The Journal of Risk and Uncertainty, 30:3; 241–259, 2005

Informational Asymmetries and Observational Learning in Search

LIRAN EINAV Leinav@stanford.edu
Department of Economics, Stanford University, Stanford, CA 94305-6072, USA

Abstract

As economics modeling moves from super rational decision makers to considering boundedly rational agents, some economic problems deserve a second look. This paper studies the effects of learning on the efficiency of search. Once learning is taken into account, the structure of information flow becomes important. In particular, I highlight the truncated information structure in the search problem. Agents stop searching once a sufficiently attractive price is found. Therefore, they observe the performance of shorter searches, but do not directly observe the performance of longer searches. I design and conduct an experiment to test the hypothesis that this asymmetric flow of information leads agents to search too little. I find strong evidence in its favor. This suggests that in the presence of learning, the provision of a more symmetric information structure will make search more efficient.

Keywords: search, learning, experiment, regret

JEL Classification: C91, D83

Nowadays it is widely agreed that learning and other behavioral regularities are important in many economic settings. The importance of learning raises a full set of new questions related to the speed at which learning takes place. In a repeated interaction, if learning is fast, equilibrium analysis is of great importance. If learning is slow, however, many real world observations may be of off-equilibrium actions, and hence our prediction and analysis should take a different path.[1]

As economics modeling moves from super rational decision makers to considering boundedly rational agents, some economic problems deserve a second look. This paper studies the effects of learning on the efficiency of search. The solution to the search problem is straightforward for super rational decision makers, who compute the optimal reservation price and search accordingly.[2] Boundedly rational agents, however, cannot compute the optimal reservation price in advance, and hence adjust their search strategy over time and learn from experience.

Once learning is taken into account, the structure of information flow becomes important. I suggest that strategies used by agents are more likely to be reevaluated once it was actually observed that a different strategy could have done better. A more frequent reevaluation leads, in turn, to faster learning. Hence, learning is likely to be faster when it is possible to observe hypothetical earnings from strategies that were not actually played. Whether or not agents observe these hypothetical earnings depends on the specific information flow.

This paper highlights the truncated information structure in the search problem. When agents search for the highest price, the search ends once a sufficiently high price is found. The typical cutoff stopping rule makes agents observe only part of the price distribution, truncated from above by the final transaction price. Such information does not allow agents to symmetrically assess and evaluate their strategies, and hence slows their learning process. Therefore, as agents search, they obtain information about the performance of shorter searches, but not about the performance of longer searches. Consequently, agents who search too much will observe the benefits of shorter searches; they will be able to adjust their behavior and to search less in subsequent searches. In contrast, agents who search too little do not get the same opportunity to observe their errors, and will not be as likely to correct it and search more in subsequent searches. The hypothesis suggests that a more symmetric information flow will significantly increase the speed of learning and therefore improve the efficiency of search. This idea is consistent with directional learning theory, as suggested by Selten and Buchta (1994).

I test this hypothesis and its implications in an experimental setup. As discussed in more detail in Section 1, there is one common phenomenon that stands out from the vast majority of experimental papers about search behavior: subjects in experiments search too little when compared to the optimal search strategy.[3] Section 1 discusses three potential reasons for this result. The first is risk aversion, which provides an incentive to avoid risky search and stop "too early." The second is an asymmetric cost structure around the optimal search strategy: searching too much would typically be much more costly than searching too little. The third potential reason for the too little search result is a direct implication of the hypothesis discussed above, which is the focus of the paper. It suggests that the truncated information structure of the search problem may be a driving force behind the robust "too little search" result.

Section 2 describes the experimental design that tests this possible explanation. The control group uses the standard search environment, in which subjects do not get any further feedback once they stop searching.[4] In contrast, subjects in the other experimental treatment face a more symmetric information structure by observing post-purchase information. While post-purchase information is useless for subjects who search optimally, it may be used as a learning device for boundedly rational agents. In particular, if an asymmetric information flow is important, post-purchase information may provide a more symmetric feedback, thereby generating longer searches.

The results are described in Section 3. They show that subjects who face a more symmetric information flow search significantly more, although still less than the optimal level. In addition, I document the important effects of information flow on learning and on strategy adjustments. I find strong evidence for the learning mechanism described above: subjects are significantly more likely to adjust their strategies in the direction of the information they actually observe. In particular, after subjects observe that a shorter (longer) search would have done better, they are more likely to shorten (lengthen) their search in a subsequent search. This result is robust, statistically and economically significant, and is present in both experimental treatments. This is precisely the mechanism which makes post-purchase information so effective: subjects in the control group cannot observe such information and therefore can never actually observe the performance of longer searches. In contrast, subjects in the experimental treatment, who face a more symmetric information structure, will sometime observe that longer search would have done better. Thus, they are more likely to revise their search strategy towards longer searches.

As discussed in Section 4, these results have important implications regarding the efficiency of search. They suggest that once these behavioral regularities are taken into account, search rules should be revised to adjust for the truncation of information flow. These implications may be more important for search markets in which post-purchase information is less common, i.e. in which it is uncommon for participants to observe prices when they are not actively searching. For example, it implies that post-purchase, supposedly irrelevant information may be actually valuable in improving search strategies in subsequent purchases; by making mistakes more salient, economic agents are more likely to reconsider their search strategy and search more efficiently in the future.

1. Motivation and background

The literature on consumers' search and its implication on market equilibrium dates back to Stigler's (1961) seminal paper "The Economics of Information." Most of the literature deals with homogeneous good markets, in which there exists price dispersion. The general setting is that consumers have information only about the distribution of prices and they search for the lowest price. Since search is costly (fixed cost per search) and consumers cannot afford to search as much as they would like, higher prices survive in equilibrium.

In this paper I focus on the simple sequential search problem, in which an agent sequentially obtains prices from a known price distribution. The agent searches for the *highest* price at which to sell a good, and has to decide when to stop; this is the prominent setup in the literature. In such a case, we can think of the agent as a seller (or a job seeker), who searches for the highest bid (wage) for her product (services).

I first outline the theoretical predictions for the simple search model. If the agent perfectly knows the (stationary) distribution from which prices are sampled, then her optimal search employs a strategy of a reservation price. The *risk-neutral* agent continues to search as long as the expected value of an additional bid is higher than the cost of obtaining this bid, i.e. the reservation price R is the solution to

$$c = (1 - F(R))E_F(p - R \mid p \geq R) \tag{1}$$

where c is the search cost, p is the sampled price, and $F(\cdot)$ is the cumulative distribution function of prices, which is known to the agent.[5]

Throughout I do not allow the possibility of recall, i.e. whether the agent can change her mind and accept previously rejected offers. While the possibility of recall is favorable to the agent, note that in this standard setting—stationary price distribution and identical cost per search—reservation price strategy is optimal, so a rational agent would never recall. In addition, by construction, there is no recall when using any other, sub optimal, reservation price, so recall is not an issue for the present study, as will become clear later on.

Experimenters in economics started studying individual search behavior in the early 1980's. The main question they tried to answer is how individuals search, and how close they are to optimal behavior. As already mentioned, the vast majority of the experiments share a common result: subjects tend to search too little compared to the optimal rule.[6] This result is the main motivation for this paper.

There are different competing explanations for the "too little search" result. The *first*, which is often discussed in the literature, is risk aversion. Once a subject is faced with the decision whether to stop or to continue, the choice is equivalent to choosing a lottery. The higher the level of risk aversion, the less likely the subject will accept the "lottery" and continue searching. There are reasons to believe, however, that risk aversion is not enough to explain the "too little search" result. First, all experiments offer very low monetary prizes, over which one may assume that subjects are *locally* risk neutral (Rabin, 2000). In addition, in a repeated search problem (which is the case of most experiments) subjects are diversified over time and bear lower risk.[7] The risk aversion problem is further reduced once subjects are restricted to use predefined search strategies. Indeed, Sonnemans (1998) shows that risk aversion can be a possible explanation for only twenty percent of his subjects who search too little.[8]

The *second* possible explanation, which is suggested by Sonnemans (1998), is that stopping rules that give rise to too little search perform rather well in most cases. On the other hand, stopping rules that give rise to too much search perform quite poorly. Hence, subjects who search too much are more likely to revise their rules downwards, so on average subjects will search too little. In Einav (2002) I test for this explanation, and show that, indeed, it partially explains the too little search result.

The *third* explanation, which has not been tested in the literature, arises from the structure of information flow in the search problem, and the notion of observational regret as a learning device. Subjects can experience regret in this setting when they observe a previously rejected price that could have made them better off.[9] When the stopping rule is of a reservation price, then *only* prices that are lower than the reservation price can be a source for regret. A price higher than the reservation price *cannot* be a source for regret, because the only such price observed will be the transaction price, which is the last in the sequence of offers (after which search ends). Consequently, if observational regret enhances learning and affects subjects when they revise their stopping rules, then we should observe more revisions downward, so subjects will end up searching too little. It is important to note that the mechanism described is not based on *anticipated regret* that enters into the individual's utility function, but on *observational regret* that serves as a learning device that "reminds" the individual to reevaluate her strategy.[10] This explanation is the focus of the experiment described below.

2. Experimental design

The experimental design has two different experimental conditions, a control group and a treatment group. Subjects in the control group face the experimental design of Sonnemans (1998), while subjects in the treatment group face a slightly different information structure. Differences between the groups will allow me to test my central hypothesis.

The setup uses a repeated search for the highest price (bid), in which subjects are restricted to use reservation price strategies. The reservation price can be revised from period to period, but not between bids within the same period. Subjects have full information about the price distribution, and they pay a fixed amount at the beginning of each period.[11]

There are several reasons why this setting is attractive. First, the distribution is known to the subjects, so there is no room for prior beliefs that subjects might form, for which the experimenter cannot control. Complete information about the distribution also controls for heterogeneity among subjects in information-gathering techniques, and provides a unique optimal strategy that can serve as a benchmark for comparison. If information is not complete, such optimal strategy depends on the prior, which may vary across subjects.[12]

Second, the search problem is repeated, so search is observed for consecutive periods and learning can take place. As has been widely shown in the experimental literature,[13] learning is key in many experiments, and many subjects need experience in order to better understand the setup and the environment.

Third, while making the problem somewhat less realistic, the use of predefined strategies allows the experimenter to avoid dealing with dynamically inconsistent behavior by subjects, which is likely to occur in a search problem. Finally, strategies are simple, and consist only of a reservation price. This allows for a simple and convenient quantitative prediction and analysis, without losing too much generality.

The price distribution is uniform over a discrete support, which includes integers between 1 and 100 (cents), so each of these integers has a probability of one percent. The search cost is 2 cents for each new bid. The fixed cost per period is 50 cents. The optimal reservation price in such a setting is 81 cents. Sonnemans (1998) replicates the too little search result in such a setting, so it is reasonable to explore this result even with these restrictions on search strategies.[14]

Let us define the concept of (*observational*) *regret*: subjects observe regret if their profits could have been higher by stopping earlier.[15] Using this definition, Sonnemans' results show that although regret is not very frequent, it has a very important role: subjects are twice as likely to revise their reservation price after experiencing regret, and when they do, the absolute magnitude of change is much higher. Similar results are replicated in this paper and provide the motivation for a deeper look at the role of observational regret as a learning-enhancing device.

As mentioned earlier, this basic setup results in an asymmetric information structure, which leads to a one-sided regret. For that, I distinguish between two types of regret. The first, *upward regret*, is a case in which a subject observes that she could have done better by setting a higher reservation price and continuing to search. The second, *downward regret*, is the original case of regret in which a subject could have done better by setting a lower reservation price and stopping the search earlier. In the control group, by construction, subjects can only experience downward regret, because they do not observe bids after they stop searching. Therefore, the second experimental treatment allows for a more symmetric information flow by making subjects observe post-purchase bids. This allows for experiencing regret in both directions. In particular, I let subjects in the second experimental treatment observe three additional bids that they would have gotten had they continued searching (by setting higher reservation price). Other than that, subjects follow exactly the same procedure as subjects in the control group. In such a setup, a downward regret is still possible, as in the control group, but now there is also the possibility of an upward regret. This may be the case if one of the three additional bids is sufficiently high.

The three additional bids are, on average, about 60% more bids than typically observed, as the average number of bids around the optimal strategy is about five. It allows for sufficiently high probability (about 45% at the optimal strategy) of observing upward regret, without making it too tedious for the experimental subjects. Given the results described below, which suggest that these additional bids affect learning and efficiency, it may be interesting to investigate how learning changes when the number of additional bids varies. I leave this for future research.

3. Implementation and results

After a small-scale pilot run that led to a few technical changes in the experimental design, I ran five identical sessions of the experiment. The subjects were chosen (on a first-come-first-served basis) from the subject pool of the Computer Lab of Experimental Research (CLER) at the Harvard Business School. The majority (but not all) of the subjects in this pool are undergraduate students from Harvard University and Boston University.[16] Each of the runs was performed in a single session that included subjects from both treatments. All sessions were held at the CLER and lasted about 90 minutes each. The random draws from the distribution were drawn in advance, and were identical for all subjects. Subjects were not aware that they all faced identical draws.[17] The number of subjects was 35 for the control group and 40 for the experimental treatment.

In what follows I describe the experimental results. First I report the means and medians of earnings and reservation prices over time and across experimental treatments. These results give a general idea of the behavior of subjects, and are most easily compared to results of previous experiments. Next, I continue by looking at other dimensions of the experimental results, focusing on observational regret.

Using regression analysis, I obtain strong evidence in support of the hypothesis that observational regret serves as a learning-enhancing device (Tables 3 and 4). The results show that an observed *downward* regret makes subjects significantly more likely to *decrease* their reservation price, and that an observed *upward* regret makes subjects significantly more likely to *increase* their reservation price. Probabilistic regret, which is not actually observed, has no significant effect. In the end of this section I summarize my findings.

3.1. Descriptive statistics

Table 1 presents the minimum, maximum, and average earnings for each treatment. As can be seen, subjects earned on average a little less than \$30, with earnings varying between \$20 and about \$31, except for few subjects who earned very little. The earning distributions for the two groups are very similar.

At the end of each session, after subjects had received their earnings, they were asked to describe in writing the way they approached the problem. Despite the big effort to make

Table 1. Earnings.

	Control	Treatment
N	35	40
Min	11.98	7.51
Mean	26.17	26.23
Median	28.18	28.74
Max	31.29	31.02
Std. dev.	9.36	10.30
Optimal*	30.90	30.90

All figures in US Dollars.
"Optimal" refers to the earnings obtained by setting 81 at all periods.
Note that subjects were guaranteed $20 (so the few subjects who obtained less than that received $20 in the end of the experiment).

subjects understand the setup, it seems from the answers that some of them had a very weak idea of what they were supposed to do during the experiment. For comparing the means and medians of reservation prices across the two experimental conditions (and *only* for that), these subjects introduce biases and increase the variance, so I chose to drop them from the sample for this specific analysis.[18] The reason is that a random draw from the distribution has a mean of about 50, while the majority of the observed reservation prices are in the range of 70 to 85. It turned out that subjects in the experimental treatment were more likely to get confused and to think that they do not understand the instructions. This is due to the additional observed draws, which had no practical effect on the outcome (beyond their effect through learning). Therefore, comparing means with these "random" subjects in the sample is likely to bias the results downward in this treatment, when compared to the control group, and hence create problems in the interpretation of the simple comparison of means across groups. In all other statistical analyses, such "random" subjects can be treated as a "white noise," so I keep them in the sample, and they have little effect on the results.

Table 2 and Figure 1 show the average and median reservation price for each group, as it evolved over the 100 periods of the experiment. As in previous experiments, which are discussed in Section 1, it is shown that, on average, subjects search too little. For both groups in all periods the average reservation price is below the optimal reservation price of 81, making subjects stop searching too early. The results confirm the hypothesis that subjects who may experience symmetric regret by observing higher prices after they have stopped searching end up searching longer and setting higher reservation prices, closer to the optimal one. Once the "random" subjects are dropped, there is much less noise in the averages, so I am able to obtain statistically significant differences between the two groups.

There are a few additional interesting results that should be mentioned. First, even subjects in the experimental treatment, who search more, still search too little by setting average reservation prices that are significantly lower than the optimal one, 81. This can be explained

Table 2. Averages of reservation prices.

	Control	Treatment
	All subjects	
First period	63.3 (2.76)	63.1 (3.18)
All periods	68.7 (2.33)	69.4 (2.26)
Last 50 periods	70.8 (2.41)	70.8 (2.32)
Last 25 periods	71.3 (2.37)	70.8 (2.43)
Last 10 periods	72.0 (2.46)	71.3 (2.43)
Last period	72.7 (3.02)	70.8 (2.93)
No. of subjects	35	40
	All "Non-Random" subjects	
First period	66.5 (2.17)	66.7 (3.38)
All periods	72.6 (1.83)	75.2 (1.45)*
Last 50 periods	75.1 (1.48)	77.1 (1.19)*
Last 25 periods	75.1 (1.57)	77.6 (1.12)**
Last 10 periods	75.2 (1.66)	77.9 (1.23)**
Last period	76.0 (1.78)	79.1 (1.52)**
No. of subjects	30	30

***Significantly different from the control group with p-value $> 5\%$ for a one-sided t-test.
**Significantly different from the control group with p-value $> 10\%$ for a one-sided t-test.
*Significantly different from the control group with p-value $> 20\%$ for a one-sided t-test.
Standard deviations in parentheses.
All figures in all tables are significantly (at 5%) different from 81, the optimal reservation price.
The criteria for defining a subject as "random" is given and motivated in endnote 18. This table and the corresponding Figure 1 are the *only* place in the paper where such selection is applied. All other analysis of the paper is based on the full sample.

by the other factors that affect search strategy (see Section 1), such as the cost asymmetry (see also Einav, 2002). Second, it is quite clear that the results are not driven by initial conditions. The first period is completely identical for both group,[19] and subjects indeed start from almost the same average reservation price, as shown in the first row of Table 2. Finally, it is important to note that even after one hundred repetitions subjects continue to adjust their reservation prices, albeit they do so less frequently.

3.2. *Observational regret as a learning-enhancing device*

In this section I focus on the effects of observational regret, which is the main focus of the paper. In analyzing observational regret, it is useful to make comparisons with an artificial

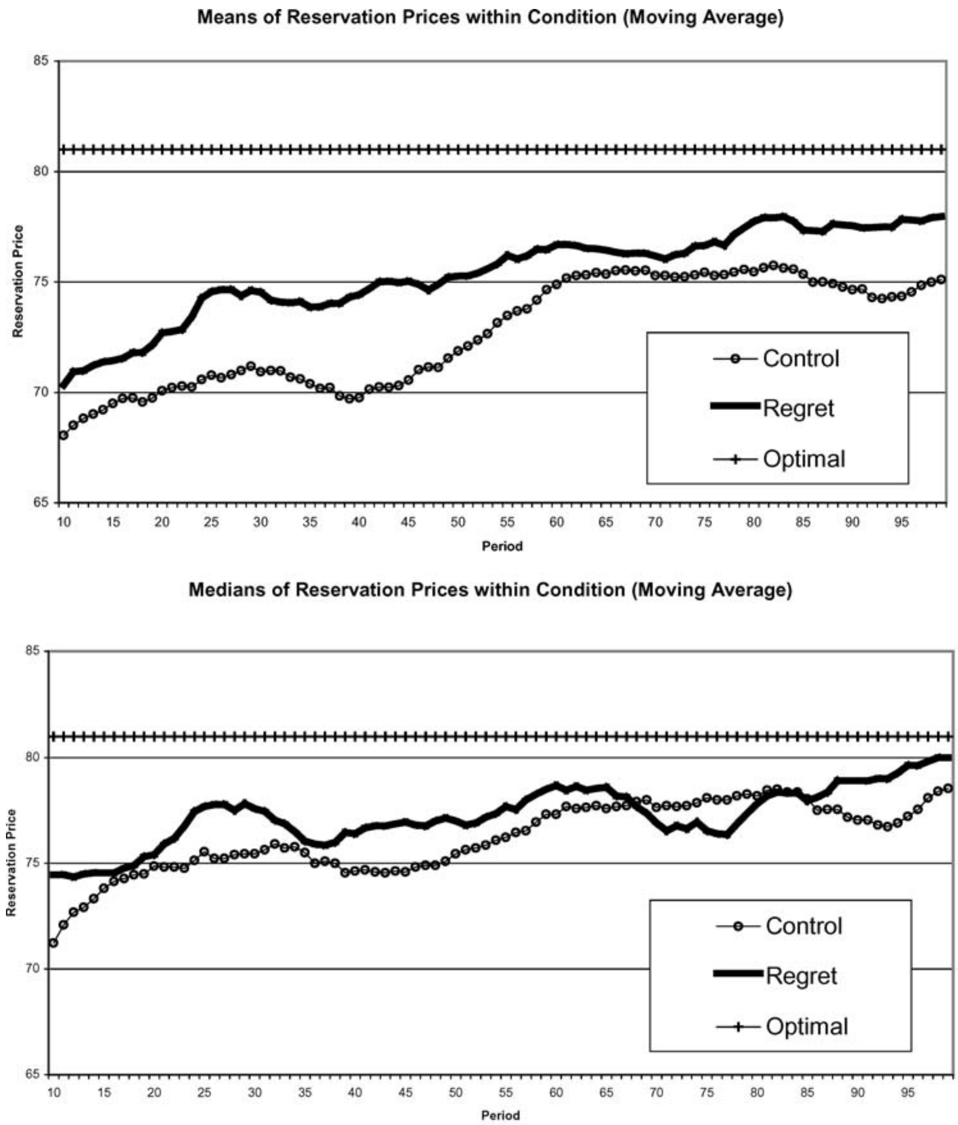

- Note that both figures are presented in moving averages in order to make the plots smoother. Each data point in the graphs is the mean of the preceding 10 averages (top figure) and medians (bottom figure).
- These figures are based only on the "Non-Random" subjects. The criteria for defining a subject as "random" is given and motivated in endnote 18. This figure and the corresponding Table 2 are the *only* place in the paper where such selection is applied. All other analysis of the paper is based on the full sample.

Figure 1. Averages and medians over time.

experimental group, the "Pseudo control group." This comparison looks at the behavior of subjects in the control group *as if* they were observing the three additional prices just as the subjects in the experimental treatment.[20] This allows me to better control for the results associated with the experimental treatment. That is, I control for the possibility that observational regret itself has no impact, but is simply correlated with some other relevant variables. For example, it is natural to assume that subjects will be more likely to adjust their reservation price downwards (upwards) after setting a relatively high (low) reservation price. Downward (upward) regret is also more likely to be observed after setting a high (low) reservation price. Hence, an adjustment upward may be positively correlated with upward regret only because of statistical reasons, without any causal effect of observational regret. By using the pseudo control group, and comparing it to the experimental treatment, I am able to obtain the perfect control, so that the only difference among subjects in the two groups is the *actual observation* of regret.

While simple empirical probabilities may provide an idea of the effect of observational regret on subjects' behavior, the cleanest test can be done using a more rigorous statistical analysis. In what follows, I use regression and multinomial logit analysis to test the idea of directional learning through observational regret.[21]

Tables 3 and 4 analyze the effects of observed regret on the subsequent adjustment of reservation prices. In particular, for each subject i in period $t(t > 1)$ I regress the change in reservation price from period $t-1$ to period t (i.e. $rp_t - rp_{t-1}$) as a function of observed regret in period $t-1$ and other controls. The control variables, which include the reservation price at period $t-1$, the number of bids in period $t-1$, and the dollar earnings (in cents) in period $t-1$ are all trying to make sure that the estimated effect of observed regret is indeed causal and driven by learning, and it is not due to some other statistical regularity, such as the one mentioned above. All regressions are run separately for each group, and cluster by subject, to correct for the likely correlation in each individual's behavior over time. The most convincing result, in my view, is the comparison between the effect of upward regret in the pseudo control group and the treatment group. The significant difference in its effect across the two groups can only be attributed to the actual observation of upward regret.

Table 3 reports a linear regression in which the dependent variable is a logistic transformation of the change in reservation price (so that the results are not driven by few extreme changes, from 1 to 100 and vice versa). In all groups the effect of a downward regret is negative and statistically significant: *ceteris paribus*, when a subject observes a downward regret, she is more likely to subsequently choose a lower reservation price. The effect of an upward regret works at the other direction: everything else equal, subjects in the treatment group who observe an upward regret are more likely to subsequently choose a higher reservation price. While this effect loses some of its statistical significance once more control variables are added (column 6; its *p-value* is about 12%), it is in sharp contrast to the absence of this effect (in terms of magnitude and significance) in the pseudo control regressions. Subjects in the control group are not affected by upward regret. This suggests that it is the actual observation of upward regret, rather than some hypothetical regret, that makes subjects adjust their strategies. Often, this adjustment would be in the right direction, towards the optimal reservation price, as most subjects use too low reservation prices. This is the primary factor for the efficiency differences between the groups shown in Figure 1.

Table 3. Linear regressions.

	Control		Pseudo control		Treatment	
Dependent variable: logistic transformation, $\exp(x)/(1+\exp(x))$, of the change in reservation price from period $t-1$ to t.						
Downward regret$_{t-1}$	-0.176^{**}	-0.146^{**}	-0.188^{**}	-0.158^{**}	-0.172^{**}	-0.092^{*}
	(−4.75)	(−3.81)	(−5.08)	(−4.18)	(−5.71)	(−2.55)
Upward regret$_{t-1}$	–	–	0.006	−0.0016	0.037^{*}	0.026
			(0.35)	(−0.10)	(2.46)	(1.59)
Constant	0.767^{**}	0.795^{**}	0.760^{**}	0.795^{**}	0.806^{**}	0.851^{**}
	(11.87)	(9.54)	(12.14)	(9.64)	(10.92)	(9.49)
Reservation price$_{t-1}$	-0.0036^{**}	-0.0031^{**}	-0.0035^{**}	-0.0031^{**}	-0.0043^{**}	-0.0036^{**}
	(−4.18)	(−3.84)	(−4.19)	(−3.84)	(−4.26)	(−3.71)
No. of bids$_{t-1}$		-0.0033^{*}		-0.0032^{*}		-0.0073^{**}
		(−2.41)		(−2.52)		(−4.15)
Earnings$_{t-1}$		−0.0006		−0.0006		−0.0008
		(−0.83)		(−0.84)		(−1.30)
No. of subjects	35	35	35	35	40	40
No. of observations	3,465	3,465	3,465	3,465	3,960	3,960
R-squared	0.054	0.056	0.055	0.057	0.089	0.095

t-stats below coefficients, clustered by subject.
**, * implies significance at the 1 and 5% confidence level, respectively.
The regressions are run for each experimental treatment separately. The control group and the pseudo control group use the same data. The only difference is that in the pseudo control an upward regret regressor is computed, *as if* subjects could observe post-purchase information. The coefficient on Downward regret suggests the existence of observational learning. The most important observation, however, is the difference in the magnitude and significance in the coefficient on Upward regret in the pseudo control group and the treatment group regressions. Its significance only in the treatment group regression implies that it is only the *actual* experience of regret that affects behavior. The rest of the regressors are controls. They all obtain the predicted signs. The regressions include all subjects.

To gain some intuition for the magnitude of the effect, a subject who would otherwise not change her reservation price will decrease (increase) it by about 0.7 (0.15) after experiencing downward (upward) regret.[22] This magnitude is quite large considering the fact that, on average, each subject increases her reservation price by about 0.08 per period. Loosely speaking, a downward regret offsets about nine periods without regret, while an upward regret is equivalent to three periods without regret.

The control variables all have the predicted signs. Higher reservation prices make it more likely to adjust downwards, and higher number of bids have similar effect. Earnings do not seem to have a significant effect on the strategy adjustment. These coefficients on the control variables illustrate why it is important to use them. Downward regret is more likely after a high reservation price and high number of bids. Thus, if we did not control for these variables in the regression, we would obtain stronger estimated effect of downward regret, which would only be driven by statistical properties of regret. After using these variables

Table 4. Multinomial logit regressions.

Dependent variable	Control		Pseudo control		Treatment	
	Downward change	Upward change	Downward change	Upward change	Downward change	Upward change
Downward regret$_{t-1}$	4.34**	1.46	4.98**	1.56	4.53**	1.28
	(3.98)	(1.05)	(4.27)	(1.18)	(4.71)	(0.60)
Upward regret$_{t-1}$	–	–	1.10	1.15	1.02	1.42**
			(1.09)	(1.41)	(0.22)	(3.64)
Reservation price$_{t-1}$	0.993	0.972*	0.994	0.973	0.987	0.962**
	(−0.46)	(−1.97)	(−0.42)	(−1.87)	(−0.99)	(−3.11)
No. of subjects (observations)	35 (3,465)		35 (3,465)		40 (3,960)	
Log-likelihood	−3,645.7		−3,642.2		−4,037.1	
Pseudo R-squared	0.032		0.033		0.058	
Downward regret$_{t-1}$	2.08*	1.12	2.38*	1.20	1.85*	1.58
	(2.00)	(0.36)	(2.23)	(0.56)	(2.00)	(1.20)
Upward regret$_{t-1}$	–	–	1.05	1.06	0.95	1.24*
			(0.44)	(0.48)	(−0.46)	(2.17)
Reservation price$_{t-1}$	0.992	0.974	0.992	0.974	0.985	0.968**
	(−0.53)	(−1.80)	(−0.53)	(−1.81)	(−1.19)	(−2.73)
No. of bids$_{t-1}$	1.052**	1.005	1.051**	1.006	1.054**	0.948**
	(2.87)	(0.25)	(2.75)	(0.32)	(4.49)	(−2.96)
Earnings$_{t-1}$	0.993	0.991	0.994	0.992	0.991*	0.988*
	(−1.45)	(−1.59)	(−1.02)	(−1.11)	(−2.08)	(−2.47)
No. of subjects (observations)	35 (3,465)		35 (3,465)		40 (3,960)	
Log-likelihood	−3,631.3		−3,630.1		−4,006.8	
Pseudo R-squared	0.036		0.036		0.066	

Estimates are reported as odds ratios. z-stats are reported in parentheses, clustered by subject.
**, * implies significance at the 1% and 5% confidence level, respectively.
In all regressions, the base category is no change in the reservation price.
The regressions reported below are similar to those reported in Table 3. Rather than using the magnitude of the change in reservation price, these regressions replicate the same pattern of Table 3, when the dependent variable is the discrete choice of whether to change the reservation price, and if so in which direction. The top panel runs the regressions with no controls, while the bottom panel adds several controls. The results are qualitatively similar. The specification is of a multinomial logit, where the reported figures are of odds ratios.
As in Table 3, the Downward regret works as predicted: it increases the probability of changing the reservation price downwards, but has no significant effect on a change upwards. As before, the most interesting contrast is between the effect of Upward regret in the pseudo control group, where it is not significant, vis-à-vis its significant effect in the treatment group. Again, this implies that it is only the actual experience of regret that affects behavior.

as controls, it seems much more likely that the remaining effect of regret is indeed due to learning.

Table 4 repeats the exercise, using multinomial logit regression. Rather than using the dependent variable as a continuous measure, Table 4 only analyzes the direction of strategy adjustment, to focus on the directional learning idea. Thus, the dependent variable is discrete: upward adjustment of the reservation price ($rp_t - rp_{t-1} > 0$), no adjustment ($rp_t - rp_{t-1} = 0$), or a downward adjustment ($rp_t - rp_{t-1} < 0$). The results are similar to those presented in Table 3. In all groups, observing a downward regret makes it more likely that subjects adjust their reservation price downwards, while observing an upward regret makes it more likely that subjects in the treatment group adjust their reservation prices upwards. In sharp contrast, (hypothetical) upward regret has no significant effect in the pseudo control group, suggesting that it is only the actual observation of regret that affects behavior.

Again, the magnitude of the effect is quite substantial. On average, each subject adjusts her reservation price downwards about 28 percent of the times. The odds ratio of 4.34 in the top panel of Table 4 implies that after experiencing a downward regret, this probability is much higher, and is about 64 percent. Similarly, the odds ratio of 1.42 for upward regret (the last column of Table 4) implies that the probability of adjusting the reservation price upwards become about 40 percent, compared to an average probability of 32 percent.

3.3. Summary

The results replicate well previous results that found that subjects search too little. In this setup it is shown that subjects in both groups set on average a reservation price which is significantly lower than the optimal reservation price of 81. In addition, I show that subjects who are able to obtain more symmetric information structure tend to search better, closer to optimal, albeit still too little. Finally, the results document quite strongly the effect of *actual* observation of regret, in both directions, upward as well as downward. The analysis suggests that observing regret can be thought of as a feedback to the subjects, which enhance their learning. Once the feedback is objective, symmetric in both directions, subjects react accordingly and better converge towards the optimal reservation price.

4. Conclusions and further research

Previous experimental search papers have consistently found that subjects in experiments search too little. In this paper I test for a possible explanation for this robust result, and find evidence in its support. Subjects, who can evaluate their strategies more symmetrically by observing post-purchase information, search more intensively.

In addition, the results show the importance of observational learning in the form of observational regret as a learning-enhancing device. Using the individual search problem it is shown that actual observation of regret leads to a directed adjustment of played strategies, while probabilistic regret has no statistically significant effect. It is evident that subjects who are able to actually observe regret end up searching in a more efficient way by setting reservation prices that are closer to the optimal one.

We can try to extend these results to strategic interactions. For example, one could look at a sequential two-player game, using both the action method and the strategy method.[23] In the strategy method subjects are informed about the full strategy of their opponents, so they can calculate their unrealized payoffs. In this sense, the strategy method is similar to the experimental treatment and allows for observational regret, while the action method is similar to the control group, in which subjects may only speculate what could have been their earnings had they chosen a different action. The results from this paper suggest that learning and convergence to equilibrium would be faster when the strategy method is used.

The observation of post-purchase information is not only a "mechanical device" to test the explanation for too little search. It may also be a closer description of certain markets in the real world. It is quite often that a consumer buys a certain good and later observes the same good being sold at a lower price. This new information can be obtained from a friend who bought the same good, from advertisement, or from visiting the same retail store for another reason. Moreover, it seems plausible that consumers who bought a certain product recently are more likely to internalize new information about it. Consequently, the results obtained in this paper suggest an important empirical implication: search is going to be closer to optimal in markets in which consumers are more likely to obtain post-purchase information. This prediction may be testable. This can be done by comparing search behavior in markets for similar products that differ in their information structure.[24] In markets where search is repeated but post-purchase information is not frequently available, a more efficient search can be achieved by designing search strategies that will make information flow in a more symmetric manner, making post-purchase information more readily available.

Appendix A: Experiment instructions

You are participating in an experiment about *individual* decision making. Hence, nothing that you do can affect the other participants, and nothing they do will affect you. In addition, the setup is not completely identical for all of you (and a lot of it is random), so trying to learn from others' choices does not make much sense.

The experiment consists of 100 identical periods. At the start of each period you will get an artificial object for 50 cents. You will have to sell this object to the computer. Each period, before the computer makes bids for the object, you will have to announce the minimum price at which you are willing to sell the object. Then, the computer will make bids for the object, until there is a bid greater than or equal to the minimum price you have announced. *The selling price will be this last computer bid.*

You can think of this setup as a search for a new job. In this analogy, each computer bid is a wage offer for the same job, and your search criterion defines a minimum wage at which you are willing to take the job.

There are three things that are very important to know:

1. Every bid costs 2 tokens. Using the wage analogy, these costs can stand for the time you spend looking for a job. Hence, your costs in a period are two tokens times the number of bids in that period, plus the 50 tokens that you had to pay to get the object.

2. Your earnings in a period are the price at which you make the sale (the first bid that is greater than or equal to the minimum price announced at the beginning of that period), minus the costs in the same period (i.e. Earning = Sale Price − 2 × Number of Bids − 50). Note that it is possible to have negative earnings in a period. If that happens, the loss will be deducted from the earnings in other periods.
3. All bids are integers between 1 to 100 tokens (1 and 100 included). Each bid is *independent* of previous bids and of the minimum price you announced. The bids have already been generated randomly, prior to the experiment, so any decision that you make during the experiment is not going to affect the bids you will get later on. In page 4 of the instructions you can take a look at a histogram of 100,000 bids that were generated by the same process that generated the bids you will get during the experiment. The height of each bar is the probability of getting the corresponding bid.

At the beginning of each period you will start by having a sample of computer bids (arranged in a circle shape) which were generated by the same process as the real ones. These sample bids are *independent* of the real ones, and are provided only in order to give you a better understanding of the distribution of bids. There are no tricks—if you feel that you understand the distribution from which the bids are generated, it is totally sensible to ignore these sample bids. The sample bids are provided only in order to give you a better "feel" of the distribution of bids. After you have looked at the sample bids, you can choose your minimum selling price for that period and bidding will take place.

At the end of each period, *some of you* will get the opportunity to observe the next three computer bids (you get it the same way you received the previous bids). These additional bids cannot be used for sale. They are provided in order to let you see how more bids would have looked like in case the minimum selling price you had set would not have made you stop where you stopped, so bidding would have continued.

The above will go on for 100 periods. Your total earnings will be the sum of the earnings in each period.

Note that by choosing your minimum selling price you choose both the expected price you will receive and the expected number of bids you will pay for. Therefore, by choosing your minimum selling price thoughtfully, you can increase your earnings.

Appendix B: Computer program description

You begin the experiment by having one of the instructors activating the program. This will make the welcome window disappear, and you will be able to start the experiment by clicking on the "Next Period" button, which appears in the right bottom corner of the screen (please raise your hand if you do not see this button).

After clicking the "Next Period" button, the minimum price window (shown below) will pop up. The circle that appears in the lower part of the window is the sample bids that are discussed in the instructions. You can then choose the minimum price for that period from the ten menus in the top of the window. After you made your choice, click on the "OK" button, which appears in the top right corner of the window.

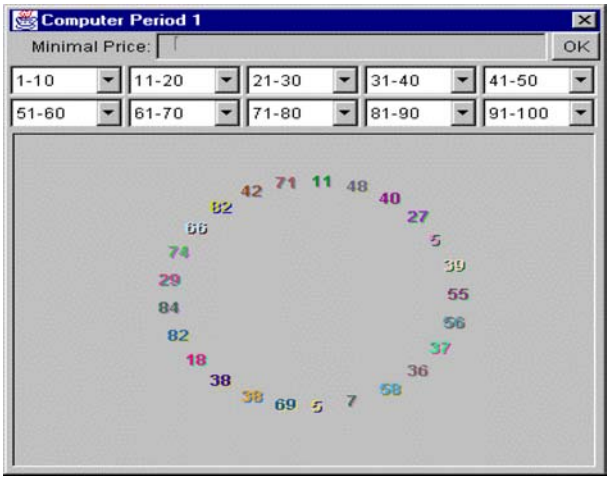

Now bidding will take place, as can be seen in the screen below. You can get new bids by clicking the "Next Bid" button (left bottom corner) for each new bid. When bids take place, they appear on the left side of your screen. There are three columns—the bid number, the bid, and the amount you would have earned if the bid was the last one for the period (calculated by: bid $-$ 2 $\times$ bid number $-$ 50). On the right side of your screen, your performance in past periods is recorded—the period, the minimum price you announced for that period, the number of bids you got, the sale price and the total earnings you had for the period. At the top of the screen, you can see also your subject number, what period it is, the minimum price you announced for the period, and your total and average earnings so far. Once a period ends, click the "Next Period" to go on.

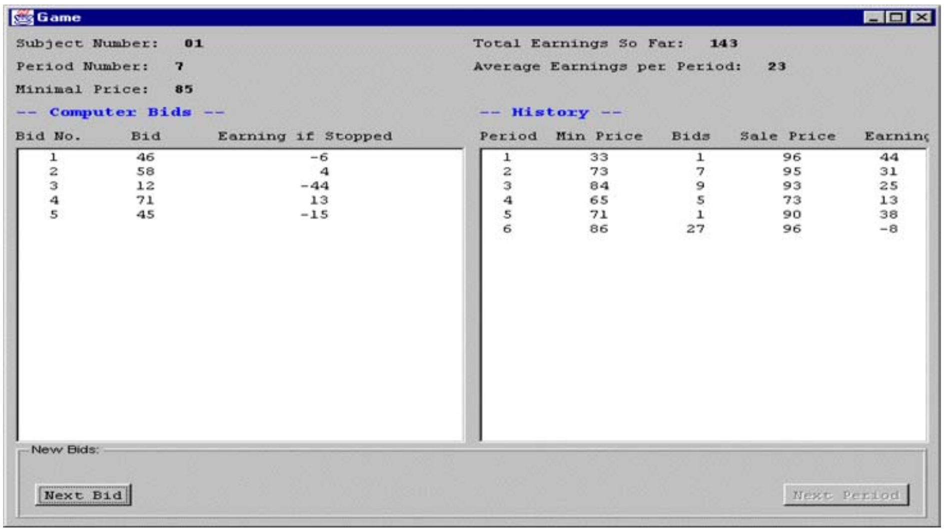

After you are done, the last screen will pop up. Fill in your name and social security number and click the "Print and Quit" button. We will collect the printouts from the printer, prepare your payment and call you to pick it up. The receipt is the only place in which your name and social security number are recorded. Other than that, only the random number assigned to you is recorded.

If there are no further questions, you are ready to start. The instructors will approach your desk and activate the program.

Enjoy!

Acknowledgments

This paper is a revised version of chapter 3 of my 2002 Harvard University dissertation. I am particularly grateful to Al Roth for continuous guidance, helpful discussions, and financial support of the pilot studies. I thank Drew Fudenberg, David Laibson, an anonymous referee, and the editor for helpful comments and suggestions. Estelle Cantilon, Guillermo Caruana, Yair Eilat, Muriel Niederle, Robert Slonim, Manuel Trajtenberg, Leeat Yariv, Elena Zoido, and seminar participants at Harvard University and at the 1999 *Econometrics in Tel-Aviv* provided suggestions and comments on earlier drafts. I thank Joep Sonnemans for sharing with me his experimental data. Financial support by the Russell Sage Foundation (RSF Project #98-00-07) is gratefully acknowledged. All remaining errors are my own.

Notes

1. See also the introduction to Erev and Roth (1998) and the references therein.
2. See Stigler (1961) and Kohn and Shavell (1974).
3. See, for example, Schotter and Brounstein (1981), Hey (1982, 1987), Kogut (1990), Cox and Oaxaca (1996), Moon and Martin (1996), Butler and Loomes (1997), and Sonnemans (1998).
4. The control group is identical to the design of Sonnemans (1998).
5. Kohn and Shavell (1974) and Telser (1973) discuss the implications of extensions to this setup.
6. See the references listed in endnote 3.
7. More formally, suppose a search strategy yields expected payoff of e with payoff variance of v. Now, instead of running one such search problem, we run N i.i.d identical problems, such that the sum of the expected payoffs remains the same (so the "size" of each problem is N^{-1} times the original problem). Thus, it is easy to see that the variance of the same strategy in the repeated problem is $N(vN^{-2}) = vN^{-1}$.
8. Although not in the scope of this paper, a possible way to directly test the power of the risk-aversion explanation is by using binary lottery techniques. These techniques use offers as probabilities of winning a binary lottery, instead of direct monetary payments. See Roth and Malouf (1979), Berg et al. (1986), and Kagel and Roth (1995) for more details. As mentioned before, Rabin (2000) suggests that this may not be necessary.
9. Sonnemans (1998) mentions this explanation, and documents that the behavior of subjects after experiencing regret is different from their behavior in other periods. He does not, however, formally analyzes this hypothesis, as it is not the focus of his paper.
10. Such ideas are consistent with directional learning theory (Selten and Buchta, 1994), and, to a lesser extent, with "regret matching" (Hart and Mas-Colell, 2000).
11. The fixed cost per period is *not* a choice and has no effect on the optimal strategy. It is useful because it increases the salience of net payoffs and directs subjects towards setting higher reservation prices, speeding

up learning and convergence. Pilot studies without fixed cost have indeed shown that the fixed cost is quite useful for this purpose.

12. Even with full information, one may wonder how well subjects understand the information they are given. In the experiment described below I try to help subjects understand what the distribution is by showing them a random sample from it.
13. See, for example, Roth and Erev (1995) and the references they cite.
14. Sonnemans (1998) had 19 subjects who played for 45 periods. They started (period 1) with an average reservation price of 67.5 and ended (period 45) with an average reservation price of 77.9, which is still below the optimal reservation price of 81.
15. As emphasized in the previous section, it is important to distinguish between the concept of observational regret, which is defined above, and the more common use of (anticipated) regret, which enters into the utility function. Throughout the paper, regret plays a role in enhancing learning, but not through any direct effect on utility.
16. The complete instructions can be found in Appendix A. Appendix B provides a description of the computer program, which was also given to the subjects before the beginning of the experiment. These, and a short pre-experiment exercise, were read aloud to subjects prior to the experiment.
17. The identical draws eliminate some of the noise that can arise from the random differences in the draws among subjects. It is also true, however, that identical draws may cause some bias in the results (if, for example, the very first rounds are crucial for learning). With a relatively small sample, the trade off is best resolved by having identical draws. With large enough sample, however, it would be definitely better to allow different draws for each individual.
18. The criterion for dropping subjects was very conservative. I drop any subject who chose reservation price lower than 50, for at least 5 times during the last 40 (out of 100) periods. In general, a reservation price lower than 50 seems somewhat unreasonable given the fixed cost of 50. I look only at the last 40 periods because it makes sense that subjects may not realize immediately that less than 50 is a bad strategy, but only later on. It goes without saying that the criterion was chosen without looking at its affect on the results. This criterion led to dropping as many as 15 subjects (5 (out of 35) from the control group and 10 (out of 40) from the treatment group). As could be predicted, more subjects got confused (and were dropped) in the somewhat "more confusing" condition, namely the treatment group, in which they get to observe post-purchase bids, which have no effect on the outcome.
19. At the very first period subjects from either group do not know which condition they are in. Only in the end of the first period, once they observe or do not observe post-purchase prices, they realize to which one of the groups they were randomly assigned.
20. These three additional prices are known to us, because the sequences of bids were previously generated. In fact, for our purposes we could also generate new i.i.d draws from the same distribution, and think of them as the three potential future bids.
21. Simpler descriptive statistics also lead to similar conclusions, and are reported in Einav (2002). The qualitative results are not sensitive to the particular specification or to whether the "random" subjects are included or not. Tables 3 and 4 provide results with *all* subjects included.
22. This is calculated by inverting the logistic transformation. For an estimated coefficient of a, its effect on the dependent variable, which would otherwise be zero, is given by $\log(0.5 + a) - \log(0.5 - a)$.
23. Take the ultimatum game for example, in the action method the responder would only accept or reject the offer he actually received, while in the strategy method the responder would have to specify a fully contingent strategy of whether he would have accepted or rejected any potential offer.
24. For example, we can think of two markets for the same product, a regular one and an online one. Alternatively, one can think of two U.S. states that differ in the law regarding advertisement of tobacco products. Post-purchase information is likely to be better in states that allow for such advertisements. It must be noted, however, that a clean test of the prediction would require the two markets to differ in the post-purchase information, but to be similar in the pre-purchase information, as well as in the mix of consumers who participate in them. Such examples may be hard to find in real world, which is one of the reasons that make experiments so useful.

References

Berg, J. E., L. A. Daley, J. W. Dickhaut, and J. R. O'Brien. (1986). "Controlling Preferences for Lotteries on Units of Experimental Exchange," *Quarterly Journal of Economics* 101, 281–306.

Butler, D. and G. Loomes. (1997). "Quasi-Rational Search under Incomplete Information: Some Evidence from Experiments," *The Manchester School* 65, 127–144.

Cox, J. C. and R. L. Oaxaca. (1996). "Testing Job Search Models: The Laboratory Approach," *Research in Labor Economics* 15, 171–207.

Einav, L. (2002). "Essays in Industrial Organization," Ph.D. Dissertation, Harvard University, Chapter 3.

Erev, I. and A. E. Roth. (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review* 88, 848–881.

Goeree, J. K., C. A. Holt, and T. R. Palfrey. (2002). "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory* 104, 247–272.

Hart, S. and A. Mas-Colell. (2000). "A Simple Adaptive Procedure Leading to Correlated Equilibrium," *Econometrica* 68, 1127–1150.

Hey, J. D. (1982). "Search for Rules for Search," *Journal of Economic Behavior and Organization* 3, 65–81.

Hey, J. D. (1987). "Still Searching," *Journal of Economic Behavior and Organization* 8, 137–144.

Kagel, J. H. and A. E. Roth. (1995). *The Handbook of Experimental Economics*. Princeton University Press.

Kogut, C. A. (1990). "Consumer Search Behavior and Sunk Costs," *Journal of Economic Behavior and Organization* 14, 381–392.

Kohn, M. G. and S. M. Shavell (1974). "The Theory of Search," *Journal of Economic Theory* 9, 93–123.

Moon, P. and A. Martin. (1996). "The Search for Consistency in Economic Search," *Journal of Economic Behavior and Organization* 29, 311–321.

Rabin, M. (2000). "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica* 68, 1281–1292.

Roth, A. E. and I. Erev. (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior* 8, 164–212.

Roth, A. E. and M. W. Malouf. (1979). "Game-Theoretic Models and the Role of Information in Bargaining," *Psychological Review* 86, 574–594.

Schotter, A. and Y. D. Brounstein. (1981). "Economic Search: An Experimental Study," *Economic Inquiry* 19, 1–25.

Selten, R. and J. Buchta. (1994). "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Function," Discussion Paper No. B-270, University of Bonn.

Sonnemans, J. (1998). "Strategies of Search," *Journal of Economic Behavior and Organization* 35, 309–332.

Stigler, G. J. (1961). "The Economics of Information," *Journal of Political Economy* 69, 213–225.

Telser, L. G. (1973). "Searching for the Lowest Price," *American Economic Review* 63, 40–49.