Exp Econ (2006) 9:323–341
DOI 10.1007/s10683-006-7050-y

Experimental internet auctions with random information retrieval

Doron Sonsino · Radosveta Ivanova-Stenzel

Received: 14 October 2005 / Revised: 1 December 2005 / Accepted: 5 December 2005

Abstract We run an experiment where 97 subjects could retrieve records of completed past auctions before placing their bids in current one-bid, two-bid, and auction-selection games. Each subject was asked to participate in 3 current auctions; but could retrieve up to 60 records of completed (past) auctions. The results reveal a positive relation between the payoffs earned by the subjects and their history-inspection effort. Subjects act as if responding to the average bidding-ratios of the winners in the samples that they have retrieved. They apply intuitive signal-dependent stopping rules like "*sample until observing a winner-value close to my won"* or "*find a close winner-value and try one more history*" when sampling the databases. History-inspection directs bidders with relatively high private-valuations to moderate bidding which increases their realized payoffs. (*JEL* C9 D4 D8)

Keywords Internet-auctions · observational learning · sampling rules · experiments

JEL Classification C93, D44, D83

1 Introduction

Major auction-sites on the Internet typically let potential bidders and sellers examine the results of completed past auctions with no charge. *Ebay* (http://www.ebay.com) for example keeps on site the records of thousands of completed auctions within each product-category; a detailed description of the "bidding-history" is provided upon request.

Electronic Supplementary Material Supplementary material is available in the online version of this article at http://dx.doi.org/10.1007/s10683-006-7050-y.

D. Sonsino (✉)
The School of Business Administration; The College of Management, 7 Rabin Boulevard, POB 9017, Rishon LeZion, Israel 75190
e-mail: sonsinod@colman.ac.il

R. I. Stenzel
Department of economics; Institute for Economic Theory III, Spandauer Strasse 1, 10178 Berlin, Germany
e-mail: ivanova@wiwi.hu-berlin.de

The categorical search-engine of *Yahoo! Auctions* includes a direct *closed-auctions* option (http://list.auctions.shopping.yahoo.com). Categorical-search is also provided by *Amazon.com Auctions* (http://www.amazon.com) and other auction-sites.

The possibility to retrieve records of completed auctions when planning current auction-transactions on the Internet defines a learning scheme that has not been examined in the auctions-literature. In this learning-form, agents retrieve and inspect relevant information, with no monetary charge, before they engage in current transactions. Information-retrieval is discretionary and random; different agents will, in general, observe different samples.

The implications of such "*learning by random information retrieval*" on bidder's behavior and auctions-results seem interesting to explore from a methodological perspective. In particular, it is interesting to examine whether the possibility to inspect historical records of completed auctions may affect participants' behavior in ways that were not observed (and could not be observed) in standard laboratory experiments. Using (Harrison and List, 2004) recent taxonomy, an Internet auction-experiment where subjects may randomly retrieve historical data before placing their current bids may be classified as a (different) *framed-field experiment* where the environment, task and information available to the bidders were modified to approach the internet-auctions field context.[1]

With this underlying motivation we run a controlled experiment intended to examine the impact of discretionary history retrieval on individual bidding-patterns in experimental Web-based auctions. In our experiment, 97 subjects were asked to play 3 different sealed-bid auction-games on the Internet. The experimental auction-sites were linked to databases with records of completed past auctions of the relevant types. Subjects could randomly retrieve records from the databases before placing their current bids.

In particular, our subjects were asked to play (sequentially) the following 3 types of first-price sealed-bid auction-games: (I) *one-bid auction* where each bidder submits a single bid (II) *two-bid auction* where each bidder submits two bids (III) *auction selection game* where each bidder is asked to choose his favorite auction-type: one-bid or two-bid auction, before placing his bid(s). In each of the 3 auctions a virtual good with values independently drawn from $\{50, 51 \ldots 150\}$ was sold to 2 bidders; the specific rules of each auction are described in detail in the next section.

The 3 specific auction-games examined in this paper were recently explored by (Ivanova-Stenzel and Sonsino 2004). To examine the effect of information-retrieval on bidding-patterns we have used the results of the preceding computerized lab-based experiment to construct 2 information bases: a one-bid auctions database containing the records of 100 completed one-bid auctions; and a two-bid auctions database containing the results of 100 completed two-bid auctions. Subjects could randomly retrieve (without repetition) up to 20 records from the relevant databases in each of the 3 auction-games played at the current experiment.

On the bottom-line, our results reveal a robust positive relation between the final payoffs earned by the subjects and their history-examination efforts. Subjects that have achieved relatively high payoffs are characterized as those that invested longer times in history-examination and retrieved a larger number of histories.

A closer examination of the relations between the samples drawn by the subjects and their current bidding patterns reveals that subjects act as if responding to or imitating the winners in their observed samples. The average bidding ratios of the winners in the auctions sampled by each subject significantly explain the current bidding ratios. Moreover, the experimental data

[1] Harrison and List, (2004) definition of "framed field-experiment" assumes a nonstandard subject pool; our Internet-experiment modifies the environment and the task but keeps the standard student subject pool.

suggests that subjects act as if applying intuitive stopping heuristics like "*sample until I reach a winner-value close to my current value*" or "*reach a close-winner-value and try one more history*" when sampling the databases. The bidding ratios of the "*winner with private-value closest to my current valuation*" are strongly correlated with the current bidding ratios (e.g., $\rho = 0.5296$ for the one-bid auctions). We show that these correlations directly follow from subjects' response to their individual samples and demonstrate that "imitation of the winner with closest private-value" directs bidders with relatively high private-values to moderate bidding and increased payoffs.

The motivation for studying the two-bid auction comes from new types of auctions that have recently appeared on the Internet (see Section 5). The motivation for experimenting with the auction-selection game (Monderer and Tennenholtz, 2004) follows from the idea that on the Web sellers may compete at the level of auction-mechanisms; potential bidders may thus choose the auction-type that they prefer. About 60% of our subjects chose the two-bid auction-type in the auction-selection phase of the experiment. The subjects that have preferred the two-bid auction are characterized as "more ambitious", but "not more successful", than those that chose the one-bid auction.

In the preceding study of (Ivanova-Stenzel and Sonsino 2004) 48 subjects were asked to play repeatedly 4 blocks of [6 one-bid auctions followed by 6 two-bid auctions]. This was followed by 16 rounds of the auction-selection game. A comparison of the bidding-ratios in the last rounds of the preceding experiment and the bidding-ratios in the current design reveals that (one-bid) bidding was significantly more aggressive on the Internet. We could not however reject the hypothesis that the distribution of (normalized) winner-payoffs is similar in the two experiments.

A distinctive feature of our Web-auctions experiment is that subjects respond to the experience gained by other (preceding) players. In this sense, the learning-scheme implemented in the current study has a flavor of *observational-learning* or *social-learning* (Bandura, 1986). Merlo and Schotter (2003) demonstrate that "observational learning" may outperform "learning-by-doing" even when the history examined by the subjects in the observational-learning treatment is the one generated by the subjects in the learning-by-doing treatment. Duffy and Feltovich (1999) show that the observation of others may affect the evolution of repeated play even when subjects also observe their private histories. We could not trace any preceding experimental evidence on the use of signal-dependent stopping rules or sampling rules in the retrieval of public information. Evidence on the use of optimal stopping rules in various wage-search problems, is found in the earlier studies of (Braunstein and Schotter 1981, 1982).

The paper proceeds as follows: Section 2 provides a detailed description of the experiment; Section 3 presents the bottom-line results on the positive relation between payoffs and inspection efforts. Section 4 provides a detailed analysis of subjects' behavior when playing the one-bid auction games; Section 5 proceeds with a brief discussion of the two-bid auction. In Section 6 we examine the distinctive characteristics of the subjects that chose the one-bid auction (versus those that preferred the two-bid auction) in the auction-selection game. Section 7 concludes.

2 The experiment

2.1 The auctions

Subjects were asked to participate in 3 auction-games. In each game, a virtual (imaginary) product was sold to a pair of subjects. The *private value* of the product to each subject was randomly selected from the set {50,51,52.... 148,149,150} with equal probabilities for each value. Values were independently drawn for each subject and each type of auction. Subjects were told (see the instructions in Appendix A) that their partner for each auction would be randomly selected from the pool of participants (with independent matching for each auction).

The first type of auction tested was the standard sealed-bid first price auction. In this auction, subjects were asked to submit a single offer (bid) for the imaginary product. Bids were restricted to integer-numbers between 0 and 200. The instructions explained that the highest bidder would win and receive a payoff that is equal to the difference between her private value and her bid. The loser will not pay or receive anything. In case of tie, the winner will be selected randomly.

The second type of auction studied was a similar first-price sealed-bid auction where participants are asked to place two different offers: a high-bid and a low one (integers between 0 and 200). The instructions explained that the bidder that has placed the highest bid would win the auction and receive his private value. If the low-bid of the winner was higher then the high-bid of the second bidder, the winner would pay his low-bid. In all other cases, the winner would pay his high-bid. In case of tie, the winner will be randomly selected. Following (Ivanova-Stenzel and Sonsino 2004) we term these auctions "one-bid auction" and "two-bid auction" accordingly.

The third auction studied was an auction-selection game: after observing their private values, subjects were requested to choose their favorite auction type (one-bid or two-bid auction) and to place a single bid or two bids accordingly. The instructions clarified that the random matching procedure for this phase of the experiment would guarantee that the two partners in each pair have selected the same type of auction.[2]

2.2 History databases

Subjects did not receive any feedback along the experiment. The instructions explained that the results of the three auctions and final payoffs will be reported by email at the end of the experiment. During each phase of the experiment however subjects could inspect the records of previous auctions of the relevant types. While participating in the one-bid auction subjects could access a "*one-bid auctions database*". In the two-bid auction phase, subjects could similarly retrieve past-records from a "*two-bid auctions database*". In the auction-selection stage subjects could retrieve more histories from both databases (see Appendix B for the translated auction-selection screen).

The records for each database were randomly selected from the one-bid and two-bid auctions conducted by (Ivanova-Stenzel and Sonsino 2004).[3] Upon clicking the databases a smaller window appeared with details on the values, bids and payoffs in a previous auction of the corresponding type. Subjects could inspect the history-window for as long as they chose and then move on to "*retrieve one more history*" or "*return to the current auction*". Subjects could retrieve up to 20 one-bid histories at the one-bid auction phase; up to

[2] Originally, we had 110 subjects. Thirteen subjects did not submit bids in at least one of the 3 auctions and were thus removed from the sample. Since the remaining number was odd (97), we randomly selected one subject for double matching. The identity of the balancing subject was separately determined for each auction-type; actual payoffs were determined by randomly selecting one of the two auctions.

[3] As mentioned above the preceding lab-experiment consisted of two phases. In the first, 48 subjects played 4 blocks of [6 one-bid auctions followed by 6 two-bid auctions]; in the second, they played 16 auction-selection games. The records for the databases were randomly retrieved from all auctions played at the first phase of the experiment.

20 two-bid histories at the two-bid auction phase and up to 10 more histories of each type in the auction-selection phase. Sampling was without repetitions (for each subject) and independent across subjects. Every record not yet inspected by a given subject was equally likely to be drawn at the next retrieval. The instructions emphasized that history inspection is completely discretionary (see Appendix A). Subjects could go back and forth from the current auction screen to the databases up to the point where they have inspected the maximal number of histories. In addition, subjects could repeatedly click a "*summary of auctions inspected*" link. The summary screen included two chronicle tabulated lists: one with the values, bids and payoffs in each of the one-bid auctions inspected by the subject and a second similar list for all the two-bid auctions retrieved by the subject. The summary screen also disclosed the average winner-payoff recorded for each type of auction (in the sample drawn by the subject).

2.3 Method

Ninety-seven subjects were recruited by distributing advertisements all over campus at the Technion, Israel Institute of Technology. In the ads students were asked to send an email to a given address in order to participate in an experiment in decision-making on the Internet. The ads mentioned that the experiment involves participation in 3 auctions and that participants may earn nice payoffs depending on their decisions. A minimum payoff of 15 NIS was guaranteed (the exchange rate at the time of the experiment was about 4.2 New Israeli Shekel to 1 U.S. dollar).[4] In response to their emails, subjects received an individual user name, password and the URL on which the experiment was mounted. Subjects could log on to the experiment independently from their own computer, at their convenience; no constraints were imposed on the place and time of participation.[5]

The experiment was programmed on HTML using Java script. After reading the instructions subjects were asked to fill in a short personal questionnaire. In addition to standard personal details subjects were also asked to state their preference between (1) a sure payoff of 49 NIS and (2) a lottery paying 100 NIS with probability 50% (and 0 otherwise). The binary choice problem was included to control for individual risk preferences.

No time limits were imposed in any stage of the experiment. On average, subjects have spent about 10 minutes on the 3 tasks (this does not include the time spent on reading the instructions and filling in the personal questionnaire); 9 subjects (of 97) have spent more than 20 minutes on the 3 auctions. The order with which the one-bid and two-bid auctions were presented was randomly selected for each pair: 49 subjects played the one-bid auction before the two-bid type (order 12); 48 subjects played order 21. The auction-selection game was always played last. Subjects were told that at the end of the experiment they would receive a check in the amount of their total payoffs in the three auctions (or the guaranteed 15 NIS minimal payoff when applicable). Checks were either sent

[4] We have used a guaranteed-minimum (rather than an independent participation-fee) since subjects that overbid their values may end-up with negative payoffs. Since subjects did not observe their payoffs until the end of the experiment, we do not believe that "limited liability" could significantly affect subjects' behavior in our design.

[5] See (Birnbaum, 2000) for a general discussion of experimentation on the Web. The current experiment could be run in a computerized lab; this would have the advantage of tighter control on the subjects. A lab-implementation however may generate other problems; e.g., it might direct the subjects into intensive sampling.

by regular mail to the provided address or collected personally in our offices.[6] The average payoff was 35.9 NIS (about 8.6 dollars). The maximal payoff was 138 NIS (about 32 dollars).

3 Performance and history inspection

Let *TP* denote the Total-Payoffs earned by each subject in the three phases of the experiment (before imposing the 15 NIS minimum). Let *HIT* denote the *History Inspection Time* of the subject. In particular, *HIT* is the time where windows with records of completed one-bid or two-bid auctions (retrieved from the databases) were opened on the subject' screen. Analysis of the correlation between *TP* and *HIT* reveals a positive significant relation between the two variables; the Pearson coefficient of correlation between the variables, for instance, was 0.2339 (p = 0.0211). The median *HIT* of the 48 subjects with $TP < 27$ was 92 compared to a median *HIT* of 191 to the 49 subjects with $TP \geq 27$ (Robust rank-order test; $z = 3.6027$; $p < 0.001$).[7]

A positive relation between payoffs-earned and inspection-efforts also appears when we examine the *Number of History Inspections*. The median number of histories inspected by the 48 subjects with $TP < 27$ was 5.5 compared to a median of 13 history-inspections for the 49 subjects with $TP \geq 27$ ($z = 2.8007$; $p = 0.0025$); the Pearson coefficient of correlation between the number of histories inspected and *TP* is positive but does not reach significance ($\rho = 1038$; N.S).[8]

Closer look into the data suggests that subjects have put more effort into history-examination when playing the two-bid auction. The median *HIT* in the one-bid auction phase was 46 compared to median *HIT* of 67 in the two-bid auction phase of the experiment (Wilcoxon test; $z = 2.186$; $p = 0.0144$). An intuitive explanation is that the two-bid auction is conceived more complicated and thus attracts more examination. Moreover, the inspection effect on performance appears stronger in the two-bid auction case: The Pearson coefficient of correlation between *HIT* and winner-payoff in the one-bid auction phase was 0.0859 (N = 49; N.S.) while the corresponding coefficient of correlation for the two-bid auction phase was 0.3978 (N = 49; p = 0.0023).[9]

The positive relation between inspection and performance is robust to re-matching of the subjects; i.e., similar results were obtained when subjects were repeatedly re-matched to different pairs. However, the effect may still follow from a selection bias; e.g., subjects that end up with higher payoffs might be in general "more ambitious", have higher valuations,

[6] Subjects were asked to choose their preferred payment method while providing the personal information.

[7] Note on statistical testing: since some of the between-sample comparisons that we make involve samples with different dispersion we use the robust rank-order test to compare the distributions of independent samples; the Wilcoxon signed-ranks test is used in the case of paired data (see (Siegel and Castellan, 1988)). Since most of our tests are "directional", we report the one-tail significance levels; the abbreviation $N.S.$ is used in cases where the one-tail significance level is higher than 0.05. Test-statistics are always disclosed in absolute value.

[8] Positive correlations also appear between *TP* and summary-screen inspection: The median summary-screen inspection time of the subjects in the $TP < 27$ category was 0 (only 19 of the 48 subjects in this category approached the summary screen) compared to a median summary-screen inspection time of 21 for the 49 subjects with $TP \geq 27$ (robust rank-order test; $z = 3.3914$; $p < 0.001$).

[9] To calculate these coefficients we restrict the sample to the subset of winners; a positive significant relation between history-inspection and one-bid auction payoffs appears in other examinations.

lower opportunity cost of time, or otherwise have qualities that make them use the databases more intensively than subjects with lower payoffs.

Direct examinations of observable confounding variables did not reveal any significant correlations: Let *TT* for instance denote the total time spent on the experiment; *TT-HIT* then measures the "*time not spent on history inspection.*" Examinations of relation between *TT-HIT* and *TP* does not reveal significant correlation. The Pearson coefficient of correlation between the two variables, for instance, was 0.0819 (t = 0.8; N.S.). The positive link between total payoffs and *HIT* thus did not follow from general higher-effort-level of the inspecting subjects. The coefficient of correlation between realized values (v) and phase-level *HIT*, to take another possible confound, was negative ($\rho = -0.1240$;N.S) for the one-bid auctions and close to zero ($\rho = 0.0476$;N.S) for the two-bid auctions. The positive relation between payoffs and history-retrieval thus did not follow from heavier sampling by subjects with higher valuations. Similar insignificancies were obtained for other potentially-confounding observables.

To argue that our results cannot be attributed to hidden selection-bias, note that the next sections provide a direct explanation to the positive effect of information retrieval on realized-payoffs. In particular, we show that history-retrieval has guided subjects into careful bidding that could increase their realized gains upon winning. In estimating the relation between observed historical data and current biding patterns we also run (Heckman's, 1978) two step regressions to check for possible selection effect; the estimations do not reveal a significant bias.

4 Results for One-bid auctions

To analyze subjects' behavior in the one-bid auctions we need the following notation:

Let v and b denote the realized value and corresponding bid of a representative subject. We use $BR = \frac{b}{v}$ to denote the *current bidding ratio* of the subject.

Let $h = (v_W, b_W, v_L, b_L)$ describe the records of a completed one-bid auction h; $v_W(h)$ and $b_w(h)$ are the value and bid of the winner, $v_L(h)$ and $b_L(h)$ accordingly denote the value and bid of the loser. For each such history h, let $WBR(h) = \frac{b_W(h)}{v_W(h)}$ denote the *winner bidding ratio* in h; $LBR(h) = \frac{b_L(h)}{v_L(h)}$ similarly denotes the *loser bidding ratio* in that auction.

Let $r = h_1, h_2 \ldots h_n$ be an arbitrary *sample (retrieval)* of n histories.

We use $AWBR(r) = \frac{1}{n} * \sum_{i=1}^{n} WBR(h_i)$ to denote the *average observed WBR* in r. Similarly, $ALBR(r) = \frac{1}{n} * \sum_{i=1}^{n} LBR(h_i)$ is used to denote the *average observed LBR* in r.

We call $WBR(h_n)$ the *last-observed WBR* in r; accordingly $LBR(h_n)$ is addressed as the *last-observed LBR* in r. Similarly, $WBR(h_{n-1})$ and $LBR(h_{n-1})$ are called the *before-last-observed* bidding ratios (for $n \geq 2$); $WBR(h_1)$ and $LBR(h_1)$ are addressed as the *first observed bidding ratios* (the last and first are equal when $n = 1$).

The next paragraphs provide a detailed analysis of the relation between the contents retrieved by the subjects and their current bidding patterns. The paragraphs are built sequentially where the results of each section motivate the analyses that follow.

4.1 Sampled data and current behavior

Table 1 gives the Pearson coefficients of correlation between the *current bidding ratio BR* and the *first, average, last, and before-last observed WBR* and *LBR*. The number at the upper-left corner of the table (0.0956) for example is the correlation between *BR* and *WBR*(h_1).

Table 1 Correlation Between *BR* and sample-statistics

	WBR	LBR
First observed ratios ($N = 74$)[10]	0.0956 ($N.S.$)	−0.0554 ($N.S.$)
Average observed ratios ($N = 74$)	0.3871 ($p < 0.001$)	0.1302 ($N.S.$)
Last observed ratios ($N = 74$)	0.1285 ($N.S.$)	0.1486 ($N.S.$)
Before-last observed ratios ($N = 64$)	0.3396 ($p < 0.005$)	0.356 ($p < 0.005$)

The numbers on the table suggest that subjects responded to the average observed bidding ratios of the winners in their samples. The coefficient of correlation between *AWBR* and *BR* was 0.3871 and highly significant ($t = 3.562$). An intuitive interpretation is that the subjects (act as if) trying to "imitate" the winners in their observed samples.[11] To test this interpretation, we have randomly retrieved for each subject a hypothetical sample of the same size as actually inspected by that subject. We have used the hypothetical samples to calculate a hypothetical *AWBR*. The coefficient of correlation between the hypothetical *AWBR* and *BR* was not statistically significant ($\rho = -0.0599$). The insignificance was replicated in repeated hypothetical sampling.

The *last observed* bidding ratios were not significantly correlated with the current ratios (see the numbers on the table). However, the correlation between the *before-last-observed WBR* and the current bidding-ratio was highly significant (0.3396; $t = 2.84$). An intuitive explanation is that subjects "*try just one more history*" after reaching a point where they retrieve an interesting or a relevant history.

The correlation between the current bidding ratios and the bidding ratios of the *losers* in the samples retrieved by the subjects was not statistically significant at the average level of *ALBR*. Note however the significant correlation (0.356) between the *before-last observed LBR* and the current *BR*. A possible explanation is that low bidding ratios of the losers may induce cautious current bidding by subjects that expect to win their auctions (see (Ockenfels and Selten, 2005) for the "downward impulse" associated with losing bids).[12]

Linear *ad hoc* estimations of the relation between observed bids and sample-statistics for the 74 inspecting-subjects confirm the results of the correlation analysis. OLS estimation of the log-linear model $\ln(b) = \alpha + \beta^* \ln(v)$ gives $\beta = 0.834$ with standard deviation 0.0467; the hypothesis $\beta \geq 1$ is strongly rejected for the alternative $\beta < 1$ $(t = 3.552)$ which confirms the concave relation between bids and values.[13] Adding

[10] Twenty-three subjects did not retrieve data at the one-bid auction phase and thus N = 74 for the first, last and average ratios; 10 subjects have only retrieved 1 record and thus N = 64 for the before-last sampled ratios.

[11] The median value of *BR* for the 74 inspecting subjects was 0.8408 compared to a median *AWBR* of 0.8345 (Wilcoxon signed-ranks test; z = 0.2289; N.S).

[12] The correlations between the bidding ratios observed two periods (or three periods) before the end of history-inspection and *BR* were not statistically significant.

[13] A concave relation between bids and values was observed in many preceding experiments; see for example (Kagel and Roth, 1992). See also (Battigalli and Siniscalchi, 2003) for a model with rationalizable concave bidding functions.

ln(AWBR) as an additional explanatory variable gives a positive significant coefficient $\gamma = 0.706$ ($t = 2.94$; $p = 0.0022$).[14] When *ln(AWBR)* is replaced with *ln(ALBR)*, the estimated coefficient for *ln(ALBR)* is not statistically significant: $\gamma = 0.083 \quad (t = 0.608)$. Similar results are obtained in alternative model-specifications; e.g., OLS estimation of the bidding-ratio model $BR = \alpha + \beta^* v + \gamma^* AWBR$ gives $\beta = -0.001$ (t = 3.47), $\gamma = 0.643$ ($t = 2.79$); the negative coefficient for v reflects again the less-aggressive bidding of subjects with higher-valuations.[15]

4.2 Sampling effects

Let h be an *arbitrary* retrieval from the one-bid auctions database (the emphasis on the arbitrariness of h will be made clear below). Recall that $v_W(h)$ and $v_L(h)$respectively denote the private valuations of the winner/loser in auction h. Let v denote the current valuation of the sampling-subject on the Internet. Since h is arbitrary, we assume that $v_W(h)$ and $v_L(h)$ are independent of v. A numeric calculation then shows that (when h is randomly drawn from the 100 auctions included in the one-bid auctions database and v is independently drawn from $\{50,51\ldots 150\}$) the expected value of the distance $\overline{d}(h) = |v_W(h) - v|$ is 33.56 while the expected value of the distance $\underline{d}(h) = |v_L(h) - v|$ is 33.62.[16]

Note next that each retrieval $r = h_1, h_2 \ldots h_n$ defines two sequences of distances: the sequence $\overline{d}(h_1), \overline{d}(h_2) \ldots \overline{d}(h_n)$ and a corresponding sequence $\underline{d}(h_1), \underline{d}(h_2) \ldots \underline{d}(h_n)$. It immediately follows that if the number of histories retrieved by the subject (n) is independent of the contents of the histories retrieved (e.g., the number of histories inspected is fixed or it is described by a random variable N that takes the values 1,2 … 20 with given probabilities) then the expected average-distances, first-distances, last-distances, and before-last distances should also be around 33.56/33.62 respectively.

The actual distance-statistics are provided on Table 2.

As expected from random retrieval, the distances between the first-observed winner/loser valuations and the current valuations (see the first line on the table) are not significantly different from the expected values. The median distance $\overline{d}(h_1) = |v_W(h_1) - v|$ for instance is 29; a Wilcoxon signed-ranks test of the hypothesis that distances are equal to random-sampling benchmark (33.56) gives $z = 1.3$; $N.S.$

The second line on Table 2 however reveals that the average distances $\overline{d}(r)$ are significantly lower than the 33.56 benchmark. The median average-distance $\overline{d}(r)$ is 27; the signed-ranks test gives $z = 2.08$; $p = 0.0119$ which contradicts the independent sampling hypothesis.

A possible explanation is that subjects (act as if) applying *stopping rules* or *sampling rules* under which the number of inspections depends on the contents of the retrievals. If these sampling rules direct the subjects towards samples with winner-values close to their

[14] To account for a possible selection bias we have also run the model using the (Heckman, 1978) two-step regression model where the first step assumes that the decision whether to inspect historical data depends on different instruments collected in the experiment and the second step estimates the equation of interest while taking into account the possibility of correlation in errors. The two step model was run using the QLIM procedure on SAS v9. The results of the regression did not reveal a significant selection bias ($\rho = 0.0520$; N.S); the estimated parameters were similar to the ones obtained in the standard estimation.

[15] The coefficient for the intersect α is about 0.55 and statistically significant at $p < 0.01$ in all models; The R^2 levels are 0.9031 for the log-linear mode; 0.9141 for the extended model with *ln(AWBR)* and 0.5231 for the bidding-ratio model.

[16] As a formal approximation, assume the 3 valuations are independently drawn from the interval [50,150]. It is easy to show that the expected values of $\overline{d}$ and $\underline{d}$ in this case are 33.33.

Table 2 Median distances between observed winner/loser values and current values

	$\overline{d}$	$\underline{d}$
First distances (N=74)	29 ($z = 1.3$) $N.S$	30 ($z = 0.18$) $N.S$
Average distances (N=74)	27 ($z = 2.08$) $p < 0.02$	30 ($z = -0.39$) $N.S$
Last distances (N=74)	23 ($z = 2.1$) $p < 0.02$	27 ($z = 0.77$) $N.S$
Before-last distances (N=64)	30.5 ($z = 0.307$) $N.S$	30.5 ($z = 0.314$) $N.S$

current valuations then the average distances between the observed winner-valuations and the current valuations would indeed be lower than the random-sampling benchmark. Specifically, subjects may (act as if) applying stopping rules of the form: "*sample until drawing a winner-value that is close to your current valuation*".[17] Aggregate-evidence supporting this behavior may be found in the fact that the last-distances $\overline{d}(h_n)$ are also significantly lower than 33.56 (see the third line on the table).

The distances between the observed loser-valuations and the current valuations ($\underline{d}$) were not significantly different from the 33.62 benchmark (see the right column on Table 2). The fact that the directed-sampling evidence appears for winner-valuations but does not appear for loser-valuations seems compatible with the fact that subjects appear as if responding to *WBR* rather than *LBR* in their current bidding. In both dimensions subjects appear as if trying to imitate or learn from the winners in their samples.

4.3 Closer look at individual samples

A major advantage of experimentation on the Web is the ability to measure and compare the time spent on different screens. On Table 3 we use $T(h)$ to denote the time invested in the examination of history h. $T(h_n)$ and $T(h_{n-1})$ accordingly denote the time invested in examination of the last and before-last observed histories. Inspecting-subjects are divided into 5 groups as follows: (a) those inspecting exactly 1 history (b) inspecting exactly 2 histories (c) inspecting more than two histories and spending at least 5 more seconds on the last history examined (compared to the preceding one) (d) inspecting more than two histories and spending at least 5 seconds less on the last history examined (compared to the preceding one) (e) inspecting more than two histories and spending similar times on the last two histories (up to 5 seconds of difference). We disclose the median distance between

[17] A quite similar stopping rule would be "sample until drawing a *maximal-value* that is close to your current valuation". The distinction between winner-values and max-values is required since about 11% of the auctions inspected by the subjects (78 of 702) resulted in an inefficient allocations (the winner's valuation was lower). We state the rule in terms of "winner values" rather than "max values" following the evidence on the significant relation between *AWBR* and *BR*. Note that the correlation between *BR* and the average "bidding ratio of bidder with maximal valuation", 0.3205, was similar, but lower, than the correlation between *BR* and *AWBR*: 0.3871.

Table 3 Inspection times and sampling effects

Group	N	I $\lvert v_w(h_n) - v\rvert$	II $\lvert v_w(h_{n-1}) - v\rvert$
(a) Inspecting 1 history	10	26	NA
(b) Inspecting 2 histories	9	28	45
(c)$T(h_n) - T(h_{n-1}) \geq 5$	19	19	40
(d) $T(h_{n-1}) - T(h_n) \geq 5$	8	34	18.5
(e) $\lvert T(h_n) - T(h_{n-1})\rvert < 5$	28	23.5	22

current values and last-observed winner values (Column I) and the median distance between current values and before-last observed winner values (Column II) for each group.

The table shows that subjects that have invested more time inspecting h_n (compared to the time invested in the inspection of h_{n-1}) drew a "close winner value" at the last retrieval (see the data for group c).[18] The subjects that have invested more time inspecting h_{n-1} (group d), on the other hand, had a "close draw" at the before-last retrieval. The data for the subjects that have examined exactly two histories (group b) is presented separately since these subjects drew a closer winner value in their second, last draw (see median distances on the table) but have invested longer time inspecting the first, "before last" record. Longer examination of first-retrieved histories indeed appears for 49 of the 64 inspecting subjects (76.6%).

4.4 *CWBR*

The results above motivate the definition of another statistic which we term *CWBR* (for "Closest *WBR*"). In words, *CWBR* is the bidding ratio of the winner in the history where the observed winner-value is closest to the subject's current valuation (amongst all the histories retrieved by that subject). We term the corresponding winner value: *CWV*. Formal definitions follow:

Let $r = h_1, h_2 \ldots h_n$ be the retrieval of agent i with valuation v.
Let $m = \arg\min_{i=1}^{n} \lvert v_w(h_i) - v\rvert$
Then $CWV(r) = v_w(h_m)$ and $CWBR(r) = WBR(h_m)$

An examination of the experimental data reveals a highly significant coefficient of correlation between *CWBR* and *BR*: $\rho = 0.5296$ ($N = 74; t = 5.2988; p < 0.001$).[19] Recall, for comparison, that the coefficient of correlation between *AWBR* and *BR* was 0.3871.The coefficient of correlation between our hypothetical *CWBR* (see the discussion in Section 4.1) and *BR* was 0.0855.

Since *CWBR* = *AWBR* for the 10 subjects that have only inspected one history it is interesting to restrict the sample to the 64 subjects that have inspected at least 2 histories.

[18] Recall that the coefficient of correlation between *BR* and the last observed *WBR* was not statistically significant for the complete sample ($N = 74$; see the data on Table 1). A significant coefficient of correlation between the two variables however appears when the sample is restricted to the 32 subjects for which $\lvert v_w(h_n) - v\rvert \leq 20$: $\rho = 0.3019; t = 1.73; p < 0.05$.

[19] A case-based interpretation would say that the "cases" retrieved by the subjects are weighted by the similarity of the observed winner-values to the current private-valuations (Gilboa and Schmeidler, 1995).

Table 4 Inspection-effect on subjects with highest/lowest values

	Median BR (winner-payoff) of 10 subjects with fewest retrievals	Median BR (winner-payoff) of 10 subjects with most retrievals	Robust rank-order test
20 Subjects with highest values	0.8270 (19.5)	0.7413 (32)	P < 0.05 $p \approx 0.05$
20 Subjects with lowest values	0.8524 (NA)	0.8891 (NA)	N.S

The coefficient of correlation between *CWBR* and *BR* for these subjects is 0.5367 ($t = 5.009$; $p < 0.001$).

Linear estimations (similar to those reported at the end of Section 4.1) confirm the significance of *CWBR* in explaining subjects' current bidding patterns. The extended log-linear model $\ln(b) = \alpha + \beta^* \ln(v) + \gamma^* \ln(CWBR)$ gives $\beta = 0.941$ ($t = 18$) and $\gamma = 0.506$ ($t = 3.66$). When both *ln(CWBR)* and *ln(AWBR)* are used as explanatory variables, the coefficient for *CWBR* (0.396) is significant at $p < 0.01$ while the coefficient for *AWBR* is not statistically significant.[20] Similar conclusions are drawn when the current bidding ratio (*BR*) is regressed on *AWBR* and *CWBR*.

4.5 *CWBR* and performance

What is the effect of *CWBR*-imitation on current bidding? We start with two observations:

(a) The one-bid-auction records retrieved by the subjects in the current experiment reflect the concavity of bids in values mentioned above. The coefficient of correlation between *WBR* and v_W in the 100 histories included in the *one-bid auctions database*, was negative -0.1823 ($t = -1.83$; $p < 0.05$) which suggests that bidding becomes relatively less aggressive as winner-values increase.[21]

(b) Subjects with larger number of retrievals drew winner values that were closer to their current valuations. In particular, the correlation between the number of histories retrieved and the minimal-distance $|CWV(r) - v|$ was negative -0.4490 ($p < 0.001$).

Compare now two subjects *I* and *II* with relatively low valuations; say $v_I = v_{II} = 60$. Assume that *I* has drawn a large sample and thus ended up with a relatively close $55 \leq CWV(I) \leq 65$; while *II* has retrieved a smaller number of histories and consequently ended up with *CWV(II)* that is farther away from his current valuation; say $CWV(II) \geq 100$. The concavity of bidding then implies that $CWBR(II) < CWBR(I)$ which demonstrates that intensive sampling may instruct low-valued subjects to bid aggressively.

Consider next two subjects *I* and *II* with relatively high valuations; say, $v_I = v_{II} = 140$. Assume as before that *I* has drawn a large sample and thus ended up with a relatively close $135 \leq CWV(I) \leq 145$; while *II* has retrieved a smaller number of histories and thus ended up with $CWV(II) \leq 100$. The concavity of bidding then implies that $CWBR(II) > CWBR(I)$

[20] The coefficient of correlation between *ln(CWBR)* and *ln(AWBR)* is 0.5456.

[21] The average *BR* of winners with $v > 130$ was 0.7790 ($N = 33$) compared to an average *BR* of 0.8238 for winners with $100 \leq v \leq 130$ ($N = 36$) and average *BR* of 0.9033 to winners with $v < 100$ ($N = 31$).

Table 5 CWV and performance

	N	Average number of histories retrieved	Median (average) payoff
Group A: $\mid CWV(r) - v\mid \leq 3$	37	13.56	15 (15.6)
Group B: $\mid CWV(r) - v\mid > 3$	37	5.4	0 (7.27)

which demonstrates that intensive sampling may instruct high-valued subjects to bid moderately.

Small-sample evidence supporting these conjectures is provided on Table 4.

In the first row of the table we look at the 20 inspecting-subjects with highest realized (one-bid auction) values. We sort these subjects by the number of one-bid histories retrieved and find that the 10 subjects with lowest number of retrievals had significantly higher bidding ratios than the 10 subjects with the highest number of one-bid retrievals (see the median ratios on the table). The winner-payoffs made by the subjects in the frequent-inspection category were accordingly larger than those earned by the subjects in the low-inspection group (the median winner-payoff for each group is provided in brackets).[22]

In the second row of the table we run a similar comparison for the 20 inspecting-subjects with lowest realized one-bid values. As anticipated from the explanation above, the direction is reversed and intensive sampling makes bidding more aggressive in this range; the difference here however is not statistically significant (since 19 of the 20 subjects in this category lost the auction in which they participated we could not compare the payoffs across the two groups).[23]

Returning to the complete sample, we find a significant negative relation between $|CWV(r) - v|$ and subjects' realized payoffs. The Pearson coefficient of correlation between $|CWV(r) - v|$ and one-bid payoffs for the 36 inspecting-subjects that have won the one-bid auctions in which they participated was -0.3993 ($t = -2.53$; $p < 0.01$). In Table 5, we divide the 74 inspecting-subjects into two equal-size groups: the group of subjects that have reached a *CWV* that is within 3 units from their current values (the first line on the table) and the group of those with *CWV* that is more than 3 units away from the current value (the second line on the table). The average payoff of the subjects in the first category is twice larger than the average payoff in the second category (see the data on the Table).

4.6 Comparison with the preceding experiment

Recall that in (Ivanova-Stenzel and Sonsino, 2004) 48 subjects played each auction-type for 24 rounds repeatedly. Private values were independently drawn at each round; subjects received "complete" feedback on the results of the game at the end of each auction. The

[22] We have counted 8 winners among the 10 subjects with highest valuations and "least retrievals"; 8 winners among the 10 subjects with highest valuations and "most retrievals" but only 1 winner among the 20 subjects with lowest valuations.

[23] Similar but weaker results were obtained when the samples were extended to include the 30 subjects with highest/lowest private valuations. Note that in order to obtain direct evidence on the results of auctions where both subjects have relatively low (or high) valuations one needs very large samples; e.g., to get 20 pairs with values in the lower 20% of the valuation-range, we should have run about 500 subjects.

Table 6 Median (average) data for web and lab experiments

	Web-experiment	Lab-experiment: Ratios at the last round (round 24)	Lab-experiment: Average ratios over the last block (rounds19–24)
BR	0.8413 (0.8405) $N = 97$	0.8048 (0.7989) $N = 48$	0.8204 (0.7933) $N = 48$
1-WBR	0.1639 (0.1586) $N = 49$	0.1978 (0.1849) $N = 24$	0.1862 (0.1910) $N = 47$

current experiment is drastically different as subjects played each auction only once. It seems interesting however to compare the "bottom-line" results of the two designs.

The median bidding ratio (*BR*) in the last (24-th) round of the preceding experiment was 0.8048 compared to a median *BR* of 0.8413 in the current experiment (see Table 6 for the average data). A robust rank-order test confirms that the Internet bidding-ratios were significantly higher than those recorded in the last round of the lab experiment ($z = 1.9180$; $p = 0.0278$). The conclusion is reconfirmed when we compare the current bidding-ratios to the average bidding-ratios in the last block (rounds 19–24) of the preceding experiment (see the median/average data on the Table).

The more-aggressive bidding on the Web should bring a corresponding decrease in payoff-ratios of successful bidders. The bottom line on Table 6 indeed suggests that the ratio of winner-payoff to winner-value (henceforth: *1-WBR*) was lower on the Internet.[24] It is interesting to note however that the differences are not statistically significant in this case. A robust rank-order test for comparing the *1-WBR* ratios in the last round of the preceding study to the corresponding ratios in the current study, for example, gives $z = 0.8253$ (N.S). It immediately follows that normalized sellers' revenues (*WBR*) were also similar across the two experiments. A possible explanation to the apparent paradox is that since *all bidders* (losers and winners) in the Web experiment imitate the *winners* in the lab experiment, bidding appears overall more aggressive while payoff and revenue ratios do not change significantly.[25] Note also that the efficiency-rate at the Internet experiment (i.e., the proportion of auctions where the buyer with higher valuation has won the auction), 85.7%, was similar to the efficiency rate in the corresponding laboratory experiment, 84.7%. Information-retrieval thus directed subjects into more aggressive bidding but did not affect efficiency and did not change the distribution of surplus between buyers and sellers compared to the lab-based experiment. The

[24] We compare the normalized payoffs (*1-WBR*) and normalized revenues (*WBR*) across the two experiments rather than comparing actual seller's revenues (b) and bidders' payoffs ($v - b$) in order to account for differences in distribution of realized winner-valuations across the two experiments. Comparison of (non-normalized) winning bids and buyer-payoffs also does not reveal significant difference across the two experiments.

[25] The average WBR in the 100 histories that composed the *one-bid auction database* was 0.8337 compared to an average LBR of 0.8005. A Wilcoxon signed-ranks test confirms that WBR is higher than LBR in these 100 histories ($z = 1.7653$; $p < 0.05$). We could not however reject the hypothesis that the bidding ratios of the winners are equal to the bidding ratios of the losers in the Internet study (average WBR: 0.8414; average LBR: 0.8396; robust rank-order test, $z = 1.0535$; N.S.).

results of Section 3 then suggest that history-examination increased the payoffs of motivated subjects on the expense of less-motivated bidders.

5 The two-bid auctions

The motivation for studying the two-bid auction comes from new forms of first-price sealed-bid auctions that have gained popularity in Israel over the last several years. In these auctions each bidder may submit multiple bids although winning is restricted to single units. One of the largest auctions in Israel, for example, "The State Auction", allows bidders to submit up to 3 different price proposals in each auction. The rules of the auction explain that in case where more than one of the offers made by a single bidder wins, the lowest winning offer would determine the actual price that the winner would pay.[26] This is similar to the rules of the two-bid auction where winners pay their low bid when it is higher than the maximal bid of their rival.

The CRRAM equilibrium bidding functions for the two-bid auction are represented in Appendix C. The RNNE equilibrium strategies are $h(v) = 50 + 2/3^{*}(v - 50)$ and $l(v) = 50 + 1/3^{*}(v - 50)$ where $h(v)$ denotes the high RNNE bid and $l(v)$ denotes the low RNNE bid.[27] As in standard one-bid first price auctions, risk-averse agents bid more aggressively than risk-neutral agents in equilibrium of the two-bid auction, where the more-aggressive bidding applies to both high and low bids. The two inequalities are reversed when agents are risk-seeking (see details in Appendix). Comparison of actual bidding patterns to the RNNE benchmark however reveals that 34% of the subjects (33 of 97 subjects) contradicted the equilibrium predictions in the sense of having one bid larger than the corresponding RNNE while the second bid was smaller than the matching RNNE prediction. In particular, we find that 25 subjects have submitted high-bids higher than the RNNE while submitting low-bids that were smaller than the RNNE.[28] Closer examination of the data reveals that many of these subjects chose extremely small low-bids that were close or even lower than the minimal possible valuation, 50. The overall proportion of "*bargain-bidding*" (Ivanaova-Stenzel and Sonsino, 2004) where $b(l) < 50$ on the Internet was 12.3% (12 bids of 97) which is similar the proportion observed in the preceding lab experiment (16.9%).

The strategic shading of low-bids relatively to high-bids in two-bid auctions is intended to reduce the closing price of the auction in cases where the competition for the item is not too intense. In this sense, low-bid shading is similar to "*demand-reduction*" in bidding for secondary units in multi-unit auctions. The experimental literature on multi-unit auctions (see, for example, List and Lucking-Reiley, 2000; Kagel and Levin, 2001) indeed finds strong evidence for strategic demand-reduction in field and laboratory multiunit auctions. Note however that in our experiment the low-bid of more than 25% of the subjects is inconsistently small relatively to their selected high-bid (according to the equilibrium benchmark); we are not aware of similar findings on multiunit auctions.

Analysis of history-retrieval affects in the two-bid auction phase of the Web experiment confirms the results of the analysis for the one-bid auctions. In particular, we find that the high

[26] The State-Auction is both Web-based at http://www.ibuy.co.il/ and published in a booklet that is distributed with leading newspapers. Bidders may submit bids by phone or through the Internet.

[27] The equilibrium strategies were derived for the continuous case where private valuations are independently drawn from the interval [50,150]. Recall that in our experiments values were drawn from the set of integers between 50 and 150; the equilibrium strategies for the continuous case are used as a benchmark for analyzing the experimental data.

[28] 8 subjects had high-bids lower than RNNE and low-bids higher than RNNE

and low bidding ratios of subjects on the Internet are significantly correlated with "*CWHBR*", the high bidding ratio of the winner with "closest private-valuation" in the individual samples drawn by the subjects. A summary of the analysis of bidding in the two-bid auction and at the auction-selection game is available in a separate supplement to this paper.[29]

6 Choice patterns in the auction-selection phase

Recall that in the auction-selection phase, subjects could retrieve more data from both databases. Subjects could also respond to the information retrieved in the preceding phases of the experiment (all inspected records could be re-examined by clicking the "summary screen"). To separate *HIT* across the three phases of the experiment we use *HITn* to denote the history-inspection time at the *n-th* phase; *HIT3* thus denotes the *HIT* at the auction-selection phase.

About 60% of the subjects (58 of 97) chose the two-bid auction in the auction-selection phase; the hypothesis that subjects choose the two-bid auction-type with independent probability 0.5 is rejected at $p < 0.05$ (binomial test; $p = 0.0335$). In the next paragraphs we look for an explanation for the selection patterns of individual subjects.

First we check an equilibrium-behavior type of explanation. The symmetric version of the CRRAM equilibrium model (see Appendix C) predicts that risk-averse agents should show a preference for the two-bid mechanism while risk-seeking agents should prefer the one-bid auction-type. Using the binary-choice control problem to characterize subjects as approximately-risk-seeking (those that preferred a 50% chance for a payoff of 100 on a sure-payoff of 49) or strictly-risk-averse (those that chose the certain payoff), we find that only 58/97 (60%) of the subjects revealed consistent risk-preferences in the binary-choice and auction-selection tasks.[30]

An alternative ("adaptive") explanation would be that subjects chose the auction-type that appeared more profitable in their individual samples. A look at the 65 subjects that retrieved at least one record of each auction-type at the first two-phases of the experiment however reveals that only 64% (42 subjects) chose the mechanism that produced higher average winner-payoffs.

Since subjects observed their private valuations before selecting the auction-type, it is also interesting to check whether the observed values affected the choice. The median v for the subjects that chose the two-bid mechanism was 105 compared to a median value of 98 for those that selected the one-bid auction. A robust rank-order test suggests that the differences are not statistically significant ($z = 0.9316$; N.S.).

Surprisingly, we find a major difference in *HIT1* and *HIT2* across the two groups. The median $HIT1 + HIT2$ of the subjects that chose the two-bid auction at the auction-selection phase was 144.5 compared to a median of 82 for the subjects that chose the one-bid auction (see Table 7 for the phase-level data). A robust rank-order test for comparison of $HIT1 + HIT2$ across the two groups gives $z = 2.38$; $p < 0.01$. The subjects that preferred the two-bid auction at the selection stage thus appear more "ambitious" than those that chose the one-bid auction. These subjects, however, were not more successful then those that chose the one-bid auction type: The median winner-payoffs at the first phase of the experiment

[29] See http://www2.colman.ac.il/business/doron

[30] 20 of the 40 subjects that chose the lottery in the control problem chose the one-bid auction in the selection phase; 38 of the 57 subjects that chose the risk-free alternative in the control problem, chose the two-bid mechanism in the selection game.

Table 7 Median HIT according to selection

	Subjects choosing one-bid ($N = 39$)	Subjects choosing two-bid ($N = 58$)
HIT1	37	56.5
HIT2	47	78.5

(the one-bid auction phase) were: 19 for the subjects "choosing the 1-bid auction" and 21.5 for the subjects "choosing the 2-bid auction" (robust rank-order test; $z = 1.6353$; N.S.). The corresponding figures for the second, two-bid auctions phase of the experiment were 28 and 21 ($z = 0.5992$; N.S).

7 Discussion

We have run a framed-field experiment intended to examine the effect of discretionary historical-information retrieval on bidding behavior on the Web. The examination of subjects' retrieval patterns and the effect of information-retrieval on bidding strategies and performance seems interesting from a methodological perspective as these issues were not yet explored in the literature. As suggested by (Harrison and List, 2004), such framed-field experiments may be considered "methodologically complementary to traditional laboratory experiments" and thus contribute to the understanding of bidders' behavior on the field.

Our main results reveal that subjects in experimental Web-auctions act as if searching for relevant information when sampling public information-bases. The relevance of different information-records is individually determined as it depends on the private signal of the agent. While the stylized auctions examined in this study are quite different from real Internet-auctions (e.g., private valuations are not observable in standard auctions), the directed-sampling result may have practical implications. For example, our result may suggest that bidders search for (what they conceive to be) "similar-auctions" before submitting their bids in Web-based common value auctions. This in turn may imply that bidders on the Web would appreciate sites that provide user-friendly information bases that facilitate selective search for relevant information.

Our results also suggest that the contents of samples retrieved from public databases may significantly affect the behavior of sampling agents. It seems interesting to further explore the relation between available information and actual behavior. For example, how would the results of the experiment be affected if the information database reflected a proportional (e.g., the CRRAM equilibrium bids for symmetric agents with fixed risk preferences) relation between bids and values (instead of the concave relation observed in the current study)? Pushing these issues further, our results may suggest that auctioneers on the Internet might have incentives to manipulate the information available on their completed-records screens.[31]

In the current experiment subjects did not receive immediate feedback on the auctions they have played. The results were only reported in delay after the experiment ended. This design

[31] See the literature on Web-inspired game theoretic models (e.g., Monderer and Tennenholtz 1999; Baye and Morgan 2001; Wellman et al., 2001).

was chosen to simplify the task and increase control on the type of information processed by the bidders; it may represent cases where individual participation in given-auctions is rare so that the impact of personal experience on bidding behavior is marginal. It seems interesting to further examine the integrated affects of learning-by-random-information-retrieval and learning-from-personal experience in a dynamic setting where subjects participate in several auctions sequentially and may react to their personal experience as well as randomly retrieve records from public information-bases.

The fact that winner-payoffs and efficiency rates were not significantly smaller than those observed in the preceding lab-based experiment in spite of more-aggressive bidding suggests that online information-bases may help decision-makers act affectively in the competitive markets on the Internet.

Acknowledgments We have benefited by comments made by Rachel Barkan, Ido Erev, Werner Güth, Todd Kaplan, Axel Ockenfels, Bradley Ruffle, Carsten Schmidt, Moshe Tennenholtz and participants in seminars at Bar-Ilan, Ben-Gurion, University of Exeter, Max-Planck Institute in Jena and Tel-Aviv University. We also thank 2 anonymous referees for important suggestions. The research was supported by the Israel Science Foundation grant number 191–364. Financial support from the fund for the promotion of research at the College of Management is also acknowledged.

References

Bandura, A., (1986). *Social Foundations of Thought and Action – A Social Cognitive Theory.*Prentice Hall.

Battigalli, P. & Siniscalchi, M. (2003). "Rationalizable Bidding in First-Price Auctions," *Games and Economic Behavior, 45*(1), 38–72.

Baye, M. R., & Morgan, J. (2001). "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review, 91*(3), 454–474.

Braunstein, Y. M., & Schotter, A. (1981). "Economic Search: An Experimental Study," *Economic Inquiry*, *19*, 1–25.

Birnbaum, M. H. (2000). *Psychological Experiments on the Internet.* San Diego. Ca. Academic Press.

Braunstein, Y. M., & Schotter, A. (1982). "Labor Market Search: An Experimental Study," *Economic Inquiry*, *20*, 133–144.

Cox, J. C., Smith, V. L., & Walker J. M. (1988). "Theory and Individual Behavior in First-price Auctions," *Journal of Risk and Uncertainty*, *1*, 61–99.

Duffy, J., & Feltovich, N. (1999). "Does Observation of Others Affect Learning in Strategic Environments? An Experimental Study," *International Journal of Game Theory, 1*, 131–152.

Gilboa, I., & Schmeidler, D. (1995). "Case-Based Decision Theory," *Quarterly Journal of Economics, 110*, 605–639.

Harrison, W. G, & List, J. A. (2004). "Field Experiments," *Journal of Economic Literature*, *42*, 1009–1055.

Heckman, J. (1978). "Dummy Endogenous Variables in a Simultaneous Equation System," *Econometrica*, *46*, 931–961.

Ivanova-Stenzel R., & Sonsino, D. (2004). "Comparative Study of One-bid vs Two-bid auctions," *Journal of Economic Behavior and Organization*, *54*, 561–583.

List, A. J., & Lucking-Reiley, D. (2000). "Demand Reduction in Multiunit Auctions: Evidence from a Sportscard Field Experiment," *American Economic Review*, *90*, 961–972.

Kagel, J., & Roth, A. (1992). "Theory and Misbehavior in First-Price Auctions: Comment," *American Economic Review*, *82*, 1374–8.

Kagel, J., & Levin, D. (2001). "Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickery Auctions, " *Econometrica*, *69*, 413–454.

Merlo, A., & Schotter, A. (2003). "Learning by not Doing: An Experimental Investigation of Observational Learning," *Games and Economic Behavior, 42*(1), 116–136.

Monderer, D., & Tennenholtz, D. (2004). "K-Price Auctions: Revenue Inequalities, Utility Equivalence and Competition in Auction Design," *Economic Theory*, *24*, 255–270.

Monderer, D., & Tennenholtz, D. (1999). "Distributed Games," *Games and Economic Behavior, 28*, 55–72.

Ockenfels, A., & Selten, R. (2005). "Impulse Balance Equilibrium and Feedback in First Price Auctions," *Games and Economic Behavior, 51*, 155–170.

Siegel, S., & Castellan, N. J. (1988). *Non-parametric Statistics for the behavioral Sciences.* Mac-Graw-Hill.

Wellman, P. W., Walsh, W. E., Wurman, P. R., & MacKie-Mason, J. K. (2001). "Auctions Protocols for Decentralized Scheduling", *Games and Economic Behavior*, 2001, 35, 1/2, 271–304.