich: JAI Press, 3, 105-137.

Starmer C., 2000, Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk, Journal of Economic Literature, 38, 332-382.

Selten R., K. Abbink and R. Cox, 2005, Learning direction theory and the winner's curse, Experimental Economics, 8(1), 5-20.

Selten R. and J. Buchta, 1998, Experimental sealed bid first price auctions with directly observed bid functions; in O.D. Budescu, I. Erev and R. Zwick (eds.), Games and Human Behavior – Essays in Honor of Amon Rapoport, Mahwah (N.J.): Lawrence Erlbaum Ass.

Thaler, R. H., and E. J. Johnson, 1990, Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, Management Science, 36, 643-660.

Tversky A. and D. Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, Journal of Risk and Uncertainty, 5, 297-323.

Vickrey W., 1961, Counterspeculation, auctions, and competitive sealed tenders, Journal of Finance, 16, 8-37.

Wilcox, N., 2006, Theories of learning in games and heterogeneity bias, Econometrica, 74, 1271-1292.

A. Software Screen

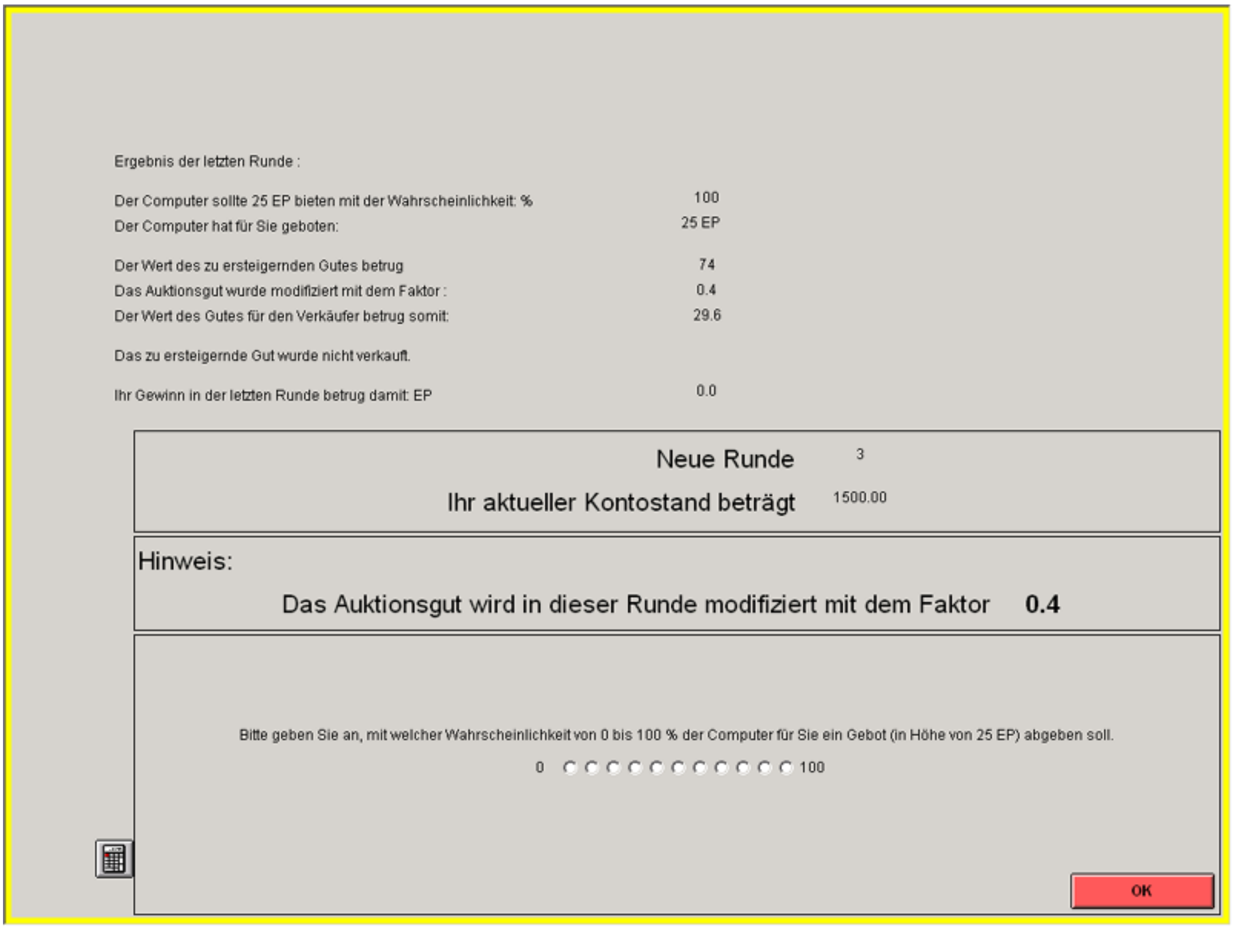

Figure 11: Computer screen

Figure 11 displays the decision screen. It features feedback information from the previous round, the round number, earnings, the prevalent q, the mixed strategy choice and a calculator symbol that opens a calculator window.

B. Instructions (originally in German)

Welcome to the experiment!

Please do not speak with other participants from now on.

Instructions

This experiment is designed to study decision making. You can earn 'real' money, which will be paid to you in cash at the end of the experiment. The 'experimental currency unit' will be 'experimental points (EP)'. EP will be converted to Euro according to the exchange rate in the instructions. During the experiment you and other participants will be asked to take decisions. Your decisions will determine your earnings according to the instructions. All participants receive identical instructions. If you have any questions after reading the instructions, raise your hands. One of the experimenters will come to you and answer them privately. The experiment will last for 2 hours maximum.

Roles

You are in the role of the buyer; the role of the seller is captured by a robot strategy of the computer. The computer follows an easy strategy that will be fully described below.

Stages

At the beginning of each round you can decide whether you want to 'bid' for a good or abstain. For this you can apply a special procedure that we will introduce in a moment. Your bid B is always 25 EP. The value of the good v will be determined independently and by chance move in every round. v is always between (and including) 0 EP and 100 EP. The chosen values are equally distributed, which means that every integer number between (and including) 0 EP and 100 EP is the value of the good with equal probability. The computer as the seller only sells the good when the bid is higher than the value multiplied by a factor q.

Mathematically, the computer only sells if $B > qv$.

The factor q takes on the values 0.4 or 0.6. At the beginning of each round you will be informed about the value of q in that round. Every 10 rounds q changes. Keep in mind that you are the only potential buyer of the good. Therefore, this is a very special case of an auction.

Special procedure to choose to bid or to abstain

The choice between bidding or abstaining is made by assigning balls to the two possibilities. You have 10 balls that have to be assigned to the two possibilities. The number of assigned balls corresponds with the probability that a possibility is chosen. Thus, each ball corresponds with a probability of 10 percent. All possible distributions of the 10 balls to the two possibilities are allowed.

Information at the end of each round

At the end of each round you receive information on the value v, whether the good was sold by the computer (or whether the good would have been sold in case you had bid). Additionally, you see your round profit.

Profit

- If you did not bid, your round profit is 0.
- If the bid was smaller than the value of the good multiplied by the factor ($B < qv$), the computer does not sell. No transaction takes place and your round profit is 0.
- If you bid and $B > qv$, then your round profit depends on the value of the good. In case of a value of the good that is higher than the bid $v > B$, you make a profit. In case of a value of the good that is lower than the bid, you make a loss in that round. Losses in single rounds can, of course, be balanced by profits from other rounds.

At the beginning of the experiment you receive an endowment of 1500 EP. At the end of the experiment, experimental points will be exchanged at a rate of 1:200. That means 200 EP = 1 Euro.

Rounds

There are 500 rounds in the experiment. In each round you have to take the same decision (assign the balls to the two possibilities of bidding or abstaining) as explained above. Only factor q will be changed every 10 rounds.

Means of help

You find a calculator on your screen. You are, of course, allowed to write notes on the instruction sheets.

Thank you for participating!

C. Estimation Results

In Table 2 each block of nine lines shows the estimation results for one subject. Line 1 is the Null-Model. In lines $2-9$ follow the different EWA models in the order discussed in the paper: In lines 2 and 3 $\delta = 0$, in lines 4 and 5 $\delta = 1$, in lines 6 and 7 $\delta = \phi$, and in line 9 all parameters are freely estimated. In the last column of lines 3, 5, 7, and 9 the p-value of the likelihood ratio test with Ho: $l = 0$ is reported. In Line 8 $l = 0$, the p-value of the likelihood ratio test with Ho: $\delta = \phi$ is reported in the last column.

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model $P(>\chi^2)$	Test for l or δ $P(>\chi^2)$
1	7.275	-0.012				-778.36		
	<0.001	<0.001	1.000			-778.37	1.000	
	0.002	0.001	1.000	-46.817		-777.96	0.667	0.367
	<0.001	<0.001	1.000			-778.37	1.000	
	<0.001	0.092	0.841	-0.857		-777.48	0.414	0.183
	0.023	0.049	0.001			-772.02	<0.001	
	0.023	0.045	<0.001	0.101		-771.91	0.002	0.630
	0.023	0.048	<0.001		0.211	-771.87	0.002	0.574
	0.022	0.033	<0.001	0.913	0.998	-771.25	0.003	0.253
2	-17.733	53.092				-13.04		
	0.009	356.008	1.000			-59.86	1.000	
	0.745	353.519	1.000	-4.036		-21.25	1.000	<0.001
	<0.001	168.643	1.000			-13.03	0.985	
	0.025	179.09	1.000	0.459		-14.92	1.000	1.000
	0.008	2.233	1.000			-36.65	1.000	
	0.517	<0.001	0.804	-1.740		-21.26	1.000	<0.001
	0.007	3.791	0.993		1.000	-14.09	1.000	<0.001
	0.617	<0.001	0.830	-1.549	0.100	-8.46	0.027	<0.001
3	7.882	0.069				-1118.74		
	<0.001	0.104	1.000			-1118.34	0.372	
	0.006	0.068	0.252	0.651		-1107.95	<0.001	<0.001
	0.014	0.021	1.000			-1118.53	0.523	
	0.002	<0.001	1.000	-12.478		-1118.37	0.695	0.572
	<0.001	<0.001	1.000			-1118.83	1.000	
	<0.001	<0.001	1.000	13.157		-1118.83	1.000	0.959
	<0.001	<0.001	1.000		1.000	-1118.83	1.000	0.992
	<0.001	<0.001	1.000	32.403	1.000	-1118.50	0.922	0.411
4	3.236	4.096				-855.30		
	1.502	<0.001	1.000			-857.21	1.000	
	1.063	<0.001	1.000	-0.645		-825.29	<0.001	<0.001
	1.249	<0.001	0.998			-831.27	<0.001	
	0.925	<0.001	0.992	-0.305		-824.02	<0.001	<0.001
	0.007	0.032	0.998			-756.61	<0.001	
	0.011	0.029	0.996	0.601		-744.25	<0.001	<0.001
	0.006	0.058	1.000		0.227	-746.98	<0.001	<0.001
	0.010	0.037	1.000	0.652	0.177	-738.23	<0.001	0.001
5	6.846	-0.130				-884.02		
	<0.001	<0.001	1.000			-884.35	1.000	
	<0.001	<0.001	1.000	45.511		-884.34	1.000	0.888
	<0.001	<0.001	1.000			-884.35	1.000	
	0.004	0.007	1.000	25.135		-883.52	0.609	0.200
	<0.001	0.001	1.000			-884.27	1.000	
	<0.001	<0.001	1.000	56.855		-883.88	0.869	0.376
	<0.001	0.001	1.000		1.000	-884.27	1.000	1.000
	<0.001	<0.001	1.000	65.241	1.000	-883.75	0.910	0.609

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model P(> χ^2)	Test for l or δ P(> χ^2)
6	3.179	1.234				-832.96		
	0.161	2.038	1.000			-795.41	<0.001	
	<0.001	1.404	1.000	0.511		-790.70	<0.001	0.002
	0.202	1.395	1.000			-798.69	<0.001	
	<0.001	1.111	1.000	0.632		-789.53	<0.001	<0.001
	0.001	0.030	1.000			-781.78	<0.001	
	<0.001	0.019	1.000	0.255		-777.58	<0.001	0.004
	0.001	0.030	1.000		1.000	-781.78	<0.001	0.991
	<0.001	0.019	1.000	0.255	1.000	-777.58	<0.001	0.996
7	5.600	0.064				-1056.00		
	0.049	<0.001	0.908			-1054.33	0.068	
	0.175	<0.001	1.000	2.401		-1051.22	0.008	0.013
	0.028	<0.001	0.718			-1054.23	0.060	
	0.152	<0.001	1.000	1.539		-1052.82	0.042	0.093
	0.012	<0.001	0.839			-1053.62	0.029	
	0.010	<0.001	0.900	0.914		-1053.39	0.074	0.498
	0.012	<0.001	0.824		0.999	-1053.61	0.092	0.884
	0.011	<0.001	0.898	0.973	1.000	-1053.27	0.141	0.615
8	8.628	0.337				-759.83		
	1.029	<0.001	0.992			-679.08	<0.001	
	1.035	<0.001	0.992	0.122		-676.48	<0.001	0.023
	0.271	<0.001	0.983			-743.59	<0.001	
	0.679	0.129	0.993	0.702		-702.09	<0.001	<0.001
	0.009	<0.001	1.000			-693.44	<0.001	
	0.044	<0.001	1.000	0.464		-672.25	<0.001	<0.001
	0.009	<0.001	1.000		1.000	-693.44	<0.001	0.963
	0.044	<0.001	1.000	0.464	1.000	-672.25	<0.001	0.990
9	63.475	7.556				<0.001		
	4.5e+27	0.102	1.000			<0.001	1.000	
	4.3e+37	0.892	1.000	-3.220		<0.001	1.000	1.000
	4.5e+27	0.102	1.000			<0.001	1.000	
	4.3e+37	0.892	1.000	-3.220		<0.001	1.000	1.000
	8.8e+28	0.008	1.000			<0.001	1.000	
	4.0e+54	0.001	0.996	-4.339		<0.001	1.000	1.000
	8.8e+28	0.008	1.000		1.000	<0.001	1.000	1.000
	4.0e+54	0.001	0.996	-4.339	0.996	<0.001	1.000	1.000
10	5.895	0.127				-413.33		
	0.030	<0.001	1.000			-413.39	1.000	
	<0.001	<0.001	1.000	129.953		-413.55	1.000	1.000
	0.015	<0.001	1.000			-413.49	1.000	
	<0.001	0.002	0.992	20.205		-413.41	1.000	0.687
	0.001	<0.001	1.000			-410.78	0.024	
	<0.001	<0.001	1.000	-35.927		-408.33	0.007	0.027
	0.001	<0.001	1.000		1.000	-410.78	0.078	1.000
	<0.001	<0.001	1.000	-48.684	1.000	-407.37	0.008	0.166

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model $P(>\chi^2)$	Test for l or δ $P(>\chi^2)$
11	6.490	3.944				-491.96		
	0.813	1.719	0.986			-498.83	1.000	
	0.763	<0.001	0.997	-1.491		-497.10	1.000	0.063
	0.609	0.847	0.999			-474.69	<0.001	
	0.377	1.236	0.991	0.796		-460.62	<0.001	<0.001
	0.005	0.019	0.994			-464.03	<0.001	
	0.009	0.015	0.993	0.457		-457.71	<0.001	<0.001
	0.007	0.027	0.992		0.558	-463.17	<0.001	0.188
	0.009	0.015	0.993	0.458	1.000	-457.68	<0.001	0.813
12	3.811	0.318				-404.82		
	0.019	0.073	1.000			-405.67	1.000	
	<0.001	0.018	1.000	3.843		-405.64	1.000	0.782
	0.020	0.120	1.000			-404.87	1.000	
	<0.001	<0.001	1.000	-429.747		-404.32	0.606	0.293
	<0.001	0.003	1.000			-398.40	<0.001	
	<0.001	<0.001	1.000	34.570		-393.05	<0.001	0.001
	<0.001	0.003	1.000		1.000	-398.40	0.002	0.998
	<0.001	<0.001	1.000	67.624	1.000	-388.74	<0.001	0.003
13	2.648	0.170				-1014.13		
	<0.001	0.035	<0.001			-1012.94	0.122	
	<0.001	0.001	<0.001	51.086		-1010.89	0.039	0.043
	0.017	0.015	0.606			-1013.70	0.354	
	<0.001	0.001	<0.001	38.704		-1011.79	0.096	0.050
	<0.001	0.035	<0.001			-1012.94	0.122	
	0.001	0.002	<0.001	26.636		-1010.94	0.041	0.046
	<0.001	0.035	<0.001		<0.001	-1012.94	0.303	0.995
	0.001	0.001	<0.001	35.254	<0.001	-1010.91	0.092	0.811
14	5.171	1.485				-999.07		
	0.227	0.349	1.000			-1002.52	1.000	
	0.223	0.046	1.000	-1.286		-1000.23	1.000	0.033
	0.180	0.114	1.000			-1004.83	1.000	
	0.106	<0.001	1.000	-2.033		-1003.39	1.000	0.090
	0.017	<0.001	0.975			-990.17	<0.001	
	0.016	<0.001	0.985	0.350		-988.65	<0.001	0.082
	0.018	<0.001	0.976		0.460	-989.70	<0.001	0.331
	0.016	<0.001	0.985	0.351	1.000	-988.65	<0.001	0.910
15	6.434	-0.233				-907.49		
	<0.001	0.719	1.000			-903.32	0.004	
	<0.001	1.373	1.000	-0.253		-900.55	0.001	0.019
	<0.001	0.213	1.000			-907.36	0.614	
	0.040	0.067	0.992	20.047		-882.76	<0.001	<0.001
	<0.001	0.011	<0.001			-908.28	1.000	
	<0.001	<0.001	<0.001	82.059		-907.67	1.000	0.268
	<0.001	0.012	<0.001		1.000	-908.18	1.000	0.645
	<0.001	<0.001	<0.001	103.425	1.000	-907.26	0.926	0.364

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model P(> χ^2)	Test for l or δ P(> χ^2)
16	5.349	0.295				-385.11		
	<0.001	0.194	1.000			-385.38	1.000	
	<0.001	0.006	1.000	28.159		-384.74	0.688	0.257
	<0.001	0.127	1.000			-385.82	1.000	
	<0.001	<0.001	1.000	-234.438		-383.87	0.287	0.048
	0.001	<0.001	1.000			-382.66	0.027	
	<0.001	<0.001	1.000	24.200		-386.06	1.000	1.000
	0.001	<0.001	1.000		1.000	-381.87	0.039	0.210
	<0.001	<0.001	1.000	24.200	1.000	-386.06	1.000	1.000
17	4.516	-0.130				-970.68		
	<0.001	1.026	0.977			-939.37	<0.001	
	0.148	3.025	0.992	-0.741		-902.03	<0.001	<0.001
	<0.001	0.580	0.977			-947.24	<0.001	
	0.055	1.507	0.993	-0.835		-919.52	<0.001	<0.001
	0.003	<0.001	1.000			-956.84	<0.001	
	0.003	0.018	1.000	-1.664		-837.14	<0.001	<0.001
	0.003	<0.001	1.000		1.000	-956.84	<0.001	0.999
	0.003	0.018	1.000	-1.663	1.000	-837.14	<0.001	0.996
18	5.404	-0.767				-897.27		
	<0.001	<0.001	1.000			-908.26	1.000	
	<0.001	0.927	1.000	-0.707		-899.39	1.000	<0.001
	<0.001	<0.001	1.000			-908.26	1.000	
	0.001	<0.001	1.000	173.948		-901.18	1.000	<0.001
	<0.001	<0.001	1.000			-908.26	1.000	
	<0.001	<0.001	1.000	-10.334		-908.26	1.000	0.989
	<0.001	<0.001	1.000		1.000	-908.26	1.000	0.992
	<0.001	<0.001	0.999	-18.105	1.000	-887.96	<0.001	<0.001
19	0.624	27.315				-39.84		
	0.066	356.131	1.000			-66.88	1.000	
	0.553	354.179	0.999	-6.767		-65.09	1.000	0.058
	0.659	2.870	1.000			-97.55	1.000	
	0.374	<0.001	1.000	-21.775		-48.06	1.000	<0.001
	0.654	0.082	1.000			-7.17	<0.001	
	0.776	1.408	1.000	0.812		-7.31	<0.001	1.000
	0.654	0.082	1.000		1.000	-7.17	<0.001	0.999
	0.686	0.122	1.000	0.489	1.000	-7.09	<0.001	0.503
20	7.131	3.960				-712.69		
	0.529	1.024	1.000			-758.12	1.000	
	0.141	<0.001	1.000	-8.595		-721.49	1.000	<0.001
	0.468	0.905	0.992			-739.90	1.000	
	0.175	<0.001	1.000	-4.170		-735.15	1.000	0.002
	0.020	0.019	0.991			-628.46	<0.001	
	0.022	0.017	0.991	0.175		-627.80	<0.001	0.251
	0.018	0.029	0.993		0.342	-625.83	<0.001	0.022
	0.019	0.027	0.993	0.047	0.352	-625.80	<0.001	0.046

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model P(> χ^2)	Test for l or δ P(> χ^2)
21	2.638	-0.494				-1153.48		
	0.171	<0.001	0.803			-1135.41	<0.001	
	0.172	<0.001	0.815	0.278		-1135.17	<0.001	0.491
	0.040	0.023	0.749			-1155.39	1.000	
	0.164	0.036	0.834	1.789		-1107.15	<0.001	<0.001
	0.037	0.006	0.116			-1152.99	0.320	
	0.049	0.016	0.880	1.859		-1104.76	<0.001	<0.001
	0.064	<0.001	0.793		<0.001	-1139.93	<0.001	<0.001
	0.051	0.017	0.865	1.793	1.000	-1102.44	<0.001	0.031
22	2.748	2.915				-976.05		
	<0.001	4.274	1.000			-954.46	<0.001	
	<0.001	4.137	1.000	0.026		-954.40	<0.001	0.729
	<0.001	1.689	1.000			-955.43	<0.001	
	<0.001	1.468	1.000	0.185		-954.24	<0.001	0.124
	0.012	<0.001	1.000			-940.23	<0.001	
	0.007	0.054	1.000	-0.750		-901.02	<0.001	<0.001
	0.012	<0.001	1.000		1.000	-940.23	<0.001	0.990
	0.007	0.054	1.000	-0.750	1.000	-901.02	<0.001	1.000
23	9.262	-0.202				-553.50		
	0.066	<0.001	0.838			-549.80	0.007	
	0.392	0.018	1.000	1.926		-535.51	<0.001	<0.001
	0.078	<0.001	0.901			-549.45	0.004	
	0.406	0.018	1.000	1.582		-536.52	<0.001	<0.001
	0.007	<0.001	0.925			-549.61	0.005	
	<0.001	0.001	1.000	-22.401		-380.07	<0.001	<0.001
	0.007	<0.001	0.927		1.000	-549.58	0.020	0.821
	<0.001	<0.001	1.000	-26.867	1.000	-380.03	<0.001	0.764
24	6.481	0.609				-908.79		
	<0.001	1.355	0.975			-879.72	<0.001	
	<0.001	1.207	0.973	0.078		-879.53	<0.001	0.533
	<0.001	0.942	0.965			-836.43	<0.001	
	<0.001	1.176	0.975	-0.257		-833.72	<0.001	0.020
	0.003	<0.001	1.000			-894.50	<0.001	
	0.003	<0.001	1.000	0.147		-894.50	<0.001	0.896
	0.003	<0.001	1.000		1.000	-894.50	<0.001	0.970
	0.003	<0.001	1.000	0.146	1.000	-894.49	<0.001	0.967
25	4.023	0.115				-327.86		
	<0.001	0.397	1.000			-325.26	0.022	
	<0.001	<0.001	1.000	1224.621		-326.83	0.358	1.000
	<0.001	0.273	1.000			-325.25	0.022	
	0.050	0.073	1.000	7.213		-321.89	0.003	0.010
	0.006	0.061	0.376			-321.64	<0.001	
	<0.001	<0.001	<0.001	124.900		-325.35	0.081	1.000
	0.022	0.086	0.017		1.000	-311.53	<0.001	<0.001
	<0.001	0.001	<0.001	48.927	<0.001	-325.03	0.130	0.424

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model P(> χ^2)	Test for l or δ P(> χ^2)
26	9.806	128.212				-31.01		
	4.455	<0.001	1.000			-75.88	1.000	
	5.4e+29	<0.001	1.000	-0.632		-31.01	1.000	<0.001
	2.468	4.357	1.000			-60.27	1.000	
	3.5e+32	<0.001	1.000	-0.536		-31.01	1.000	<0.001
	0.144	0.056	1.000			-43.08	1.000	
	1.0e+28	<0.001	0.954	-5.261		-31.01	1.000	<0.001
	0.144	0.056	1.000		1.000	-43.08	1.000	0.999
	1.0e+28	<0.001	0.954	-5.261	0.954	-31.01	1.000	1.000
27	63.475	7.556				<0.001		
	8.9e+36	0.612	1.000			<0.001	1.000	
	4.5e+64	1.285	1.000	0.462		<0.001	1.000	1.000
	8.9e+36	0.612	1.000			<0.001	1.000	
	4.5e+64	1.285	1.000	0.462		<0.001	1.000	1.000
	5.5e+106	<0.001	1.000			<0.001	1.000	
	2.9e+109	<0.001	1.000	-1.066		<0.001	1.000	1.000
	5.5e+106	<0.001	1.000		1.000	<0.001	1.000	1.000
	2.9e+109	<0.001	1.000	-1.066	1.000	<0.001	1.000	1.000
28	15.562	1.402				-33.68		
	<0.001	1.532	1.000			-26.98	<0.001	
	<0.001	1.521	1.000	0.009		-26.98	0.001	0.984
	<0.001	1.532	1.000			-26.98	<0.001	
	<0.001	1.522	1.000	0.008		-26.98	0.001	0.989
	0.022	<0.001	<0.001			-34.54	1.000	
	1206373.459	<0.001	<0.001	-18.946		-32.78	0.406	0.060
	0.022	<0.001	<0.001		<0.001	-34.54	1.000	0.998
	3777953.293	<0.001	<0.001	-17.804	<0.001	-32.78	0.614	1.000
29	5.949	1.389				-921.43		
	0.425	<0.001	1.000			-933.07	1.000	
	0.287	<0.001	1.000	-0.718		-920.41	0.363	<0.001
	0.434	<0.001	1.000			-930.64	1.000	
	0.242	<0.001	0.938	-0.858		-911.78	<0.001	<0.001
	0.006	<0.001	0.999			-932.58	1.000	
	<0.001	<0.001	0.998	-22.059		-927.71	1.000	0.002
	0.008	<0.001	0.996		0.459	-932.01	1.000	0.285
	0.028	0.030	0.912	-1.311	0.284	-915.62	0.009	<0.001
30	5.081	-0.071				-490.81		
	0.026	0.063	0.281			-486.52	0.003	
	0.143	0.256	1.000	-4.992		-475.68	<0.001	<0.001
	0.028	0.072	0.284			-484.66	<0.001	
	0.065	<0.001	1.000	2.512		-490.15	0.516	1.000
	0.022	0.048	0.395			-485.51	0.001	
	0.005	0.018	0.977	-1.593		-475.83	<0.001	<0.001
	0.023	0.053	0.327		1.000	-484.29	0.001	0.119
	0.008	0.064	0.981	-0.872	<0.001	-464.14	<0.001	<0.001

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model P(> χ^2)	Test for l or δ P(> χ^2)
31	5.949	4.335				-494.49		
	0.662	<0.001	1.000			-596.37	1.000	
	0.163	0.042	1.000	-7.671		-501.93	1.000	<0.001
	0.433	0.303	1.000			-588.14	1.000	
	0.173	0.050	1.000	-6.328		-511.32	1.000	<0.001
	0.008	0.015	1.000			-439.01	<0.001	
	0.009	0.014	1.000	0.146		-438.59	<0.001	0.359
	0.008	0.015	1.000		1.000	-439.01	<0.001	1.000
	0.009	0.014	1.000	0.146	1.000	-438.59	<0.001	1.000
32	7.061	1.762				-466.98		
	0.415	<0.001	0.99			-455.57	<0.001	
	0.531	<0.001	1.000	0.769		-455.43	<0.001	0.589
	0.241	0.681	0.990			-451.20	<0.001	
	0.298	1.114	0.991	-0.647		-447.71	<0.001	0.008
	0.010	0.017	0.975			-457.97	<0.001	
	0.010	0.005	0.990	0.920		-454.36	<0.001	0.007
	0.002	0.020	0.988		1.000	-454.21	<0.001	0.006
	0.004	0.005	0.995	1.758	1.000	-444.44	<0.001	<0.001
33	11.093	1.191				-384.41		
	0.379	0.210	1.000			-390.16	1.000	
	0.033	<0.001	1.000	-12.326		-384.69	1.000	0.001
	0.275	0.179	1.000			-387.94	1.000	
	0.019	<0.001	1.000	-15.445		-384.97	1.000	0.015
	<0.001	<0.001	<0.001			-399.38	1.000	
	<0.001	0.061	<0.001	-1.763		-397.50	1.000	0.052
	<0.001	<0.001	0.001		0.001	-399.38	1.000	0.999
	<0.001	0.074	<0.001	-1.493	0.574	-396.47	1.000	0.151
34	3.073	1.920				-905.29		
	0.169	0.348	1.000			-917.22	1.000	
	0.088	<0.001	1.000	-12.164		-898.70	0.001	<0.001
	0.116	0.246	1.000			-917.36	1.000	
	0.007	<0.001	1.000	-74.548		-909.70	1.000	<0.001
	0.009	<0.001	0.991			-878.00	<0.001	
	0.014	<0.001	1.000	1.150		-871.33	<0.001	<0.001
	0.013	0.004	0.988		<0.001	-874.16	<0.001	0.006
	0.012	<0.001	1.000	1.368	0.022	-870.16	<0.001	0.126
35	5.663	1.935				-441.40		
	<0.001	1.201	0.985			-434.55	<0.001	
	<0.001	0.917	0.987	0.288		-434.09	0.001	0.336
	0.068	0.623	0.984			-441.59	1.000	
	0.045	<0.001	0.912	-13.658		-445.55	1.000	1.000
	0.006	0.003	0.992			-431.96	<0.001	
	0.003	0.005	0.998	0.584		-432.43	<0.001	1.000
	0.003	0.006	1.000		<0.001	-425.30	<0.001	<0.001
	0.003	0.013	1.000	-0.250	<0.001	-425.16	<0.001	<0.001

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model P(> χ^2)	Test for l or δ P(> χ^2)
36	3.653	2.468				-887.44		
	0.039	1.601	1.000			-886.44	0.158	
	0.050	4.613	1.000	-0.486		-879.49	<0.001	<0.001
	0.184	0.692	1.000			-908.45	1.000	
	0.134	<0.001	0.978	-4.269		-874.15	<0.001	<0.001
	0.011	<0.001	0.986			-903.45	1.000	
	0.007	<0.001	0.956	-4.881		-881.38	0.002	<0.001
	0.011	<0.001	0.986		1.000	-903.41	1.000	0.765
	0.011	<0.001	0.949	-3.921	0.284	-879.13	0.001	0.034
37	4.912	5.346				-160.06		
	1.195	<0.001	1.000			-217.91	1.000	
	<0.001	<0.001	1.000	-4822.192		-382.81	1.000	1.000
	0.595	0.898	1.000			-195.69	1.000	
	0.850	1.683	1.000	-1.563		-134.06	<0.001	<0.001
	0.002	0.038	1.000			-85.42	<0.001	
	0.002	0.143	1.000	-0.715		-82.87	<0.001	0.024
	0.003	0.037	1.000		1.000	-85.00	<0.001	0.359
	0.002	0.143	1.000	-0.715	1.000	-82.87	<0.001	1.000
38	10.228	-2.658				-420.13		
	<0.001	<0.001	0.296			-478.73	1.000	
	<0.001	<0.001	<0.001	567.432		-478.72	1.000	0.971
	<0.001	<0.001	0.029			-478.73	1.000	
	<0.001	<0.001	1.000	28.945		-478.73	1.000	1.000
	<0.001	<0.001	0.619			-478.73	1.000	
	<0.001	<0.001	<0.001	44.321		-478.72	1.000	0.957
	<0.001	<0.001	<0.001		1.000	-478.72	1.000	0.945
	<0.001	<0.001	0.034	41.516	<0.001	-478.72	1.000	0.973
39	14.644	41.497				-35.41		
	1.6e+30	<0.001	1.000			-35.41	1.000	
	8.3e+36	<0.001	1.000	0.397		-35.41	1.000	0.994
	3.4e+23	<0.001	1.000			-35.41	1.000	
	9.0e+29	<0.001	1.000	0.368		-35.41	1.000	0.993
	3.5e+37	<0.001	1.000			-35.41	1.000	
	2.9e+53	<0.001	1.000	-0.955		-35.41	1.000	0.998
	3.5e+37	<0.001	1.000		1.000	-35.41	1.000	1.000
	3.9e+58	<0.001	1.000	0.397	1.000	-35.41	1.000	0.995
40	5.645	0.567				-659.27		
	0.453	0.077	0.992			-632.88	<0.001	
	0.544	0.140	0.997	-0.087		-633.22	<0.001	1.000
	0.323	<0.001	0.989			-639.77	<0.001	
	0.445	0.033	0.992	0.241		-635.88	<0.001	0.005
	0.004	<0.001	1.000			-629.67	<0.001	
	0.007	0.001	1.000	0.334		-627.09	<0.001	0.023
	0.004	<0.001	1.000		1.000	-629.67	<0.001	0.979
	0.007	0.001	1.000	0.334	1.000	-627.09	<0.001	0.997

Table 2: Estimation Results for EWA Models

ID	β_1	β_2	ϕ	l	δ	log-likelihood	Test vs. 0-Model $P(>\chi^2)$	Test for l or δ $P(>\chi^2)$
41	7.134	5.115				-337.32		
	0.953	1.966	1.000			-359.90	1.000	
	0.577	<0.001	1.000	-2.169		-345.83	1.000	<0.001
	0.629	2.407	1.000			-334.89	0.028	
	0.171	2.618	1.000	0.399		-333.08	0.014	0.057
	0.024	0.013	0.985			-364.17	1.000	
	0.043	0.016	0.947	-1.252		-354.19	1.000	<0.001
	0.024	0.015	0.985		0.849	-364.15	1.000	0.843
	0.044	0.016	0.946	-1.273	0.876	-354.18	1.000	0.863
42	4.195	0.342				-724.80		
	0.022	0.271	0.965			-721.63	0.012	
	0.003	0.004	0.976	113.727		-716.53	<0.001	0.001
	<0.001	0.329	0.955			-717.93	<0.001	
	0.006	0.009	0.959	56.966		-707.82	<0.001	<0.001
	<0.001	0.020	0.959			-717.39	<0.001	
	<0.001	<0.001	0.997	27.839		-723.43	0.254	1.000
	<0.001	0.021	0.958		1.000	-717.24	0.001	0.587
	<0.001	0.001	0.935	40.947	1.000	-712.14	<0.001	<0.001
Rep Agent	-8.226	0.594				-39155.90		
	0.268	<0.001	0.992			-38052.18	<0.001	
	0.154	0.201	0.995	-6.500		-35458.16	<0.001	<0.001
	0.110	0.333	1.000			-38144.13	<0.001	
	0.031	0.298	1.000	0.746		-38075.52	<0.001	<0.001
	0.006	0.002	0.987			-37660.29	<0.001	
	0.006	0.003	0.975	-1.336		-37568.47	<0.001	<0.001
	0.009	<0.001	0.986		<0.001	-37256.02	<0.001	<0.001
	0.009	<0.001	0.986	0.125	<0.001	-37256.01	<0.001	<0.001