Informational Asymmetries and Observational Learning in Search*

Liran Einav
Harvard University
einav@fas.harvard.edu

First version: May 8, 2000
This version: March 6, 2002

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, this paper highlights 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 asymmetric flow of information leads agents to search too little. I find strong evidence in its favor. I conclude that in the presence of learning, the provision of a more symmetric information structure will make search more efficient.

KEYWORDS: Experiment, Search, Directional Learning, Observational Learning, Regret
JEL classification: C91, D83

*I am thankful to Al Roth for continuous guidance, helpful discussions, and financial support of the pilot studies. I thank Drew Fudenberg and David Laibson 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 "Econometrics in Tel-Aviv" provided suggestions and comments on earlier drafts. I am grateful to 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.

1 Introduction

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.

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. 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. It is suggested 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 (e.g. sellers) 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 subjects 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. There is one common phenomenon that stands out from all experimental papers about search behavior–subjects in experiments search too little, when compared to the optimal search strategy.[1] The hypothesis discussed above suggests that the truncated information structure of the search problem may be the driving force behind this robust result. Another possible explanation for the "too little search" result is related to the typical shape of the expected profit curve[2] in the search context–with typical price distributions, over-searching is more costly than under-searching, thus making agents more likely to search less when they over-search than to search more when they under-search.

I design an experiment that looks at these two possible explanations for the "too little search" result. The

[1]See, for example, Kogut (1990), Moon and Martin (1996), Butler and Loomes (1997), and Sonnemans (1998). These papers are surveyed in more detail in the next section.

[2]Sonnemans (1998) calls this curve the "Efficiency Curve".

control group is the standard search environment, in which subjects face an asymmetric expected profit curve, and do not get any further feedback once they stop searching. In the experimental conditions, one group of subjects faces a symmetric expected profit curve, while another group faces a non-truncated information structure. I find that both experimental conditions partially ameliorate the too little search result. In addition, I document the important effects of information flow on learning and on strategy adjustments. Once subjects are able to actually observe strategies that could have done better, they are more likely to adjust their strategies accordingly. This result is robust and statistically significant, and is present in both groups, the group that faces truncated information flow, as well as the group that observes the full distribution of prices.

These results have important implications about the efficiency of search. It may suggest that once these behavioral forces 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.

The paper continues as follows. Section 2 reviews the experimental search literature. Section 3 describes the experimental design. Section 4 describes the implementation and the results, and Section 5 concludes.

2 Background

2.1 Existing Theory

It is hard to imagine a totally homogeneous good market. Usually goods differ in their characteristics, and it is very rare to find a single price for the same good. Because of this diversity, in many markets consumers are not fully informed about the goods available in the marketplace, their prices, and their characteristics.

One way in which consumers can learn about the different products is by searching. The literature on consumers' search and its implication on market equilibrium date 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 a consumer sequentially obtains prices from a known price distribution. The consumer searches for the lowest (highest) price at which to purchase (sell) a good, and has to decide when to stop. This problem can be easily extended to other environments.[3] To simplify the analysis, from now on I consider the case in which the consumer searches for the *highest* price; this is the prominent setup in the literature. In such a case, we can think of the consumer as a seller, who searches for the highest bid on his product.[4]

[3] For example, in the labor market workers search for jobs, looking for the job that pays the highest wage, so search is on the supply side and not on the demand side, but everything else remains the same. In differentiated goods markets, we can assume that consumers have a utility function defined over the different goods available in the market, and their prices, so consumers search for the product that gives them the highest utility.

[4] Although analytically equivalent, one may argue that searching for the highest price is behaviorally different from searching for the lowest price. For example, having prices non-negative means that searching for the lowest price has some bound, while searching for highest price may not. A simple experiment can be easily designed to address such a critique. In this paper I follow the experimental literature and use a search for the highest price, although the more intuitive example that comes to mind is of a consumer, who searches for the lowest price.

Before looking at previous experimental results, I outline the theoretical predictions for the simple search model. If the consumer knows perfectly the (stationary) distribution from which prices are sampled, then her optimal search strategy is a strategy of a reservation price. The risk-neutral consumer continues to search as long as the expected value of an additional bid is higher than the cost of obtaining this bid. Hence, the reservation price R is the solution to:

$$c = (1 - F(R))E_p(p - R|p \geq R) \tag{1}$$

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 consumer. Note that if the consumer has only partial information about the price distribution,[5] or no information at all, then the optimal strategy depends on the individual prior distribution.[6] If consumers have different priors, then their optimal strategies are not the same.

Another important issue in search problems is the possibility of recall, i.e. whether we allow the consumer to change his mind and accept previously rejected offers. In many markets we observe both: in a job market or real estate market recall is generally not feasible,[7] while in food or consumer durable goods markets it usually is. Clearly, the possibility of recall is favorable to the consumer. Note, however, that in the standard setting–stationary price distribution and identical cost per search–reservation price strategy is optimal, so a consumer who behaves optimally 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 clearer later on.

2.2 Previous Experiments

Experimenters in economics started studying individual search behavior in the early 1980's (there are also earlier search experiments designed by psychologists). The main question they tried to answer is how individuals search, and how close they are to optimal behavior. *All* the experiments mentioned below share a common result–subjects tend to search too little compared to the optimal rule. This result is the main motivation for this paper.

One of the first experiments was published by Hey (1982). He found that individuals use several rules of thumb, which take into account several variables: the last price, the highest price so far, the last 2-3 prices, the number of searches, and the search cost. Although none of the subjects used the optimal search strategy, most of these rules of thumb performed fairly well. Other experiments show other aspects of non-optimal behavior. It was shown that many subjects exercise their option to recall once it is possible (Hey, 1987, and Kogut, 1990), and that many subjects focus on total earnings rather than on marginal earnings when they decide whether to stop or to continue searching (Kogut, 1990, Butler and Loomes, 1997, and Sonnemans, 1998). Butler and Loomes (1997) suggest an aspiration-satisfying model that fits their data rather well. Moon and Martin (1996) show that most subjects behave consistently, even though not necessarily optimally.

Sonnemans (1998) restricts subjects to use a predefined stopping rule for their search. In one experiment, the stopping rule can use many different variables, and in the other subjects are restricted to use the (optimal) reservation price strategy, but are free to choose their reservation price. In both experiments the same search

[5]Partial information can appear in different forms: the shape of the distribution, without all the parameters, the first few moments of the distribution, but not all, or a sample from the distribution. Whatever it is, the main idea remains the same.

[6]Kohn and Shavell (1974) and Telser (1973) discuss this, as well as other related issues, in more detail.

[7]One may also think of the case of gas stations spread along the highway. No U-turns allowed.

problem is repeated for many periods, and subjects are allowed to update their stopping rules from period to period.[8]

2.3 Explaining the "Too Little Search" Result

There are different competing explanations for the "too little search" result. The first, which is perhaps the first that comes to mind, 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. However, there are reasons to believe 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. In addition, in a repeated search problem (which is the case of most experiments) subjects are diversified over time and bear lower risk.[9] 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 only for 20 percent of his subjects who search too little (the other 80 percent use strategies that are strictly dominated when analyzed using a mean - standard-deviation space). 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.[10]

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.[11] Hence, subjects who search too much are more likely to revise their rules downward, so on average subjects will search too little.

The *third* explanation, which has not been mentioned in the existing 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.[12] When the stopping rule is a reservation price, then *only* prices that are lower than the reservation price can be a source for regret.[13] A price higher than the reservation price *cannot* be a source for regret, because the only such price observed will be the selling 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

[8] It is important to emphasize that each period consists of a complete search problem that may include *several* prices (bids) offered. Subjects are not allowed to adjust their stopping rules between bids, but only between periods.

[9] 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 $N(vN^{-2}) = vN^{-1}$.

[10] This 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.

[11] This can be seen by looking at the dashed line at the top panel of Figure 2. Figure 2 is discussed in further detail in Section 3.

[12] Sonnemans (1998) documents that the behavior of subjects after experiencing regret is different from their behavior in other periods.

[13] When subjects use a different stopping rule, other than a reservation price, the same effect still holds, but only on average.

device that "reminds" the individual to reevaluate her strategy.[14]

In the experiment described below I investigate the power of the latter two explanations, emphasizing the one related to observational regret.

3 Experimental Design

In this experiment, I rely on the third experiment in Sonnemans (1998), and use it as my control condition. Sonnemans's 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.[15]

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. Full 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 full, such optimal strategy depends on the prior, which may vary over subjects.[16]

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,[17] learning is a key part 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.[18] 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, which is known to the subjects, is discrete over all integers between 1 and 100 (cents), with equal probabilities (of 1 percent) for any price in the support of the distribution. 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.[19] Sonnemans (1998) replicates the too little search result in such a setting, so it seems reasonable to investigate this result even with these restrictions on search strategies.[20]

[14]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).

[15]The fixed cost per period is *not* a choice and has no effect on the optimal strategy. It is useful in directing subjects towards setting higher reservation prices, which speeds up learning and convergence. Pilot studies without fixed cost have shown that indeed the fixed cost is very useful.

[16]Still, even with full information, we should ask ourselves 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 sample from the distribution.

[17]See, for example, Roth and Erev (1995) and the references they cite.

[18]One may argue that a free search, without predefined strategy, is more realistic. Although true, such a procedure adds noise to the experimenter analysis of the result, once trying to infer the actual strategies (and adjustments) subjects use.

[19]Recall equation (1)–given this setting, the first element on the right hand side is $(100 - R + 1)/100$, and the second is $(100 - R)/2$. The right hand side is monotonically decreasing in R, surpassing the search cost of 2 between 80 and 81. Hence, the optimal reservation price is 81, the first integer that makes the expected gain from obtaining a new bid smaller than its cost.

[20]Sonnemans 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.

In addition, we define the concept of (*observational*) *regret*–subjects observe regret if their profits could have been higher by stopping earlier.[21] Using this definition, Sonnemans's results show that although regret is not very frequent (it occurs in less than 10 percent of the search problems), 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.

In the experiment described below, I build on Sonnemans's results, and test for the power of the two competing hypotheses for the "too little search" result: the asymmetric efficiency of the strategies around the optimal reservation price, and the asymmetry of information structure that leads to a one-sided regret. To do that, I form three experimental conditions. The first condition (the **Control Condition**) is identical to Sonnemans's, using the parameters described above and the same restrictions on predefined strategies–subjects have to set a reservation price before the beginning of each period. The second condition is similar to the first, but its subjects observe bids even after they have stopped searching. This allows for experiencing regret in both directions. The third condition is similar to the first, but has a different price distribution, which gives rise to the same optimal reservation price, but with a symmetric expected profit curve around the optimal strategy.

In the second condition (the **Symmetric Regret Condition**) I use a setup that gives rise to a more symmetric occurrence of observational regret. For that, I distinguish between two types of regret. The first, which is termed *upward regret*, is a case of regret, where a subject observes that he could have done better by setting a higher reservation price and continuing to search. The second, termed *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 Condition, subjects can only experience downward regret, because they do not observe bids after they stop searching.[22] To let subjects experience an upward regret as well, I let subjects in the Symmetric Regret Condition observe three additional bids that they would have gotten had they continued searching (by setting higher reservation price).[23] Other than that, subjects follow exactly the same procedure as subjects in the Control Condition. 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.[24]

In the third condition (the **Symmetric Distribution Condition**) I use a distribution that gives rise to a symmetric curve of expected profits. The equation for the curve is given by:

$$E\pi(R) = E_p(p|p \geq R) - \frac{c}{1 - F(R)} \tag{2}$$

In order to make it symmetric we should either change the cumulative distribution function $F(\cdot)$ or the search cost c. A nice potential property for the Symmetric Distribution Condition is to have the same

[21] 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.

[22] One may argue that in a repeated setting, subsequent periods may provide a source of upward regret. However, it seems reasonable to assume that subjects do not treat early bids in subsequent periods as potentially late bids in earlier periods, but do have in mind an independent sequence of potential bids for each period.

[23] Recall from the bottom part of Figure 2 that the expected number of bids received by setting the optimal reservation price is 5, so 3 is a reasonable number to choose for the additional post-purchase bids.

[24] A simple extension would be to form a fourth condition, in which subjects may choose whether they want to observe the additional prices or not. Their choice may indicate, for example, whether they use the additional prices for learning (and regret), or not.

optimal reservation price and the same expected profits as in the control group. This would allow us to compare easily the deviations from optimality in the different experimental conditions, because we have the same interpretation for "distance".[25]

These properties cannot be achieved by changing the search costs, so the change is made to the price distribution. One of the simplest non-uniform distributions is triangular. In addition, in order to have a symmetric expected profits curve, we must have a mass point at the highest point of the distribution (at a price of 100).[26]

Thus, the distribution used for the second condition is a triangular distribution over the same support (all integers between 1 to 100), with more weight given to lower values, but with a mass point (with probability of 9 percent) at a price of 100. Figure 1 shows the probability mass function of this distribution, together with the uniform probability mass function, which is used for the two other experimental conditions. The top part of Figure 2 shows the curves of the expected profits, as a function of the reservation price, for both distributions. Both curves peak at the same reservation price (81), giving rise to the same expected profits (80.5), and almost coincide (in the relevant range) below the optimal reservation price. However, the new distribution provides a much flatter expected profits curve to the right of the optimal reservation price.[27] Hence, we would expect that this distribution would make subjects who set prices above the optimal reservation price remain there, and not adjust downward as in the case of the uniform distribution. On average (and only on average) this should make subjects set reservation prices closer to the optimal. The bottom part of Figure 2 tries to provide the intuition for what may be expected to drive the individual behavior. The figure plots the expected number of bids for each reservation price. For the uniform distribution the expected number of bids starts shooting up from a certain point, and so do the search costs. For the second distribution, however, the mass point at 100 keeps the expected number of bids (and consequently the search cost) at a reasonable level, even for high reservation prices.[28]

4 Implementation and Results

After one 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

[25] Keeping the same optimal reservation price also helps to avoid psychological issues. For example, we might think that subjects converge better to, say, 61 than to 81.

[26] The Intuition for the need for a mass point at 100: set $R = 100$, then by lowering it to 99 we loose very little in expected price (first term in equation (2)), but earn much from the increase in probability of success (second term)–with "mass point" at 100, the earning in the second term is reduced, and the curve becomes flatter.

[27] It may be tempting to suggest a different distribution, that would give rise to a symmetric expected profits curve, which is steep (and not flat) in both directions, but I do not think that such a distribution exists.

[28] One potential problem of the Symmetric Distribution Condition is that it is not clear how well subjects can interpret the triangular distribution, compared to the uniform distribution used for the other conditions. To remedy this problem I give subjects, in all conditions, a sample from the distribution prior to each period. This sample is given on screen in a circle shape (see the instructions in the appendix). With many repetitions, subjects should get eventually a good feel for the distribution. Similarly, one may be worried that the additional bids that are observed by the Symmetric Regret Condition have also an informational effect, so that subjects in this condition obtain a better sense of the (known) distribution because they observe a larger sample from it. In order to deal with it, I let subjects in the Symmetric Regret Condition observe a smaller sample from the distribution prior to each period (15 bids instead of 25). I believe that after enough repetitions this kind of information differences does not make any difference for the results.

majority (but not all) of the subjects in this pool are undergraduate students from Harvard University and Boston University.

The complete instructions can be found in Appendix A. In Appendix B I present two exercises that were performed in order to make sure that the subjects understood the setup. Appendix C provides a description of the computer program, which was also given to the subjects before the beginning of the experiment. All of these were read aloud to subjects prior to the experiment.

Each of the runs was performed in a single session that included subjects from all conditions.[29] All sessions were held at the CLER, lasted about ninety minutes each, and included 29, 12, 25, 16, and 25 participants, adding up to a total of 107 subjects.

The random draws from the different distributions were drawn in advance, and were identical for all subjects in the same condition, and across sessions and conditions where possible (i.e. for the Control Condition and for the Symmetric Regret Condition, where the same distribution is used).[30]

In what follows I describe the experimental results. First I report the means and medians of earnings and reservation prices across experimental conditions. These results give a general idea of the behavior of subjects, and are most easily compared to results of previous experiments. However, the interpretation of these results is not straightforward and they can be driven by different sources. Therefore, I continue by looking at other dimensions of the experimental results. The next part looks at convergence issues, examining whether subjects converge throughout the 100 periods of the experiment, in an absolute or in a conditional sense. Next, I focus attention on observational regret, and give descriptive statistics of higher statistical moments of the results, those that are most relevant for investigating the effects of observational regret. While these descriptive statistics are quite informative, they lack the clean statistical test that can explicitly support the hypothesis regarding the effect of observational regret. Hence, the next part of this section looks at the same moments using multinomial logit analysis. This allows us to obtain statistically significant results, controlling for all other potential sources of variation.

The results I obtain from the multinomial logit analysis (Table 4 and Table 5) provide the strongest evidence in support of the hypothesis that observational is a learning-enhancing device. They 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 all my findings and discuss them.

[29] The instructions for all conditions are identical, except for the histogram of the price distribution, so I could read the instructions aloud even though the audience included subjects from different conditions.

[30] The identical draws eliminate some of the noise that can arise from the random differences in the draws between subjects. However, it is also true that it may cause some bias in the results, because of the sensitivity of the subjects to the different draws. Even if one thinks that given some kind of law of large numbers 100 periods are enough to make the randomness unimportant, it may still be the case that the initial periods are much more important in determining the final outcome, so the bias is not totally eliminated even after increasing the number of periods. 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.

4.1 Earnings

Table 1 presents the minimum, maximum and average earnings for each condition. As can be seen, subjects earned on average a little less than 30 dollars, with earnings varying between 20 dollars to about 32 dollars,[31] except for few subjects who earned very little. The earning distributions for the three different experimental conditions is very similar, and is driven mainly from the similarity in the expected profits curve (see Figure 2) and from its flatness around the optimal reservation price of 81. It is also interesting to observe that two subjects did better than they would have done by always using the optimal reservation price.

One might have expected that the flatness of the distribution of the Symmetric Distribution Condition will make earnings vary less in that condition. As can be seen in Table 1, the standard deviation of earnings is quite similar over the different conditions. This is because subjects in the Control Condition and in the Symmetric Regret Condition are all operating in the flat part of the expected profit curve, almost all the time. This observation can rule out in advance any concern that the Symmetric Distribution Condition has a fundamental difference, in terms of variance in payoff, because of the flatter expected profits curve.

4.2 Means and Medians

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 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, or that they just did not care. For comparing the means and medians over the different 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.[32] 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. Because of the experimental setup, subjects in the Symmetric Distribution Condition and in the Symmetric Regret Condition are more likely to get confused and to think that they do not understand the instructions. Therefore, comparing means with these "random" subjects in the sample is likely to bias the results downward in the Symmetric Regret Condition and Symmetric Distribution Condition, when compared to the Control Condition, and hence create problems in the interpretation of the simple comparison of means across the 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 Figures 3 and 4 show the average and median reservation price for each condition, as it evolved over the 100 periods of the experiment. As in previous experiments, which are discussed in Section 2, it is shown that, on average, subjects search too little. For all conditions, in all periods, the average reservation price is below the optimal reservation price of 81, making subjects stop searching too early. The

[31] Note that there is no point in comparing the earnings of subjects from the Symmetric Distribution Condition to the others. While their expected optimal earnings are the same, this is not the case for other reservation prices. As can be easily seen from Figure 2, being in the Symmetric Distribution Condition is almost never worse than being in the other conditions, and in particular it is better for high reservation prices.

[32] The criterion for dropping subjects was to drop any subject who chose reservation price lower than 50, for at least 5 times during the last 40 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 24 subjects (5 from the Control Condition, 9 from the Symmetric Distribution Condition, and 10 from the Symmetric Regret Condition). As could be predicted, more subjects got confused (and were dropped) in the "harder" conditions, the Symmetric Distribution Condition and the Symmetric Regret Condition.

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 results, when comparing the reservation prices set by both conditions to those set by the subjects in the Control Condition.

There are a few additional interesting results that should be mentioned. First, even subjects in the Symmetric Distribution Condition and in the Symmetric Regret Condition, who search more, still search too little, by setting average reservation prices that are significantly lower than optimal one, 81. Second, it is quite clear that the results are not driven by initial conditions. At least for subjects in the Control Condition and in the Symmetric Regret Condition, for whom the first period is completely identical,[33] subjects indeed start from almost the same average reservation price, as shown in the first row of Table 2. Note also that once we look at medians instead of averages, the subjects in the Symmetric Distribution Condition are much more like the ones in the Control Condition, suggesting that the high averages are driven mainly by a few subjects who set very high reservation prices (and do rather well, because of the flatness of the expected profits curve at the top of the distribution).

4.3 Convergence

Do subjects converge to some reservation price towards the end of the experiment? Here I want to discuss separately the concept of absolute convergence, in which all subjects (within the same condition) converge to a unique reservation price, and the concept of conditional convergence, in which each subject converges to her own (personal) reservation price.

Figure 5 tries to address the question of absolute convergence. It plots the standard deviation of reservation prices within condition over time. It can be seen that there is no convergence after the first 10 periods, and that the standard deviation remains pretty much constant, and significantly positive. It is interesting to note, however, that the standard deviation in the Symmetric Distribution Condition is significantly greater than in the other conditions. This is what we could expect–because of the flatness of the distribution, subjects in the Symmetric Distribution Condition do not tend to converge to the optimal reservation price, so they remain very spread. This is not the case for subjects in the Control Condition and in the Symmetric Regret Condition, who get better feedback on their strategies and react accordingly. Indeed, consistent with the hypothesis, subjects in the Symmetric Regret Condition receive the best feedback, and hence have the smallest standard error. This is important in analyzing the implications of the results. The fact that subjects in the Symmetric Distribution Condition search better *on average* does not imply that they search better on an individual basis. Subjects in the Symmetric Regret Condition, who search better on average but are also more concentrated around this better average, are undoubtedly doing better than those in the Control Condition. We can also see this better outcome by looking at the earnings, as reported in Table 1.

Figure 6 tries to address the question of conditional convergence. Here I plot the compute the standard deviation for each subject's last 10 reservation prices, and plot the averages of these standard deviations, within condition. For example, if all subjects fixed their (possibly different) reservation prices for the past 10 periods, the value for the last period would have been zero. Figure 6 shows that there is a tendency for conditional convergence, although it never goes all the way to zero, so even in the last 10 periods subjects

[33] At the very first period subjects from the Control Condition and subjects from the Symmetric Regret Condition 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 conditions they were selected.

continue to adjust their reservation prices.[34] This result may be interpreted in two ways. One interpretation is that 100 periods are not enough, and conditional convergence can be obtained eventually, if the experiment lasted longer. A different interpretation, which seems more appealing, is that at least some of the subjects do converge to some kind of mixed strategy, in which they keep changing their reservation prices, according to some rule of thumb.

4.4 Observational Regret as a Learning-Enhancing Device–Descriptive Statistics

So far, I analyzed and discussed the general results of the experiment. From now on, I want to focus more on the effects of observational regret. Because of the flatness of the distribution in the Symmetric Distribution Condition, subjects in this condition, as is also noted in their post-experiment reactions, behave much more randomly, and do not react too much to information obtained from previous periods. Hence, whenever we look at learning, there is very little learning going on in the Symmetric Distribution Condition. In what follows I report all the figures for all conditions, but concentrate in the discussion on the comparison between the Control Condition and the Symmetric Regret Condition.

In analyzing observational regret, it is useful to make comparisons with an artificial fourth experimental condition, which is named the **Pseudo Control Condition**. This comparison looks at the behavior of subjects in the Control Condition as if they were observing the three additional prices like the subjects in the Symmetric Regret Condition.[35] This allows us to better control for the results associated with the Symmetric Regret Condition. That is, we control for the possibility that observational regret itself has no effect, but is simply correlated with some other relevant variables. For example, it is natural to assume that subjects will adjust their reservation price downward after setting a relatively high reservation price, and upward after setting a relatively low reservation price. However, downward regret is also more likely to be observed after setting a high reservation price, and upward regret is more likely to be observed after setting a low reservation price. Hence, an adjustment upward may be positively correlated with upward regret only because of statistical reasons, without any real effect of observational regret. By using the Pseudo Control Condition, and comparing it to the Symmetric Regret Condition, I am able to obtain a perfect control, so that *actual observation* of regret is the only difference between the subjects.

Figures 7 and 8 provide a descriptive statistical "story" of the subjects' behavior, with emphasis on observational regret. I look at two key variables–the change of the reservation price and the observational regret–and measure the conditional and unconditional frequencies in each condition, over all subjects and periods. Each quadruplet stands for the empirical frequencies for the four conditions–the Control Condition, the Symmetric Distribution Condition, the Symmetric Regret Condition, and the Pseudo Control Condition.

The idea of the structure of the figure is to mimic a decision tree for the subjects. After obtaining a bid sequence, there are three mutually exclusive possibilities–no observational regret, downward regret, or upward regret. Then, the subject first decides whether to adjust (change) her strategy or not. If she changes, she decides whether to change it up or down, and in each case by how much. The empirical frequencies of these hypothetical decisions are depicted in Figure 7 and 8. Figure 7 reports the pooled statistics for the

[34] These convergence results are qualitatively similar to those obtained from the experimental data of Sonnemans (1998).

[35] 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.

whole 100 periods, and Figure 8 for the last 50 periods, where it is reasonable to assume that subjects have already understood the setup and adopted some behavioral rule.

First, it is interesting to look at the unconditional frequencies. Downward regret is experienced much less often than upward regret. In addition, subjects often adjust their reservation prices, almost every other period.[36] Looking at the conditional frequencies, I find some confirmation for the effects of regret, at least for the Control Condition and for the Symmetric Regret Condition. After observing a downward regret, subjects in these conditions are much more likely to change their reservation price (compare the 73% and 64% to the 55% and 46% in Figure 7). When they do, it is mostly downward. If they do change upward, the magnitude is much smaller. After experiencing an upward regret the story is reversed - subjects in the Symmetric Regret Condition (the only condition that can experience an upward regret in the experiment) are more likely to change their reservation price upwards.

One may argue that observational regret makes subjects search more, but not necessarily better. It happens that they start from too little search so searching more is better, the argument may continue. One way to address it is by looking at the way subjects in the Symmetric Regret Condition react to downward regret.[37] It is suggested by Figure 8 that in the last 50 periods not only do subjects in the Control Condition seem to react to upward regret, but also to better react to downward regret. 93% of the times that they observe it and decide to change their strategy, they change their reservation price downward. Subjects in the Control Condition do it only 82% of the times. This may suggest that once subjects have the opportunity to observe regret symmetrically, they may become better learners in *both* directions. From looking at the individual behavior, it can be observed that after setting a reservation price above the optimal one, subjects in the Symmetric Regret Condition do adjust downward quite fast, in about 10 periods. In sharp contrast, subjects in the Symmetric Distribution Condition do not adjust back, and the few subjects that shoot "too high" remain there and never adjust downward towards the optimal reservation price. This phenomenon is exactly the one that makes the subjects in the Symmetric Regret Condition converge, in an absolute sense, much better towards the optimal reservation price, when compared to subjects in the Symmetric Distribution Condition.

As argued before, without trying to control for other variables, the interpretation of the results is not transparent. The results may stand for some other variables that are statistically correlated with regret. This leads us to look at the Pseudo Control Condition, which is supposed to control for all these other variables. The only difference between the Symmetric Regret Condition and the Pseudo Control Condition is the actual observation of regret. Looking at Figure 7 and 8, I do find indeed that subjects in the Pseudo Control Condition are also more likely to change their reservation price upward after (pseudo) observing a (pseudo) upward regret, probably exactly because of this statistical reason. However, subjects in this condition are much closer to the no-regret 52-48 split than the subjects in the Symmetric Regret Condition (compare the 70.33% with the 60.05% in Figure 7).

All the results discussed above remain more or less the same in Figure 8, which replicates them for the last 50 periods. Moreover, some of the results, which are highlighted by the arrows in the figure, become even more extreme in distinguishing between the groups. In particular, it is shown that subjects in the

[36] However, it is important to note that the individual behavior shows that subjects vary a lot in that aspect. Some of the subjects adjust their reservation prices almost every period, while others almost never adjust.

[37] We can also design a new experiment to address this critique. In such an experiment we can use two settings, such that in one it is more likely that subjects initial reservation price would be below the optimal (as in our setting), and in the other subjects will start searching from a higher than optimal reservation price. Note, however, that finding a setup that fits the latter condition is not an easy task.

Symmetric Regret Condition become even better learners and react more systematically to both types of regret after obtaining experience of 50 periods, while subjects in the (pseudo) Control Condition get closer to the unconditional behavior. In particular, after (pseudo) observing a (pseudo) upward regret and changing the reservation price, subjects in the Pseudo Control Condition adjust upward only 55% of the times, and when they do, the magnitude of the change is quite low.

Table 3 tries to provide another layer of behavioral effect to observing regret. It shows the average time (in seconds) that it takes for subjects to decide on their new reservation price. While for the control group this average time varies very little once conditioned on observing regret in the previous period, it does somewhat vary for the Symmetric Regret Condition. Subjects in this condition react much faster once they do not observe any regret, suggesting that the potential observation of regret teaches them also to obtain a better feedback from not observing any regret.

4.5 Observational Regret as a Learning-Enhancing Device–Multinomial Logit Analysis

In the last part of the analysis, I provide stronger statistical evidence about the effect of observational regret. I run multinomial logit regressions[38] in which the dependent variable is categorical, and can take three possible values: a change upward in the reservation price, a change downward or no change.[39] The explanatory variables include a constant and the last reservation price as controls, and dummy variables for downward and upward regret (whenever relevant). I continue to use the Pseudo Control Condition as a (pseudo) forth condition in order to make sure that there are no other possible variables for which the regret may proxy.

Table 4 reports the regression results and Table 5 reports the corresponding odds ratios. The regressions were run separately for each condition, pooling all periods together, but correcting the reported standard errors for possible dependence among observations of the same subject over time.[40] Looking at the results, it is first important to observe that the reservation price, which is used as a control, almost always has the expected sign in the sense that higher reservation price are likely to be adjusted downward and lower ones are likely to be adjusted upward.[41]

The interesting and strong evidence comes from looking at the coefficients on the regret dummy variables, or at the corresponding odds ratios as reported in Table 5. Downward regret, which is observed by all conditions, always increases the probability of changing the reservation price downward, and is *always* significant at the 1% level. After experiencing downward regret, subject in the Control Condition and in the Symmetric Regret Condition are three to four times more likely to adjust their reservation price downward (compared to not making any adjustment). Consistent with the descriptive evidence, this effect is lower (although still very significant) for the Symmetric Distribution Condition, for which the corresponding odds

[38]Results from ordered probit regressions of the same specification are very similar, and lead to identical conclusions.

[39]I also tried to use the magnitude of change as a continuous dependent variable, but this specification turned out to be very sensitive to a few outliers. In addition, I think that any behavioral model can try to explain the direction of change, but will have much more difficulties trying to give structure to the magnitude of change, which is much more arbitrary.

[40]Alternatively, all conditions can be pooled together, adding dummy variables for some of the conditions, and interactions of these dummies with the other explanatory variables. Such regressions were run and provided very similar results. The tables are available from the author upon request.

[41]Similar regressions were run with additional controls, such as the number of bids obtained in the last period and the earning, and produced quite similar results. However, the significance levels of the coefficients on the regret dummies were lower, because of probable multicolinearity among the explanatory variables.

ratio is below 2. In addition, it is shown that downward regret has no significance effect on the probability of changing the reservation price upward, in all experimental conditions.

Similarly, the coefficients (and odds ratios) on upward regret give the mirror image. Upward regret has no significant effect on the probability of changing the reservation price downward, but once it is experienced it increases significantly the likelihood of changing the reservation price upwards, by more than 50%. For upward regret we also have the perfect control group, the Pseudo Control Condition, for which it has no statistically significance effect. This result provides another layer of evidence that it is the *actual* observation of regret that makes subject revise their strategies, and not any probabilistic anticipated regret.

4.6 Summary of Findings and Discussion

The preceding section obtained several important results. First, they replicate well previous results that subjects search too little. In this setup it is shown that subjects in all conditions set on average a reservation price which is significantly lower than the optimal reservation price of 81. Second, the results document quite strongly the effect of *actual* observation of regret, in both directions, upwards as well as downwards. The analysis suggests that experiencing regret can be thought of as a feedback to the subjects, which enhance their learning. Once the feedback is objective, in the sense that it is symmetric in both directions, subjects react accordingly and converge better towards the optimal reservation price. It is also suggested that the possibility of obtaining a symmetric feedback may make subjects better learners, by interpreting feedbacks better in both directions. In addition, I find support for the hypothesis regarding the asymmetric expected profits curve. Once it is symmetric subjects do not adjust their reservation prices downwards as often, and hence end up setting higher reservation prices on average, closer to the optimal one. However, as may be expected, once the expected profits curve is flat, the hypothesis works only for the aggregate, but does not make subjects search better at the individual level.

It may be useful to take the results further, by constructing a reinforcement-learning model and using the data to fit its parameters, in the spirit of Roth and Erev (1995). Such a model should take into account the significant effect of observational regret as a learning-enhancing device. It should be noted, however, that this search model is quite complex, and hence there are many modeling choices for the exact strategies, and for the exact process by which subjects update their score. Hence, it is not clear how one can take these modeling choices into account when trying to use them for predicting behavior in similar, but not identical, situations.[42] Appendix D provides a simple implementation of such a model. However, when trying to fit the model to the data, it turned out that the results are quite sensitive to different modeling choices, and that the parameters vary a lot in response to small changes in the specification of the learning model. I conclude from this result that the experimental data in this setup does not identify well the exact ways by which observational regret affects learning, and that there is a need for a much simpler environment for testing such learning models.

Although the evidence presented here does not identify exactly what kind of learning is going on, it strongly suggests that learning, through observational regret, plays an important role in determining behavior.

[42] Similar arguments can be made for other learning models.

5 Conclusions and Further Research

Previous experimental search papers have consistently found that subjects in experiments search too little. In this paper I test for two possible explanations for this robust result, and find evidence that both explanations contribute to less than optimal search. Subjects who observe a price distribution that makes the cost of sub optimal strategies more symmetric increase their search intensity. Similarly, subjects who can evaluate their strategies more symmetrically, by observing post-purchase information, search more intensively as well. All subjects, however, still search less than the optimal strategy suggests.

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 very 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 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.[43] In the strategy method subjects are informed about the full strategy of their opponent, so they can calculate their unrealized payoffs. In this sense, the strategy method is similar to the Symmetric Regret Condition and allows for observational regret, while the action method is similar to the Control Condition, 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 implications–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.[44]

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.

Another related empirical question regards the practice of "rebate-clause" or "best-price-clause", according to which a seller guarantees the consumer to compensate her if a cheaper price is found for the product later on. Typically, economists regard these practices as ones that are developed in order to facilitate col-

[43] 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.

[44] For example, we can think of two markets for the same product, a regular one and an Internet one. Post-purchase information is likely to be more widely available in the former. Alternatively, one can think of two 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.

lusion. However, the same practices can be rationalized from a regret perspective. Sellers may offer such terms to consumers in order to reduce consumer search, thereby softening price competition and allowing higher prices to prevail in equilibrium.

References

[1] Berg J., Daley L., Dickhaut J., and O'Brien J. (1986), "Controlling Preferences for Lotteries on Units of Experimental Exchange," *Quarterly Journal of Economics*, 101(2), 281-306.

[2] Butler D. and Loomes G. (1997), "Quasi-Rational Search under Incomplete Information: Some Evidence from Experiments," *The Manchester School*, 65(2), 127-144.

[3] Hart S. and Mas-Colell A. (2000), "A Simple Adaptive Procedure Leading to Correlated Equilibrium," *Econometrica*, 68(5), 1127-1150.

[4] Hey J. (1982), "Search for Rules for Search," *Journal of Economic Behavior and Organization*, 3, 65-81.

[5] Hey J. (1987), "Still Searching", *Journal of Economic Behavior and Organization*, 8, 137-144.

[6] Kagel J. and Roth A. (1995), *The Handbook of Experimental Economics*, Princeton University Press.

[7] Kogut C. (1990), "Consumer Search Behavior and Sunk Costs," *Journal of Economic Behavior and Organization*, 14, 381-392.

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

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

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

[11] Roth A. and Malouf M. (1979), "Game-Theoretic Models and the role for Bargaining," *Psychological Review*, 86(6), 574-594.

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

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

[14] Stigler G. (1961), "The Economics of Information," *Journal of Political Economy*, 69, 213-225.

[15] Telser L. (1973), "Searching for the Lowest Price," *American Economic Review*, 63(2), 40-49.

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. $Earning = Sale\ Price - 2 \times Number\ of\ Bids - 50$). 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.

B Pre-Experiment Exercises

Before starting the experiment, we want to make sure that the instructions are clear, and that the setup is well understood. For that purpose, you should try to answer the following questions. Your earnings in the experiment are not affected by the correctness of your answers. However, a good understanding of the setup will help you do better in the experiment. After you hand in your answer sheet, a solution set will be distributed, so you can check yourselves.

Example: Consider the following potential sequence of bids:

$$65, 24, 70, 87, 67, 54, 100, ...$$

Suppose you set in the beginning of the period a minimum price of 50, then bidding would stop after one bid, the last bid would be 65, costs would be 2, and earnings would be $65 - 2 - 50 = 13$, which is 13 cents. Alternatively, a minimum price of 66 would have made bidding stop after 3 bids, with earnings of $70 - 6 - 50 = 14$, which is 14 cents.

Make sure you understand the above, by going through the following questions:

Question 1: Consider the following potential sequence of computer bids:

$$43, 56, 32, 78, 12, 17, 84, 54, 67, ...$$

Complete the following, in the case that the minimum price was set to 65:

Number of bids	___
Last bid	___
Costs for the period	___
Earnings for the period	___
Money earned this period	___

Question 2: Consider the same potential sequence of computer bids as in question 1. Complete the following, in the case that the minimum price was set to 84:

Number of bids	___
Last bid	___
Costs for the period	___
Earnings for the period	___
Money earned this period	___

C 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.

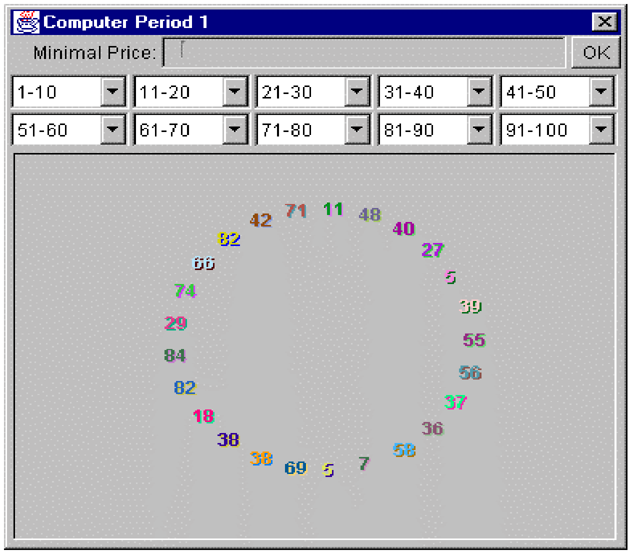

Now bidding will take place, as can be seen in the screen below. You can get new bids by clicking the “Next Bid” button (left bottom corner) for each new bid. When bids take place, they appear on the left side of your screen. There are three columns - the bid number, the bid, and the amount you would 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.

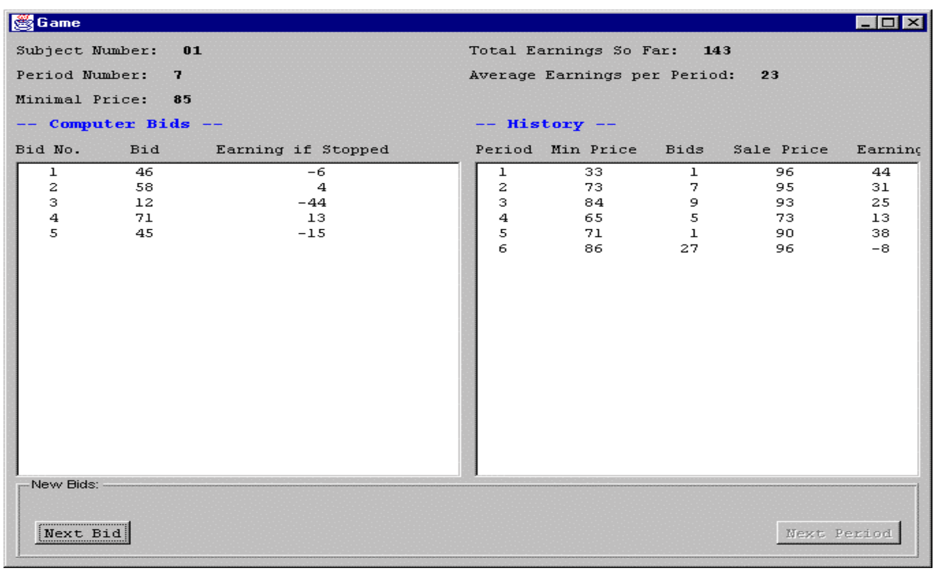

After you are done, the last screen will pop up. Fill in your name and social security number and click the “Print and Quit” button. We will collect the printouts from the printer, prepare your payment and call you to pick it up. The receipt is the only place in which your name and social security number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

D An Example for a Reinforcement Learning Model

In this appendix I want to give an example for a reinforcement model for the experimental search setting that can incorporate the effect of observational regret. It is important to note that in this setting it is not very plausible to use a simple two-parameter reinforcement-learning model, as suggested by Roth and Erev (1995). On one hand, subjects in this model have 100 possible strategies (any reservation price between 1 to 100), so one may think that reinforcement learning may take many periods. However, the main difference between the search problem and the games analyzed by Roth and Erev is that in the case of a search problem a subject that plays a certain strategy receives information about the payoff he would have earned had he played other strategies (lower reservation price). Hence, it does not seem plausible to make subjects increase the propensity score only for the strategy he played, knowing that many other strategies would have given him the same payoffs. In addition, the payoffs of a strategy in a one-period game are quite random, and in my opinion even a very unsophisticated subject realizes it. Hence, his propensity of playing the strategy should not be affected by it.

I suggest two modified models that address these issues. Both are more complex than the simple two-parameter model. The first model begins with an initial propensity, as the standard reinforcement models. The update of the propensity is not done only at one point, but over an interval, with an isosceles triangular shape of update. The triangle is centered around either the previous reservation price plus 1 (trying to capture the fact that if a subject did well, he would try to shoot for more), if no regret was experienced, or around the highest price that would have achieved the best outcome observed, if regret was experienced. In addition, while there is an update of the propensity after every period, I allow a change of the strategy only with some probability (to capture what we observe in the data - strategies are revised only about half of the times). This probability is greater after regret is experienced, as the data suggest. I do not fit parameters, but just simulate the model with calibrated parameters, for 100 subjects and 500 periods. The top part of Figure 9 presents the results. The comparative effect is clear - subjects that observe the additional prices search much better, closer to the optimal reservation price. From absolute perspective, the model needs a better fit. The model predicts too low reservation price for the Control Condition, and especially too high reservation price for the Symmetric Regret Condition. As argued in the text, fitted parameters were not robust to slight changes in the model specification.

The second model requires higher degree of rationality from subjects. In such a model, the subjects understand that there is a full interval of reservation prices that would have given them the same result, and hence update their propensities equally over the whole interval of prices that would have obtained the best outcome observed. The rest of the model is similar to the first one. Its calibrated simulations are presented in the bottom part of Figure 9. For this model, the comparative relationship is similar, and the absolute prediction is much closer to the data than that of the first model, although we still obtain higher reservation prices than those observed.

It should be noted that for both models, the results mimic the experimental results only for the aggregate (of each condition). The individual level simulations are not similar to the experimental results by any mean. The difficulties of fitting the individual level behavior are the source for the sensitivity of the reinforcement-learning model results to the exact specification.

Table 1: Earnings

	Control	Symmetric Distribution	Symmetric Regret
N	35	32	40
Min	11.98	12.26	7.51
Mean	26.17	26.56	26.23
Median	28.18	28.91	28.74
Max	31.29	32.2	31.02
Std. Dev.	9.36	8.45	10.3
Optimal*	30.9	31.99	30.9

Table 1: Earnings

* "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).
Recall that while subjects in all conditions face the same expected profits when using the optimal reservation price, Figure 2 shows that subjects in the Symmetric Distribution Condition face higher expected profits for sub optimal reservation prices (in part of the relevant range), so that it is of little interest to compare the average earnings in this condition to that in the others.

Table 2: Means and Medians

All Subjects:

	Control	Symmetric Distribution	Symmetric Regret
First Period	63.3 (2.76)	64.9 (3.21)	63.1 (3.18)
All Periods	68.7 (2.33)	68.3 (2.84)	69.4 (2.26)
Last 50 Periods	70.8 (2.41)	69.5 (2.97)	70.8 (2.32)
Last 25 Periods	71.3 (2.37)	70.2 (3.12)	70.8 (2.43)
Last 10 Periods	72.0 (2.46)	70.5 (2.17)	71.3 (2.43)

All "Non-Random" Subjects:

	Control	Symmetric Distribution	Symmetric Regret
First Period	66.5 (2.17)	72.7 (2.69)***	66.7 (3.38)
All Periods	72.6 (1.83)	76.1 (2.25)*	75.2 (1.45)*
Last 50 Periods	75.1 (1.48)	77.9 (2.21)*	77.1 (1.19)*
Last 25 Periods	75.1 (1.57)	78.6 (2.41)*	77.6 (1.12)**
Last 10 Periods	75.2 (1.66)	78.6 (2.48)*	77.9 (1.23)**
Last Period	76.0 (1.78)	78.3 (2.61)	79.1 (1.52)**

Table 2: Means and Medians

*** 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

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 footnote 32. This table and the corresponding Figures 3 and 4 are the only place in the paper where such selection is applied. All other analysis of the paper is based on the full sample.

Table 3: Decision Times

After downward regret:

	Control	Symmetric Distribution	Symmetric Regret
N	118	187	155
Mean	7.38	7.17	7.15
Std. Dev.	5.13	5.57	6.75

After no regret:

	Control	Symmetric Distribution	Symmetric Regret
N	2852	2090	2149
Mean	7.42	7.86	6.65
Std. Dev.	8.53	7.13	7.35

After upward regret:

	Pseudo Control	Symmetric Distribution	Symmetric Regret
N	744	NA	666
Mean	7.6	NA	7.54
Std. Dev.	8.75	NA	7.84

Table 3: Decision Times

The decision time is measured in seconds, starting when the reservation price screen is shown, until the subject sets the new reservation price and clicks on the OK button.
The "Pseudo Control Condition", as explained in the text, is the subjects in the control condition, treated as if they could observe the next three additional bids, after the end of the search.

Table 4: Multinomial Logit Regressions

	Control		Symmetric Distribution	
Dependent Variable	Downward Change	Upward Change	Downward Change	Upward Change
Const	-1.3	1	4.62*	7.59**
	-0.95	0.61	2.91	4.43
Reservation Price	0.01	-0.02	-0.06*	-0.10**
	0.54	-0.91	-3	-4.51
Downward Regret	1.29**	-0.02	0.65**	0.34
	3.51	-0.04	3.3	1.41
No. of Observations	2970		2277	
Log-Likelihood	-3119.32		-2248.51	
Pseudo R-Squared	0.02		0.1	

	Pseudo Control		Symmetric Regret	
Dependent Variable	Downward Change	Upward Change	Downward Change	Upward Change
Const	-1.35	2.34	-0.93	3.24
	-0.99	2.71	-0.45	1.67
Reservation Price	0.01	-0.02	0.001	-0.05*
	0.56	-0.83	0.052	-2.17
Downward Regret	1.41**	0.07	1.25**	0.1
	3.8	0.16	4.26	0.23
Upward Regret	0.09	0.18	-0.032	0.44**
	1	1.67	-0.25	4.18
No. of Observations	2970		2970	
Log-Likelihood	-3116.13		-2909.05	
Pseudo R-Squared	0.02		0.06	

Table 4: Multinomial Logit Regressions

** significant at the 1% level
* significant at the 5% level

Z-Stats below coefficients (corrected for possible dependence between observations)
Corresponding odds ratio for the coefficients of interest are provided in Table 5
In all regressions, the base category is no change in the reservation price

Table 5: Multinomial Logit Regressions–Odds Ratios

Dependent Variable	Control: Downward Change	Control: Upward Change	Symmetric Distribution: Downward Change	Symmetric Distribution: Upward Change
Downward Regret	3.63**	1.02	1.92**	1.40

Dependent Variable	Pseudo Control: Downward Change	Pseudo Control: Upward Change	Symmetric Regret: Downward Change	Symmetric Regret: Upward Change
Downward Regret	4.09**	1.07	3.48**	1.11
Upward Regret	1.10	1.20	0.97	1.55**

Table 5: Multinomial Logit Regressions–Odds Ratios

** significant at the 1% level
* significant at the 5% level

The above is the mirror image of Table 4, providing odds ratio instead of estimated coefficients for the coefficients of interest. In all regressions, the base category is no change in the reservation price
The figures in this table show the change in the odds ratio after observing a regret, compared to the odds ratio after a period with no observational regret. The odds ratio is the probability of a downward/upward change, divided by the probability of no change (the base category). For example, for the Symmetric Regret Condition, the average subject is about 3.5 more likely to revise her reservation price downwards after observing a downward regret.

Figure 1: Price Distributions

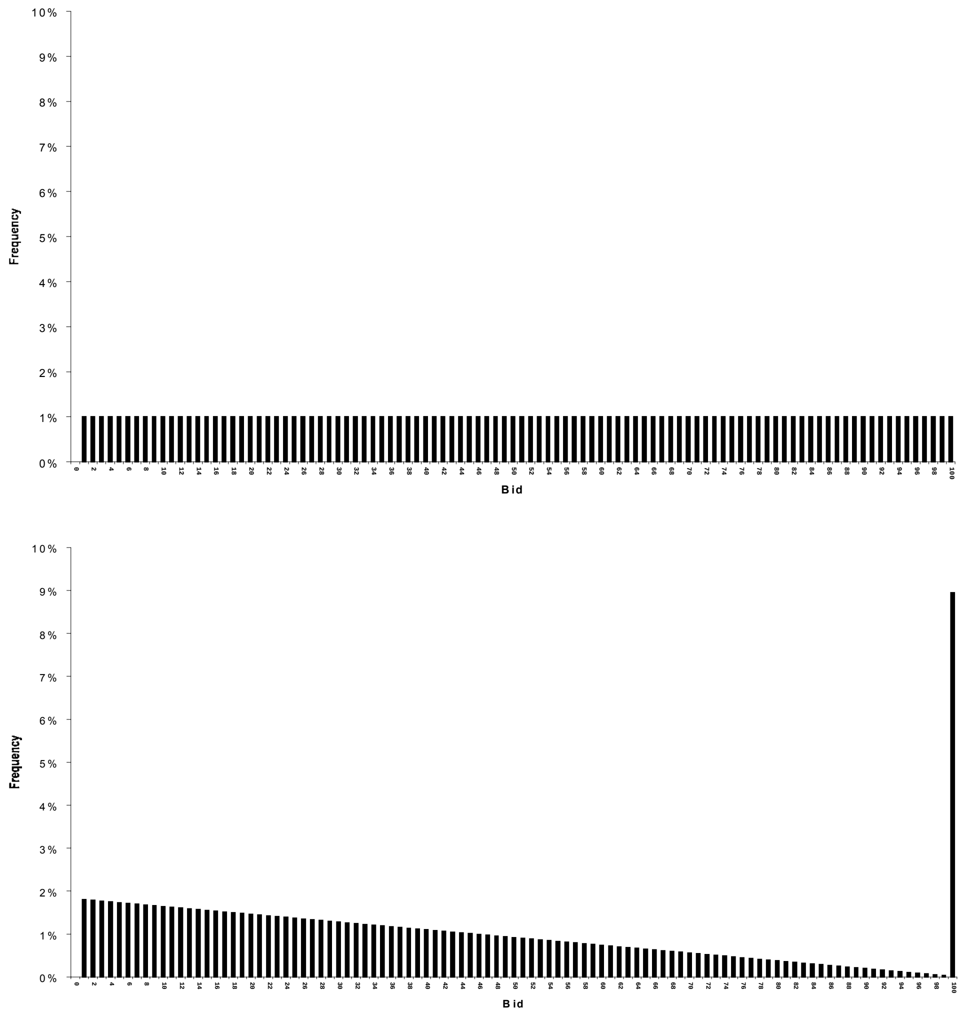

Figure 1: Price Distributions

The top panel shows the uniform price distribution used for the Control Condition and for the Symmetric Regret Condition. The bottom panel shows the price distribution used for the Symmetric Distribution Condition. These figures are identical to those provided to subjects prior to the experiment.

Figure 2: Properties of the Distributions

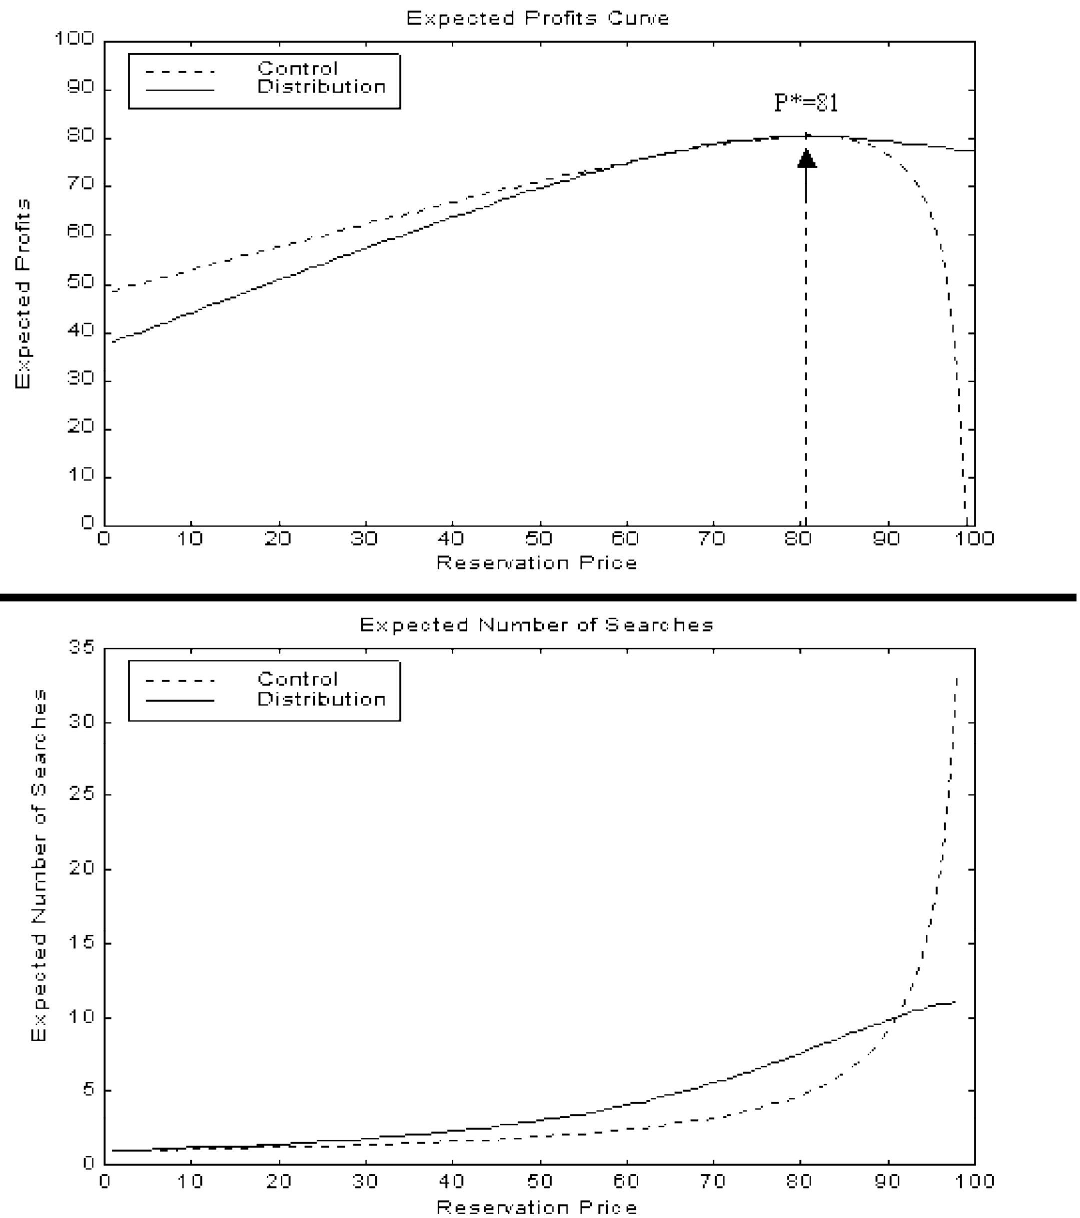

Figure 2: Properties of the Distributions

The dashed lines ("Control") correspond to the uniform price distribution (used for the Control Condition and for the Symmetric Regret Condition). The solid line ("Distribution") corresponds to the triangular distribution (used for the Symmetric Distribution Condition).

For each distribution, the top panel shows the expected profits as a function of the reservation price, and the bottom panel shows the expected number of bids as a function of the reservation price.

The optimal reservation price for both distributions is 81 (cents).

Figure 3 and Figure 4: Avergaes and Medians over Time

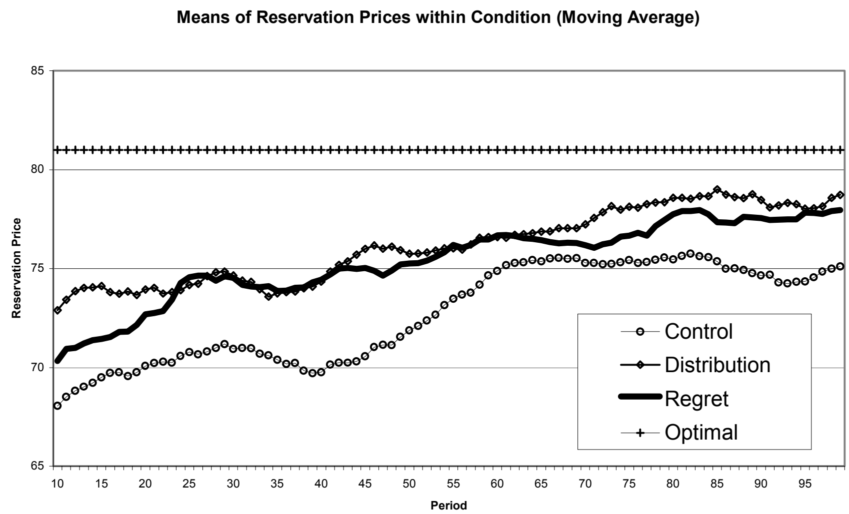

Figure 3:

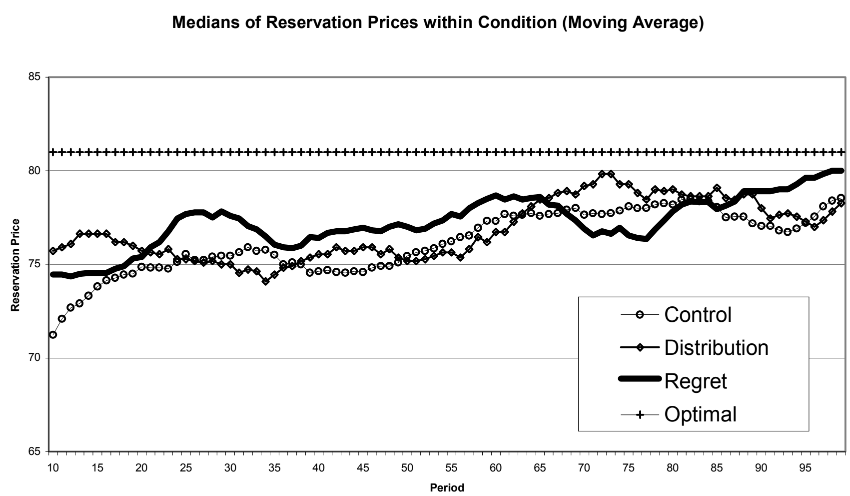

Figure 4:

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 footnote 32. 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 5 and Figure 6: Measures of Convergence

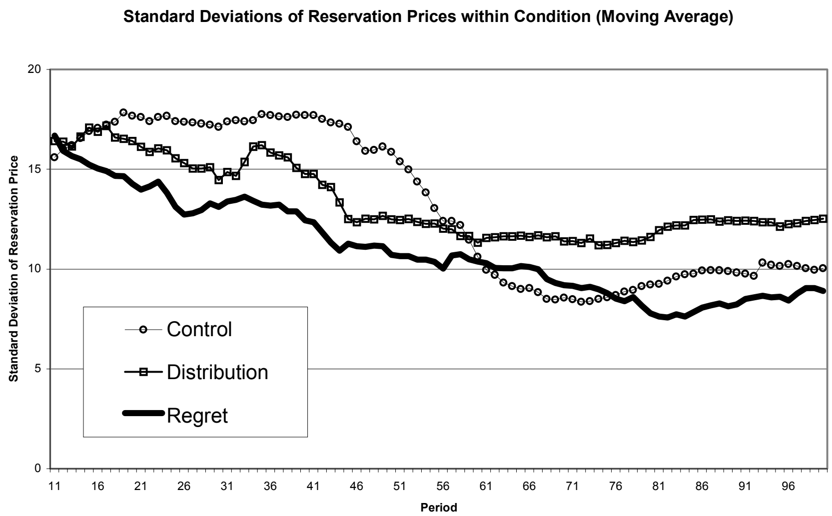

Figure 5:

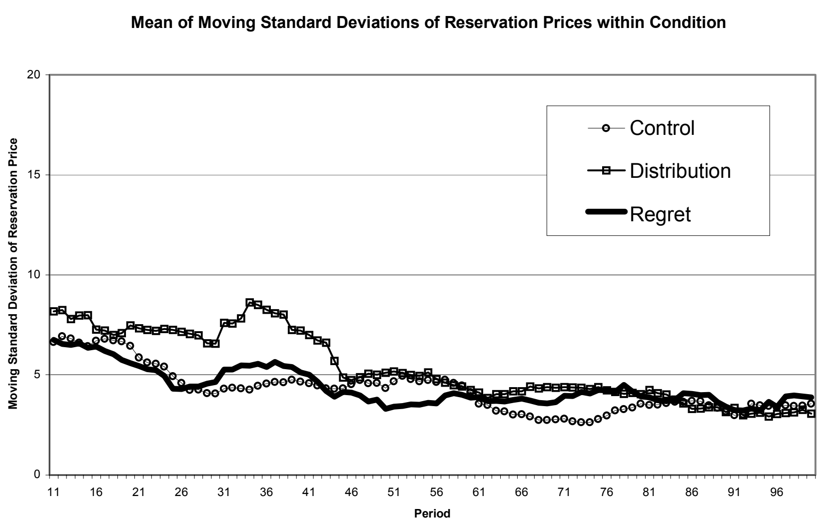

Figure 6:

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 standard deviations.

The top part graphs the standard deviation of reservation prices across subjects within condition. The bottom part graphs the mean (across subjects within condition) of the standard deviation of the reservation prices of each subject over the last 10 periods.

Figure 7: Descriptive Statistics–All Periods

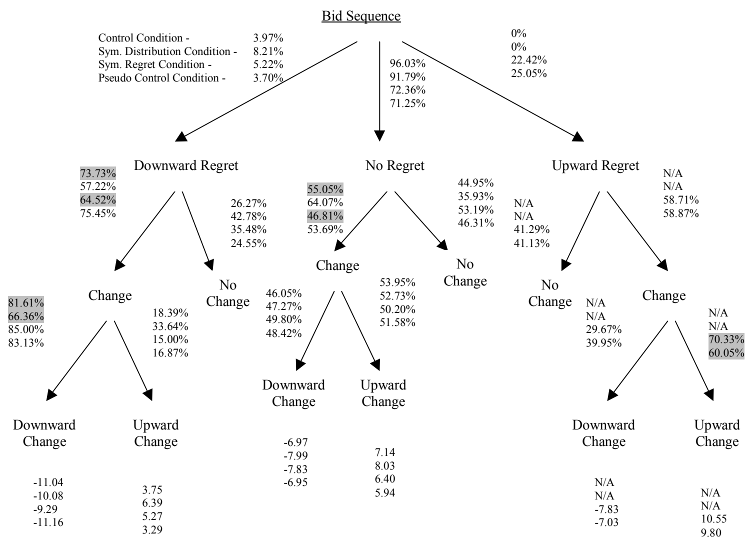

Figure 7: Descriptive Statistics–All Periods

Each Quadruplet stands for the four different conditions–Control Condition, Symmetric Distribution Condition, Symmetric Regret Condition and Pseudo Control Condition–as shown in the upper left corner of the figure.
The figure shows the frequency or actual observation of regret, the frequency of changes to reservation price in each case and its direction, and the mean change at the bottom part.
The highlighted numbers are those that show the most important differences between the different conditions.
More thorough discussion of the figure is given in the text (Section 4.4).
Recall: Downward Regret–profit could have been higher by stopping earlier (setting lower reservation price), Upward Regret–profit could have been higher by stopping later (setting higher reservation price).

Figure 8: Descriptive Statistics–Last 50 Periods

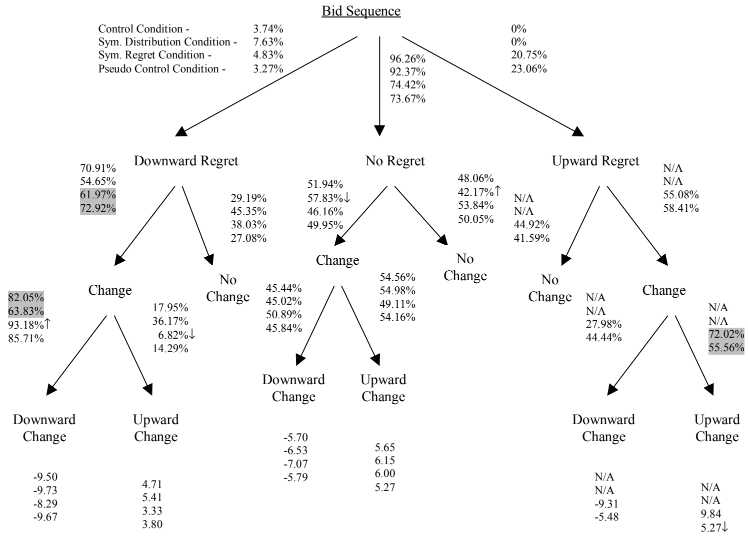

Figure 8: Descriptive Statistics–Last 50 Periods

This figure is identical to Figure 7, but is based only on the last 50 periods, in order to capture the behavior of "experienced" subjects. The major differences from Figure 7 are given by the small arrows.

Each Quadruplet stands for the four different conditions–Control Condition, Symmetric Distribution Condition, Symmetric Regret Condition and Pseudo Control Condition–as shown in the upper left corner of the figure.

The figure shows the frequency or actual observation of regret, the frequency of changes to reservation price in each case and its direction, and the mean change at the bottom part.

The highlighted numbers are those that show the most important differences between the different conditions.

More thorough discussion of the figure is given in the text (Section 4.4).

Recall: Downward Regret–profit could have been higher by stopping earlier (setting lower reservation price), Upward Regret–profit could have been higher by stopping later (setting higher reservation price).

Figure 9: Simulations of Reinforcement Models

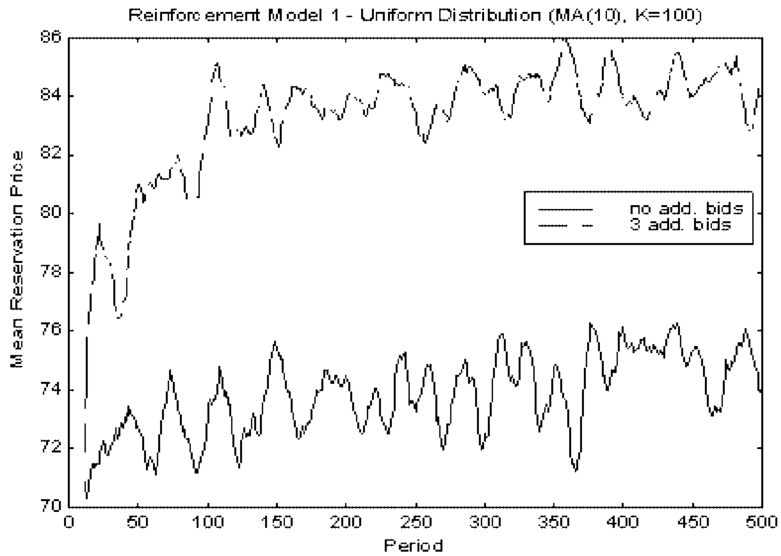

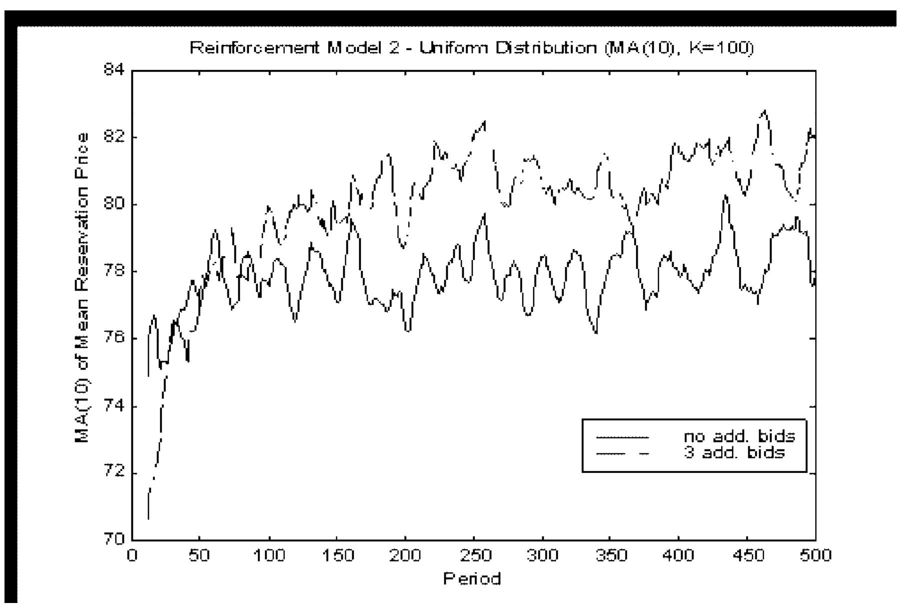

Figure 9: Simulations of Reinforcement Models

See Appnedix D for a discussion of this figure.