Sampling for Information or Sampling for Imitation: An Experiment on Social Learning with Information on Ranks*

Theo Offerman
University of Amsterdam

Andrew Schotter
New York University

July 28, 2005

Abstract

In this paper, we present the results of an experiment on social sampling. People are asked to make a risky decision in a situation where an idiosyncratic luck term affects their performance. Before they make their decision, people have the opportunity to sample others who have done exactly the same problem before them. These previous participants (subjects in our "Seed" treatment) are ranked on the basis of their success. A subject may sample a person in the Seed with the goal to retrieve information that is valuable to a rational decision maker. Alternatively, she may sample successful people in the Seed, who tend to be lucky risk takers, and simply imitate what they did. We find that, by and large, subjects sample successful people and imitate what they do. Thus, social sampling makes people look much more risk seeking than the people who do not have the opportunity to sample. We obtain similar results for the case where we allow subjects to sample only once and for the case where we allow subjects to sample three times.

*The authors would like to thank CREED-programmer Jos Theelen for programming the experiment. The research of Theo Offerman has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

1 Introduction

Consider the following two problems.

Problem 1: A CEO of a major corporation is contemplating a decision. It could be whether to enter a new market, whether to expand the production capacity, or whether to start a new line of products. At a meeting of the board they decide to form a committee to investigate the issue and one of the first things the committee does is to sample other firms and investigate those who have faced this problem and have been successful. In other words, the strategy is sample and imitate the best, i.e. find those who have faced a similar problem and found a successful solution

Problem 2: An entrepreneur wants to open a gift store in a newly gentrified neighborhood. In order to prepare for his entry he samples successful gift stores in similarly situated neighborhoods and copies the format and goods sold there. Again his strategy is imitate the best or the successful.

While imitating those who have been successful certainly appears to make sense on the surface, as a strategy it fails to appreciate one fact which is that by imitating the successful you may actually be imitating the foolhardy but lucky. For example, if we were to rank people by their degree of success then we would expect that on the top of the list would be those people who took big gambles and were lucky. On the bottom of the list might be others who made the same choices but were unlucky. Those above the middle but below the top of the list are very likely the ones who chose optimally in the sense of making that choice which was ex ante optimal given the chances of success. So if one would sample intelligently and had limited resources with which to do so, sampling and imitating just above the middle of the list would be optimal since that is where most of the relevant information can be found. While few of us would think it is smart to imitate failures and sample on the bottom of the list, those observations may also be informative since at least we would learn that those same risky choices that paid off well for some lead others to disaster.

In evolutionary terms the problem may be more severe since firms that make bad decisions go bankrupt and hence are not available to be learned from or imitated. This leaves the set of extant firms to be mostly composed of those who made risky choices and were lucky. Imitating the behavior of these firms exposes the imitator to a high variance distribution of returns whose ex ante mean is below those who chose more conservatively, yet the ex post realized profits are the best. The best rule is copy those whose decisions are ex ante optimal and not those who are ex post lucky. The problem is how to find out who they are.

So when our CEO instructs his committee as to what to do he may pursue either of two strategies. One is to sample other firms for information about the decision under investigation and make one's own choice based on what it learned. This strategy is in line with the paradigm of expected utility maximization and we refer to it as "sampling for information".

Gilboa and Schmeidler (1995; 2001) provide an alternative paradigm to model decision making under uncertainty. According to their case-based decision theory, people make decisions by analogies to past cases. Decision makers are inclined to make choices that performed well in previous similar situations and to avoid choices that fared badly. Gilboa and Schmeidler suggest that case-based decision theory provides an accurate description of the way decisions are made when the decision maker faces an unfamiliar problem, such as whether to start a war, whether to invest in a politically unstable country and whether to get married.[1] In such circumstances, decision makers may search their memory for past cases that are similar to theirs.[2] Each case is weighted by its similarity to the current problem, and the decision maker chooses the act that had the highest (average) past performance. The strategy in line with case-based decision theory is to sample and imitate the most successful predecessors. If our CEO employs this strategy, we say that he is "sampling for imitation".

The consequences of imitation myopia may be far ranging. First, it is no surprise that a substantial majority of all new business fail if entrepreneurs insist on only sampling those businesses in the population who chose risky plans and were lucky.[3] That percentage might possibly be cut dramatically if business owners sampled for information in a more intelligent way. Put differently, if people imitate the successful but fail to realize that those are exactly the lucky in society, then those decision makers are suffering from a type of winner's curse in their inability to adjust their behavior for the fact that they are sampling only the highest order statistics of success.[4]

Second, it is often observed that people are schizophrenic in their relationship to risk. While some are apparently risk averse in one realm of their life (for instance, when they buy insurance against bike theft), they may appear to be highly risk seeking in others (for instance, when they decide about their stock portfolio). Our paper offers an explanation for this which we test experi-

[1] Biological experiments show that females often copy the mate choices of other females. Dugatkin and Godin (1992) offer female guppies the opportunity to express a preference between two males. Then the female observes a second female displaying a preference for the male she herself did not prefer. When given a second opportunity to select between the same two males, the females reverse their mate choices significantly more often than the females in the control group who do not observe the mate choices of other guppies. Likewise, female sticklebacks have a preference to spawn with males whose nests contain eggs (Ridley and Rechten, 1981). There is also evidence for female copying of mate choices amongst lekking birds and mammals (Gibson and Hoglund, 1992).

[2] In addition, Gilboa and Schmeidler argue that cased-based decision theory is plausible in situations where the decision maker faces the same decision problem frequently enough, such as whether to stop at a red traffic light. In such cases, decisions become almost automated. Expected utility theory then covers the middle ground between the two extremes of repetitive and unfamiliar problems.

[3] Camerer (1997) notes that 80% of all new business fail in their first three years.

[4] There are, of course, other pitfalls when decision makers evaluate the performance of decision makers who did the same problem before them. For instance, Rabin (2002) gives the example where a decision maker observes a list of several performances of each person in the sample. A believer in the "law of small numbers" will conclude that some decision makers are superior and others are inferior, even in a situation where a Bayesian decision maker would eventually figure out that there are no differences in performance. In agreement with this example, Offerman and Sonnemans (2004) show in an experiment that people tend to believe that series that are actually produced by fair coins are instead produced by false autocorrelated coins.

mentally. The explanation is simple. If people imitate success and if those who are successful are exactly those who have made the most risky choices and were lucky with them, then imitation is very likely to lead to what appears to be a population of risk seekers despite the fact that these same people exhibit a large degree of risk aversion when tested or in other contexts where there is no possibility to sample others. Imitation leads them to act as risk seekers since it masks the riskiness of the choices they are following.

Finally, as stated above, our results have evolutionary consequences. If imitators copy the successful and only those that have taken big risks are the successful ones, then sooner or later those choosing optimally will fail to exist and hence fail to be available for imitation. Those that remain will look exceedingly risk seeking and we can expect to continue to observe a large fraction of businesses failing since only high variance businesses will be imitated.

In this paper, we report the results of an experiment designed to investigate how people sample for information before making a decision. Our design allows us to cleanly separate those who are imitators from those who are learners. Overwhelmingly we find that people fail to take advantage of the opportunity to learn in preference for imitation despite the fact that it does not take a rocket scientist to sample for information in our experiment. More importantly, our experimental results exhibit many of the perverse results mentioned above. First, we find that when subjects have the opportunity to sample the results of other subjects, they tend to make *less* money at a higher variance. Second, we find that imitation leads people to act as if they are more risk seeking than they are. We infer this by measuring subjects' risk aversion using the Holt-Laury (2002) procedure after they perform the experiment and then by inferring it from the decisions they make after sampling. We find that clearly after sampling people behave as if they have a coefficient of risk aversion far smaller than that implied by the Holt-Laury scale. This leads us to conclude that sampling leads to as-if risk seeking behavior. Finally it appears as if our subjects learn relatively little from being given the chance to observe the behavior of those that proceeded them. We support this claim by measuring the beliefs of our subjects about the main uncertain variable in the problem, the price of the good they will "produce" after they have sampled what others have done. We find their posteriors to be particularly diffuse and not significantly different from the uniform priors we induced in them.

Our paper contributes to the emerging literature on imitation. One branch of the theoretical and experimental literature investigates the role of imitation in oligopoly games (Vega-Redondo, 1997; Huck, Normann and Oechssler, 1999; Selten and Ostmann, 2001; Offerman, Potters and Sonnemans, 2002; Selten and Apesteguia, 2004; Abbink and Brandts, 2004; Apesteguia, Huck and Oechssler, 2004). This literature shows that with the right kind of information feedback, the industry moves into the direction of the Walrasian outcome. The crucial aspect that differentiates our paper from this work is that we investigate the role of imitation when an idiosyncratic luck shock affects the success of decision makers. Our paper is closer to the theoretical work by Ellison and Fudenberg (1995) and Schlag (1998; 1999) who study what happens when people imitate while an idiosyncratic term affects their payoffs. Ellison and Fudenberg show that word of mouth communication may lead to more efficient outcomes when each agent samples only a few others. In their model, it is assumed that each player hears of the experiences of a random sample of N other players. The fraction of players who listen to what they hear pick the action that gave the highest average payoff (those who do not listen do not change their choices). Schlag (1998) considers a situation where people choose between actions yielding uncertain payoffs (the multi-armed bandit problem). Schlag allows people to obtain a random sample of one other person. He shows that the rule where an individual imitates the action of the observed individual with a probability proportional to the difference of the other's payoff and the own payoff in the previous round outperforms all other learning rules with limited memory. Schlag (1999) extends the analysis to the situation where each person receives information about a random sample of two others. The major difference between the setup in these papers and our setup is that we endogenize the sampling phase. Instead of presenting the player with the information of a random sample, we let players decide for themselves who they want to sample.

This paper makes a contribution to the field of social learning since it deals with a problem that has not been dealt with before. More precisely, in the typical social learning experiment when it is a person's turn to act she has access to either all the decisions that have gone before her or at least a subset. In other words, the information available to a person is exogenous and all the decision maker needs to do is to incorporate this information into her prior and make a decision. Our experiment combines elements of search theory with social learning since our subjects must decide from whom they want information and then sample them. In this sense it adds a new dimension to the social learning problem [5]

In this paper we will proceed as follows. In Section 2 we will describe the problem faced by our experimental subjects. The experimental implementation of this problem and our design will be described in Section 3. Section 4 presents our results while Section 5 describes offers some extensions. Section 6 concludes.

2 The Problem

Consider the following decision-theoretic problem. A firm has to decide how much it wants to produce of a product it sells. Assume that the price is uncertain and uniformly distributed between $[\underline{p}, \bar{p}]$. (In the experiment $\underline{p} = 10$ and $\bar{p} =$ 90). The firm has two options. Given its location it can decide to limit its sales to the market that is local to its business, i.e. only produce in the state where its factories are, or it can produce nationally. Producing for the local market differs from producing nationally in two ways. First, the local market is smaller and hence the amount produced, q, is constrained such that q $\in$ [$\underline{q}$,q"] where

[5]The only other paper we know of in the social learning literature that makes the information structure endogenous is the work by Celen, Choi and Hyndman (2005).

$\underline{q} \leq$ q" $\leq \bar{q}$ where $\underline{q}$ and $\bar{q}$ are the lower and upper limits on production. (In the experiment q" = 40 with so production levels of 40 or below were in the local market, while $\underline{q}$ = 10 and $\bar{q}$ = 90.) Second, because the firm knows the local market it can easily judge what its production costs will be so there is no uncertainty there. If the firm decides to produce nationally, then it can choose to produce an amount in the interval (q", $\bar{q}$] but it faces a stochastic cost of production. More precisely let the profit of the firm be,

$\pi_l = 2 \cdot p \cdot q - c(q)$ if q $\leq$ q" (i.e. if the firm produces for the local market), and

$\pi_n = 2 \cdot p \cdot q - c(q)((1 + 0.01 \cdot \varepsilon)$ if q > q" (i.e. if the firm produces nationally),

where ε is a random variable that is uniformly distributed on the interval [-60,60] and

c(q) = q^2.

Note that the price faced by the firm will be the same whether it is sold in the local or the national market. However, costs are stochastic if one sells in the national market where production levels are greater than q".

Given the assumed functional forms for the distribution of prices and costs, expected profits can be written as,

E(π) = $2 \cdot E(p) \cdot q - q^2((1 + 0.01 \cdot E(\varepsilon))$. The first order conditions show that $2 \cdot E(p)$ = 2q or q = 50. So without any information about price or cost shocks the optimal risk neutral choice is q = 50. If price were known, then the optimal risk neutral price setting rule will be q=p.

Now consider that this problem has been faced by a set of 60 firms in the past who vary in their risk attitudes (and perhaps also in their cognitive skills) and therefore have made choices that are distributed over the interval [10,90]. Some will choose high q's and get good realizations while others will choose high q's and get bad and negative realizations. Others will choose q's with middling or low values. Finally, assume that unbeknownst to the firms the actual realized price, p^R is p^R = 38. In other words, while the firms only know that the price is drawn uniformly from the interval [10,90], we, as outside observers, know its realized value is 38.

Given these assumptions, if we were to rank firms by profits and could see what they did we would see that those on the top of the list would be the high-q low-ε firms while those on the bottom would be the high-q high-ε firms. In other words, those on the top and the bottom would be those firms that chose to produce for the national market (i.e. chose high q's) and were either lucky or unlucky. Those with middling levels of profits would be those who chose in local market where uncertainty about costs don't matter. So the question is if a decision maker were given the rankings of those who have performed this problem before him or her and was told nothing about either their output choices or their profits but was told whether they sold in the local or national market, and if such a decision maker was allowed to sample either once or three times before making his or her output choice, (i.e. sampling means seeing the output choice and profit of the firm sampled) what would be the best place to sample and what would be the best output given the information received?

The answer to this question is clear. It is optimal to sample a firm (any firm) that produced in the local market, find out the output and profit of this firm, invert the deterministic profit function for the implied price and then set the quantity equal to that price. So, the optimal risk-neutral q is then 38 since that is the realized price in the example above, and in the experiment below. It is important to note that the optimal sampling procedure is one where you sample for information and then use that information to set your output optimally. An imitator might behave differently and sample the firm that received the highest profit and copy its output. So sampling for information and sampling for imitation are two very different things, imply different sampling procedures, and different ultimate outputs.

3 The Experiment and the Experimental Design

The experiment performed was a fairly straightforward implementation of the problem described in Section 2 above. All experiments were performed at the Experimental Lab of the Center for Experimental Social Science (CESS) at New York University. All subjects were undergraduates at New York University and were recruited by E-mail and signed up on a first-come first-served basis. The experiment lasted almost one hour and average payoffs were $16. At the end of each experiment, but before they received the information about their own earnings for the production decision, all subjects were asked to take the Holt-Laury (2002) procedure for eliciting their risk attitude (see the appendix for details).

All instructions were presented to the subjects on computer terminals. The procedures were then reviewed and questions answered. In total there were three treatments called the Seed, Sample-One and Sample-Three. In all treatments subjects earned points that were exchanged at the end of the experiment at a rate of 500 points to $1. We provided each subject with a starting capital of 5000 points from which we deducted points in case a subject made a loss. No subject actually left the experiment with a loss.

The Seed was an experiment run to provide the ranked list of subjects that later subjects in the Sample-One and Sample-Three treatments would use for sampling. In this experiment subjects simply came into the lab and were presented with the problem described above and asked to choose an output level. They did not have an opportunity to sample an output-profit pair of any firm. The problem was described to them using the terms "firm" and "production level" and they were presented with the profit formula. Their costs of production were presented to them in table form from which it was obvious that the cost of production level q was q^2.

In the Sample-One treatment subjects would come into the lab and read the instructions describing the problem but would be told that 60 in the Seed had done exactly the same problem before them. They would then be presented on their computer screens with a list of those people ranked from the one who did best to the one who did worst and next to each subject in the Seed was an indication of whether they chose an output level in the national or local market (we placed "Nat" or "Loc" next to their names). In order to find out their production level and profit the subject had to click on the rank of the person they wanted to inspect and when they did they would see what production level this person had chosen and what her profit was. After doing so, they would be asked to choose a production level or output for themselves.

The Sample-Three treatment was identical to the Sample-One except that here the subjects could sample three times before choosing their production level. Actually the subjects were not told that they had exactly three sampling opportunities. Rather they were told that the number of times they could sample was not revealed to them. We did this so as to get some insight into what they thought was the best place to sample in their first, second and third samples. The idea here was that if they knew they could sample three times, there would be no premium on sampling that person they thought was best to sample first and then continue in order of priority since knowing that they had three samples they could sample in any order and simply make their decision after they collected all their data.[6] With the uncertainty, however, each sampling should be their best expected sample conditional on the information they had gathered before. So whom they sampled first, second and third should be revealing. In all treatments except the Seed subjects did the experiment once and only once. In the Seed we added one extra task which we will describe below.

The above treatments were embellished in several ways. First of all in the Seed we ran the experiment in two stages. Stage 1 was described above. After stage 1 was over we told the subjects that there would be a stage 2 in which the price of the good was given to them. They then had to choose a production level knowing this. This part of the experiment would be the proper benchmark for the Sampling treatments if subjects were able to deduce the price level through sampling. The price shown to every subject was 38 so 38 was the optimal risk neutral choice for all subjects. (**In all treatments the realized price was 38**).

In the Sample-One and Sample-Three treatments we also added an extra task after they had finished sampling and choosing their production level. The details of this task were not mentioned to them before they completed choosing a production level. In this part of the experiment, we wanted to find out what subjects had learned by their search. We did so by asking them, in light of their sampling, to report what they believed the price for the product was in the investment decision experiment.[7] Subjects were asked to fill in numbers in 8 boxes on the bottom of their computer screen indicating what they thought the probability was that the price in the production problem fell into 8 different intervals 10-19, 20-29, 30-39, 40-49, 50-59, 60-69, 70-79, and 80-90. They had

[6] An alternative design would have been to impose a cost for each search.

[7] When they reported their beliefs, the sampling results were still listed on the screen so they had all the information available to them that they accumulated.

to allocate 100 probability points across these intervals. We rewarded them for their beliefs by a payment generated by a quadratic scoring rule The quadratic scoring rule is an incentive compatible mechanism, i.e., it induces subjects who want to maximize their expected payment to report their beliefs truthfully The appendix lists the details of how we made use of the quadratic scoring rule besides a general overview of the instructions. Subjects completed this task before they received the information about their own profit.

The exact experimental design is given in Table 1.

[Table 1 here]

3.1 Hypotheses

Our discussion leads to a number of hypotheses which we will test in Section 4 below.

Our first set of hypotheses concern behavior in the Seed treatments. To begin, since in stage 1 of the Seed treatment prices are assumed to be distributed uniformly over the interval [10,90] and since subjects in the Seed can not sample for information, we expect risk neutral subjects to choose a production level of 50 which is equal to the mean of the distribution. However, in stage 2, after they are informed that the price is 38, we would expect them to choose 38. These expectations furnish us with the following hypotheses.

Hypothesis 1: Seed Behavior

Subjects in the stage 1 of the Seed treatment choose a production level of 50.

Hypothesis 2: Seed with Price Behavior

Subjects who receive the price in stage 2 of the Seed treatment choose a production level of 38.

Notice that even if subjects are not risk neutral, we would expect to observe lower production levels in stage 2 than in stage 1. In stage 2 subjects know that the price is lower than the expected price of the prior distribution which would lead expected utility maximizing subjects to choose lower production levels.

Our next set of hypotheses concern sampling behavior. In the experiment, optimal sampling amounts to the following. Since there is no random cost elements in the local market, observing the production level and profit of a subject choosing locally in the Seed allows a subject to solve for the realized price and then set his or her production accordingly. In line with optimal sampling, we posit Hypothesis 3 which states that the (risk neutral) production level in the Sample-One treatment should be 38. Note that sampling three times in our Sample-Three treatment offers no new information. The price can be inferred exactly if one were simply to sample once and do so in the local market. This leads us to posit Hypothesis 4 which states that production levels should be the same in the Sample-One and Sample-Three experiments.

Hypothesis 3: Sample-One Behavior - Sampling for Information versus Sampling for Imitation

Subjects in the Sample-One treatment sample a subject (any subject) in the local market, and set a production level of 38.

Hypothesis 4: Sample-Three Behavior;

The production level set by subjects in the Sample-Three experiment is not different from that set by subjects in the Sample-One experiment.

In contrast, subjects who sample for imitation will search people at the top of the list and choose higher production levels. It is interesting to compare Sample-One behavior with Sample-Three behavior, because such a comparison reveals whether potential biases observed in Sample-One are robust. It is possible that with multiple searches subjects sample broadly over the list and find out that imitating the top is risky and not necessarily optimal.

The next hypothesis is a very important one. In the stage-1 Seed treatment the ex ante optimal choice is 50 since no price information beyond the prior information is available. In the sampling treatments, if subjects sample for information, they can find out the price. So we expect rational risk neutral subjects to choose 38. However, if they sample for imitation then they would sample the person in the Seed who got the highest profits level (who happened to choose 63) and copy him. Hence, if we see significantly higher levels of production in the Sample treatments than we do in the Seed, then we know that subjects sample for imitation and not for information.

Hypothesis 5: Seed-Sample Comparisons

Subjects in both the Sample-One or the Sample-Three treatments set lower production levels than those in the stage-1 Seed treatment. (The null hypothesis here is that there is no difference in the production levels).

Our next hypothesis concerns itself with risk taking behavior and the impact of imitation on it. The idea here is that if subjects imitate others and the ones they imitate are lucky risk takers, they too will appear to be similar in type despite the fact that their underlying risk preference is different. To test this hypothesis we will use the fact that we test each subject for their level of risk aversion using the Holt-Laury test after the experiment and therefore can compare this measures with the one implied with the production choice after sampling.

Hypothesis 6: Risk Taking

Subjects in the Sample-One and Sample-Three treatments exhibit a lower degree of risk aversion implied by the production choice than their degree measured in the Holt-Laury test. (Null hypothesis is that there is no difference).

As economists we are always interested in the welfare implications of what we do. Therefore, we ask whether the ability to imitate or sample is welfare increasing or not. On the surface, if subjects sample for information it should increase welfare (as measured by subject payoffs) since only in the sampling treatments does the subject have a chance to find out the price. If they sample for information and choose 38, they earn a higher expected profit than the subjects in the Seed who choose the ex ante optimal choice of 50. If they imitate however, and choose higher, profits will fall. This is a classic example of how more information can be damaging if we compare the profits obtained in the Seed treatment with those of the Sampling-One and Sampling-Three treatments.

Hypothesis 7: Welfare

Payoff levels for subjects are higher in the sampling treatments than the Seed. (Null hypothesis is that there is no difference).

Our final hypothesis concerns learning. We evaluate learning by the shape of the distributions of beliefs elicited from subjects after their sampling is done. Note that full learning is available to subjects in our experiment if they sample locally so complete learning would imply that subjects have a degenerate belief distribution after sampling placing their full mass on the interval that contains the true price. Even if they do not learn in exactly this way we can still ask if they significantly change their priors as to what the realized price is in the experiment and answer this question by looking at their reported posterior price distributions elicited after their information search and comparing it to the induced uniform prior.

Hypothesis 8: Learning

After either sampling once or three times subjects learn the true realized price in the experiment.

4 Results

We will present the results of our experiment by first looking at the behavior of subjects in the Seed and then move on to the more important question of sampling behavior.

4.1 The Seed

Tables 2A and 2B and Figures 1-2 present the results of the choices of subjects in the Seed treatment:

Tables 2A and 2B and Figures 1-2 here

There are several interesting things about subject behavior in the Seed. First, as can be seen in the first four columns of Tables 2A and 2B, which show the rank of each person in the Seed, the market he or she produced in, the production level chosen and the profit made, after 60 subjects made their choices in the Seed Treatment, the people who chose nationally were ranked simultaneously on the top and the bottom. More precisely, the top 11 subjects chose nationally as did the bottom 9. The top 3 choices were production levels of 63, 90 and 60 while the bottom three production levels were 67, 70, and 75. The mean production level chosen in stage 1 of the Seed was 42.6. Figure 1 shows that the modal choice of subjects was to select the output level which was the highest in the local market, i.e. 40 with 13 subjects choosing that. The lowest production level chosen was 15. On the basis of this data we can reject

Hypothesis 1 that the median choice made was 50 using a sign test (p=0.00) since only 14 choices were strictly higher than 50 while 42 choices were below.

In stage 2 of the Seed treatment, after they were told the price was 38, production levels actually went up to a mean of 47.0. This mean was not significantly different from the stage-1 mean of 42.6 using a Wilcoxon test (p = 0.25). Again, the median choice here was 40, very close to 38. Still, a sign test (p=0.00) rejects hypothesis 2 in favor of the hypothesis that the price in the Seed treatment is higher than 38.

Figure 2 shows that in stage 2 of the Seed 7 subjects chose 38 (compared to none in stage 1). Some subjects decreased their production level and selected the risk neutral optimal choice once they have learned that the price is 38. At the same time, some other subjects revealed in the post-experimental questionnaire, that once they had learned that the price was 38, they chose a higher production level than before, because now they were certain that they could afford a higher production level without being exposed to the danger of running into a loss (which might occur if the price were below 38). It seems that these opposing forces have offset each other on average.

Later we will see that subjects did not learn the true price when they sampled people who did the problem before them. This means that stage 1 behavior of the Seed is the natural benchmark for the Sampling treatments and in the remaining part of the paper we will compare the Sampling treatments with stage 1 of the Seed. It is reassuring, however, that the distribution of choices is the same with or without information about the price, so that empirically it does not matter which one is used as the benchmark. What the Seed data show is that if left to their own devises, subjects make rather conservative choices (mostly less than 50). As we will see later, being able to sample makes them more adventurous.

In summary, the Seed treatment produced the results we expected. Those subjects who chose high production levels (risk seekers) made simultaneously the most and the least profits. Those who chose ex ante optimally, 50 (or approximately 50), made less extreme amounts of money but more on average. In the experiment we now ranked these subjects, without revealing their profits or choices, and let subjects in the Sample-One and Sample-Three treatments sample. We now turn our attention to those treatments.

4.2 The Sample-One and Sample-Three Treatments: Choice Behavior

Table 3 and Figures 3 and 4 present the results of the Sampling treatments.

Table 3 and Figures 3, 4 and 5

In the Sample-One and Sample-Three treatments we have 32 and 25 subjects respectively. Figures 3 and 4 present the histograms of production levels for subjects in the Sample-One and Sample-Three treatments. Notice that allowing subjects to sample dramatically leads them to choose high production levels. In terms of Hypothesis 3, it should be clear that we can reject the hypothesis that the median choice of subjects in the Sample-One and Sample-Three treatments was equal to 38 since the median choice in both the Sample-One and Sample-Three treatments is 60 (compared to a median of 40 in the Stage-1 Seed).

In contrast to Hypothesis 5, Figure 5, which presents a "smoothed" histogram of production levels, shows that the distribution of production levels chosen in the Seed is to the left of the distribution of production levels chosen in the pooled Sampling treatments. Table 3 shows that the production levels in the Seed rose from 42.6 to 54.5 and 55.3 for the Sample-One and Sample-Three treatments, respectively. Mann-Whitney rank tests reveal that the Sample-One as well as the Sample-Three treatment produce significantly higher production levels than the Seed treatment. In accordance with Hypothesis 4, the difference in the production level of Sample-One and Sample-Three is far from significant.[8] The effect of sampling is also apparent in the market that subjects use for their production. While only 38.3% of subjects invested in the national market in the Seed experiment, for the Sample-One and Sample-Three treatments these same percentages are 71.9% and 76.0%, respectively.

From observing production choices it appears that we can reject the hypothesis that subjects sample for information since if they did then we would find production levels in the Sample-One and Sample-Three treatments to be below those of the Seed but we find just the opposite. So, it would appear that either subjects did not sample for information or at least if they did, they did not choose the appropriate production level that corresponds to a true price of 38. As we will see below, the answer is the former – people do not sample for information but rather imitate the successful.

4.3 The Sample-One and Sample-Three Treatments: Sampling Behavior

The main focus of this paper is on sampling behavior. While we have no preconceived ideas as to what would happen in the experiment we did feel that imitative behavior is so ingrained in us that a convincing conjecture would be that subjects in the Sample-One and Sample-Three treatments would tend to not sample in the local markets but rather at the top of the profit rankings where people produced nationally.

The right four columns of Tables 2A and 2B present the percentages of time that subjects sampled in the local and national markets. Remember that if they were sampling for information they would only sample in the local market. The results here are striking. For example, in the Sample-One treatment 75.6.% of the samples were in the national market with 83.6% of those samples being samples of the top ranked person in the Seed. Overall, 63.2% of subjects sampled the top ranked person with an additional 5.3% sampling the subject ranked

[8] Comparing the treatments with a series of Kolmogorov-Smirnov tests leads to the same conclusion: the production levels of both the Sample-One and the Sample-Three treatments are significantly different from those of the Seed at the 1% level, while they are not significantly different from each other.

either second or third. The second most frequent sampling pattern is for subject to sample that subject who was ranked highest amongst subjects who chose locally. 8.8% of the subjects did this. A binomial test rejects the hypothesis that sampling behavior was random between the local and national market (binomial probability of $p = 0.5$) at the 1% level in favor of the hypothesis that sampling the national market is more popular than sampling the local market. People are clearly biased toward sampling in the national market.

The same behavior carried over to the Sample-Three treatment. Here 76.0% of all first samples were on people who chose nationally in the Seed with 94.7% of those being on the person who received the highest rank. Overwhelmingly the person sampled first for subjects with three sampling opportunities is the top-ranked subject in the Seed. On the second sample the most favorite person to sample is the person who produced locally and who received the highest profit. 36% of subjects sampled here (this person chose a production level of 40).The second most popular person to sample on the second sample opportunity was the lowest ranked Seed subject. 28% of subjects did this. On the third sample people concentrated most on the second-ranked subject (who produced nationally). Over the entire set of three samples 66.7% of those sampled chose to produce nationally so overwhelmingly the information gathered was from those who had produced at the national level and hence chose high production levels.

Given their sampling pattern, we can ask how subjects transformed what they learned during their search into a production level choice for themselves. While we will soon capture this process in a regression, let us first look at some descriptive statistics.

[Table 4 here]

As we can see from Table 4, which presents the sampling and production behavior for subjects in the Sample-One and Sample-Three treatments, in the Sample-One (Sample-Three) treatments 75% (76%) of subjects sampled a subject who produced for the national market. (For the Sample-Three treatment we are looking only at the first sample). Of those who sampled in the national market in the Sample-One (Sample-Three) treatments 87.5% (84.2%) ultimately produced in the national market themselves. In addition, where subjects sampled had a great impact on what they ultimately chose to produce. For example, when a subject sampled the national market in the Sample-One (or first sample in the Sample-Three) treatment they chose an average production level of 61.6 (59.0) while those sampling in the local market chose 33.4 (43.7). So subjects' sampling behavior has consequences for their production levels.

To give some structure to this behavior we ran the following regression:

$\text{prod} = c + \beta_1 \cdot \text{prod}_1 + \beta_2 * d_1 * \text{prod}_1 + \beta_3 * d_2(\text{prod}_2 - \text{prod}_1) + \beta_4 * d_3 * (\text{prod}_3 - \text{prod}_1) + \beta_5 * d_4 * (\text{prod}_2 - \text{prod}_1) + \beta_6 * d_5 * (\text{prod}_3 - \text{prod}_1) + \varepsilon$

where

prod is the production level set by a given individual

prod_i = production level observed on the i-th sample.

d_1 =1 if (treatment = Sample-Three) and $d_1 = 0$ if (treatment = Sample-One)

$d_2 = 1$ if (treatment = Sample-Three and profit second sample > profit first sample) and $d_2 = 0$ otherwise

$d_3 = 1$ if (treatment = Sample-Three and profit third sample > profit first sample) and $d_3 = 0$ otherwise.

$d_4 = 1$ if (treatment = Sample-Three and profit second sample < profit first sample) and $d_4 = 0$ otherwise.

$d_5 = 1$ if (treatment = Sample-Three and profit third sample < profit first sample) and $d_5 = 0$ otherwise.

and

ε is a random disturbance term with mean zero.

The regression results are summarized in Table 5:

[Table 5 here]

This regression suggests that what subjects see on their first sample, whether that is the first and only sample or the first of three samples, is the key determinant of production choice. For example, subjects tend to choose production levels which are 93% of those they observe on their first sample. This is true for both the Sample-One and Sample-Three experiments. (Notice that β_2 is not significantly different from zero). Whatever they see on their second and third samples, if they have any, does not influence their choice in a significant manner.[9]

These results are indicative that subjects are sampling for imitation rather than using their sample for its information content. As noted before, they tend to sample the top person in the national market (whose production level was 63) and virtually copy what he or she did. Little that they find out in their second or third samples changes this.

4.4 Risk Taking

In our experiment, we test all subjects for their level of risk aversion using the Holt-Laury (2002) procedure and compare it to the level inferred from their production level. If we posit that subjects maximize a CRRA utility function of the form $U(x) = \frac{x^{1-r}}{(1-r)}$ if $x > 0$ and $U(x) = \frac{-(-x)^{1-r}}{(1-r)}$ if $x < 0$, then we can calculate which r would rationalize the production level chosen. We call this level r_{prod}. Notice that we need a utility function that handles negative amounts, because to compute expected payoffs we need to integrate over all possible outcomes including negative ones. The utility function proposed above provides a natural way to work with negative payoffs (Wakker, 2005). There

[9] Interestingly, subjects who first sample the top and then sample the bottom on their second or third sample are not scared away from high production levels: the 7 subjects who sample top and bottom produce 60.0 (std. dev.16.6), while the 11 subjects who sample the top but not bottom produce 55.6 (std. dev. 16.6). The difference in prioduction levels between these two groups are not significant (Mann-Whitney test: p=0.49).

is an upperbound on the concavity of the utility function, r<1, to ensure that utility is increasing.[10]

The Holt-Laury procedure asks subjects to make 10 choices between 2 risky lotteries. The Table in the appendix lists the choices. The choices are constructed such that the crossover point for switching from risk averse lotteries A to risk seeking lotteries B provides an interval of the estimate of a subject's relative risk aversion coefficient. We set the estimate for a subject equal to the middle of this interval and denote it by r_{HL}. The majority of subjects started with choosing A , switched to B at some choice and then never returned to the A choices, like it is expected from someone maximizing expected utility. For subjects who switched back to A choices, we used the total number of A choices as a measure of the subject's risk aversion (similar to Holt and Laury). 14.7% of the subjects (17 out of 116) switched back at least once.[11] 3 of these subjects switched back 3 or more times and we drop these people from the analysis because we felt that these people made more or less random choices in the lottery procedure. (This does not affect the analysis in an important way).

The Holt-Laury procedure only deals with positive payoffs. Therefore, the part of the utility function that deals with negative payoffs and the restriction r<1 are not needed if one only wants to explain behavior in their problem. In particular, if subjects make 8 or more safe choices before they switch to risky choices then the implied r is larger than 1. Given our constraint r<1, needed to infer the risk aversion coefficient from the production decision where payoffs may be negative, we chose to set r_{HL} equal to 0.99 if subjects made 8 or more safe choices. We had to downgrade the risk aversion levels of 6 subject in this way. This does not affect the main result of this section which is that sampling leads subjects to behave as if they were more risk seeking. In fact, if anything, this choice made it only harder to show that the risk coefficient inferred from the production decision is smaller than the risk coefficient resulting from the Holt-Laury procedure.

Table 6 presents an overview of the r_{HL} and r_{prod} coefficients that we derived for our subjects Figures 6-9 present cumulative density functions of r_{HL} and r_{prod} broken down by treatment.

[**Table 6 and Figures 6-9 here**]

Figure 6 shows that the distribution of r_{HL} does not differ for subjects who engaged in different treatments. This is as it should be since it only indicates that subjects doing our various experiments were drawn from the same population. We have no a priori reason to suspect otherwise. To bolster this claim we tested the hypothesis that the distribution of r_{HL} measured by the Holt-Laury procedure did not differ depending on which treatment a subject performed and

[10] We infer the risk attitude from the production level assuming that a subject's belief about the price is represented by the induced uniform distribution. In section 4.6, we will present evidence showing that in Sample-Three subjects' reported beliefs by and large coincide with the uniform distribution.

[11] The risk aversion data of 1 subject were lost (this was the only part of the experiment that was run by hand).

found that we could not reject it at the 20% level for any binary comparison (Kolmogorov-Smirnov tests). What is more interesting is that we can also not reject the hypothesis that the r_{prod}'s inferred by subject decisions in the Seed treatment were the same as those measured by the Holt-Laury procedure (see Table 6). In other words, if we infer what r_{prod} coefficients are from observing production decisions in the Seed experiment and then compare them to what we measured subjects' r_{HL} coefficients to be in the Holt-Laury procedure, then there appears to be nothing obvious in the task offered subjects in the Seed treatment that causes them to behave as if they have a different r_{prod} from the one we tested them as having. Figure 7 confirms that the Seed treatment appears not to have affected subjects' revealed levels of risk aversion.

The same can not be said for a comparison of r_{prod} and r_{HL} for the Sample-One and Sample-Three experiments. Here, the distributions of r_{prod} implied by production choices reveal that engaging in these experiments leads people to behave as if they were more risk seeking. Figures 8 and 9 show the shift of the cumulative density functions to the left in both cases when compared to the r_{HL}'s for the same people revealed through the Holt-Laury procedure. Table 6 presents the test results that are significant for both comparisons. This should not be a surprise since if we look at the actual production decisions made in the Seed treatment we see that over 61% of them were in the local market, representing a risk averse choice, while in the Sample-One and Sample-Three experiments only 28% and 24% of the choices were in the local markets, respectively. This is a huge shift in risk taking behavior seen when subjects are given a chance to sample. The evidence is clearly in favor of Hypothesis 6.

4.5 Welfare

As economists we ultimately care about welfare. In this experiment that translates into a question of whether allowing subjects to sample and imitate others increases their welfare (as measured by their payoff in the experiment) over what they would achieve if we simply asked them to choose without any additional information as we did in the Seed.

Figure 10 presents the cumulative frequency distribution of payoffs in the of our subjects in the Seed and all treatments where they were allowed to sample (i.e., Sample-One and Sample-Three).

[**Figure 10 here**]

As we can see, subjects in the sampling treatments were more exposed to the danger of making losses than subjects in the Seed. In fact, payoffs decrease when sampling is allowed. While the mean payoff for subjects in the Seed was 1289.58 (std. dev. 852.38) it was only 798.32 (std. dev. 1485.04), in the combined sampling treatments, so subjects earned less profit at a higher variance when they sampled. A two-sided Kolmogorov-Smirnov test rejects the hypothesis that the distributions are equal at p=0.07. (Separated across the Sample-One and Sample-Three treatments subjects earn 776.16 and 826.68 respectively). This is a good example of where more information can be a bad thing. Hypothesis 7 is rejected, since subjects in the Sampling treatments earn less money instead of more money compared to the subjects in the Seed.

4.6 Beliefs

The evidence reported in the preceding sections clearly suggests that people used the possibility to sample for imitative rather than informational purposes. The most direct evidence on the question whether people learn is provided by the beliefs that subjects reported in the Sampling treatments after they had made their production decision. If people do not learn anything from their sampling about the price, their posterior distribution should coincide with the uniform prior distribution on [10,90]. If on the other hand, people use their sampling to recover the underlying price, they would assign 100% probability to the interval [30,39] containing the true price 38. Table 7 and Figure 11 present the results on subjects' beliefs.

[Table 7 and Figure 11 here]

Table 7 shows that on average subjects spread out their probabilities over the whole range of prices. In fact, the average reported distributions are practically indistinguishable from the induced uniform prior distribution. This is true when the data are split across the Sample-One and the Sample-Three treatments, and it remains true if the data are split across people whose first (or only) sample was in the local market and people whose first sample is in the national market. The latter finding is interesting because it shows that people who sample the local market do not do so to recover the true price. Instead, it seems likely that they sample in the local market because they wanted to imitate the most successful person in that market. (Recall that the samples in the local market are focussed on the most successful person of the local market).

Table 7 reports averages only and it masks some differences in subjects' reported beliefs. To get an idea about the variation in beliefs, Figure 11 shows the histogram of the probabilities assigned to the correct price interval [30-39]. Most subjects assign a probability between 0% and 20%, close to the probability of 12.5% induced by the uniform prior distribution. The maximum probability assigned to the true price interval is 50%, chosen by three subjects. One of these three subjects revealed in the questionnaire that she had exactly derived the true price of 38. (Surprisingly, she only assigned 50% to the true interval, as if she was still not completely sure).

It is clear, however, that the overwhelming majority did not use the sampling possibility to learn about the true price. Therefore, we confidently reject Hypothesis 8.

5 Extensions

Our experiment may have implications that extend beyond the circumstances described here. More precisely, there may be many other contexts where imitation can have similar welfare decreasing and risk seeking consequences. For example, consider markets with asymmetric information such as the Samuelson-Bazerman (1985) takeover game. In this game two firms exist called the target and the bidder. The target firm's true value is drawn for a uniform distribution over the interval [0,100] whose realization is known only to it. The bidder does not know the true realized value but does know that whatever the value is the target's firm is worth 3/2 more in its hands - it is a better manager. The theory demonstrates that such markets fail since the risk neutral equilibrium bid is zero. When this game is played in the lab, and presumably in the real world, the predictions of the theory fail. Bidder subjects make bids dispersed all over the support of values.

Now assume that a new bidder firm is contemplating making a take-over bid for a firm under similar circumstances and is instructed to sample how other firms go about making their bid. If this game had been played by a large number of firms before and bidders bid over the whole support of values as they usually do then we would expect that the firms who did best would be those who bid 100 and happened to be lucky enough to find a firm whose value was in a small interval near 100. (The payoff function for the game is $\pi = \frac{3}{2}v - p$. Since any bid below value will be rejected the highest (ex post) payoff results when the bidder bids 100 and the target realized value is also 100, in which case the firm's profit will be 50.) So the most successful firms make the most extreme and irrational bids (as do the biggest failures). If they are imitated the consequences are obvious.

This result is not surprising since we know from the Samuelson-Bazerman experiments that people fail to realize that the expected profit maximizing strategy is to bid 0 and we know from our experiments that people have no problems imitating risky behavior when they are only able to see the payoff ranks of those who went before. So here we would expect to see exactly the results we achieved in our experiments. Only the foolhardy and lucky are ranked on the top (and on the bottom) of the Seed, behavior appears to be risk seeking, and the ex post successful are imitated.

Another application is to imitation in oligopolistic markets where one firm has a first-mover advantage and must choose its output before others (see, for instance, Huck, Muller and Normann, (2001)). Imagine however, that before it moves it is able to see the output choice of other first movers in other identical markets, i.e. markets with the same number of firms and the same demand and cost functions. Clearly if all firms were rational and risk neutral then they would choose the Stackelberg output and sampling other firms behavior would be useless. However, if we expect some variance in firms' decision making, as is true in many experiments including ours, then in the long run those first movers who chose more aggressive outputs than the one dictated by the Stackelberg theory but were lucky to have met up with a meek competitor who limited his or her output choice will be the most successful ones. Such an aggressive choice for our imitator, however, may be profit decreasing if he is matched with more aggressive cohorts.

Finally, our results have implications for social learning. While most social learning problems involve a decision maker being able to see only the action of his predecessors and not their realized profits before she takes an action, our situation combines elements of search theory with social learning - a decision maker must decide whom to learn from before updating his beliefs. This adds another layer of complexity to the social learning problem and also opens another avenue through which outcomes can be socially inefficient.

We could go on and offer other potential applications but the point should be clear. Any situation that involves risk and where extreme choices can be ex post successful but ex ante foolish is one subject to what we might call the "fallacy of imitation" which implies the belief that we should, on average, be successful if we imitate those who have been successful in the past. Our results suggest that this is not true.

6 Conclusions

This paper has demonstrated that a common heuristic of "imitate the best" (or sampling for imitation) can lead economic agents to make decisions that are welfare decreasing. It does so because it fails to take into account the fact that those who have done well may have chosen irresponsibly but happened to be lucky. Copying their recklessness may be a blueprint for disaster. In addition, such a heuristic seems to make economic agents appear risk preferring when in fact their underlying preferences are quite the opposite. This result is striking because in the experiment we perform, another heuristic, "sample to learn", is readily available to subjects if they think about the task at hand. The fact that so few subjects avail themselves of this strategy makes us believe that imitation of successful others is a dominant behavioral principle when decision makers face an unfamiliar task. That fact, coupled with the result that imitation can lead to socially undesirable consequences is something that should be considered by policy makers interested in improving the efficiency of markets where new innovations or businesses are being considered. In such markets, entrepreneurs should be discouraged from merely imitating their successful predecessors.

References

[1] Abbink, Klaus and Jordi Brandts (2004). "24", University of Nottingham and IAE. Mimeo.

[2] Apesteguia, Jose, Steffen Huck and Jorg Oechssler (2004). "Imitation: Theory and Experimental Evidence", Public University of Navarre, University College London and University of Bonn. Mimeo.

[3] Camerer, Colin F. (1997). "Progress in Behavioral Game Theory". *Journal of Economic Perspectives* 11 (4), 167-188.

[4] Celen, Bogachan, Choi, Syngjoo and Hyndman, Kyle (2005) , Endogenous Networks, work in progress, New York University

[5] Dugatkin, Lee Alan and Jean-Guy J. Godin (1992). "Reversal of Female Mate Choice by Copying in the Guppy (Poecilia Reticulata)". *Proceedings: Biological Sciences* 249, 179-184

[6] Ellison, Glenn and Drew Fudenberg (1995). "Word-of-Mouth Communication and Social Learning". *Quarterly Journal of Economics*, 440, 93-125.

[7] Gibson, Robert M. and Jacob Hoglund (1992). "Copying and Sexual Selection". *Trends in Ecology and Evolution* 7, 229-231.

[8] Gilboa, Itzhak and David Schmeidler (1995). "Case-Based Decision Theory". *Quarterly Journal of Economics* 110, 605-639.

[9] Gilboa, Itzhak and David Schmeidler (2001). "A Theory of Case-Based Decisions". Cambridge University Press.

[10] Holt, Charles and Susan Laury (2002). "Risk Aversion and Incentive Effects". *American Economic Review* 92, 1644-1655.

[11] Huck, Steffen, Wieland Muller and Hans-Theo Normann (2001). "Stackelberg Beats Cournot: on Collusion and Efficiency in Experimental Markets". *Economic Journal* 111, 749-765

[12] Huck, Steffen, Hans-Theo Normann and Jorg Oechssler (1999). "Learning in Cournot Oligopoly: An Experiment". *Economic Journal* 109, C80-C95.

[13] Offerman, Theo, Jan Potters and Joep Sonnemans. (2002). "Imitation and Belief Learning in an Oligopoly Experiment". *Review of Economic Studies* 69, 973-997.

[14] Offerman, Theo and Joep Sonnemans (2004). "What's Causing Overreaction? An Experimental Investigation of Recency and the Hot-hand Effect". *Scandinavian Journal of Economics* 106, 533-553.

[15] Rabin, Mathhew (2002). "Inference by Believers in the Law of Small Numbers", *Quarterly Journal of Economics* 117, 775-816.

[16] Ridley, Mark and Catherine Rechten (1981). "Female Sticklebacks Prefer to Spawn with Males Whose Nests Contain Eggs". *Behaviour* 76, 152-161.

[17] Samuelson, William and Max Bazerman (1985). "The Winner's Curse in Bilateral Negotiations". In V.L. Smith, ed., *Research in Experimental Economics*, Vol. 3.. Greenwich, Conn.: JAI Press, 105-137.

[18] Schlag, Karl (1998). "Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-Armed Bandits". *Journal of Economic Theory* 78, 130-156.

[19] Schlag, Karl (1999). "Which One Should I Imitate?". *Journal of Mathematical Economics* 31, 493-522.

[20] Selten, Reinhard and Jose Apesteguia (2004). "Experimentally Observed Imitation and Cooperation in Price Competition on the Circle". *Games and Economic Behavior*, forthcoming.

[21] Selten, Reinhard and Axel Ostmann (2001). "Imitation Equilibrium". *Homo Oeconomicus*, 43, 111-149.

[22] Vega-Redondo, Fernando (1997). "The Evolution of Walrasian Behavior". *Econometrica*, 61, 57-84.

[23] Wakker, Peter (2005). "The Logpower (CRRA) Parametric Utility Family Graphically Explained", University of Amsterdam, Mimeo.

Appendix: INSTRUCTIONS EXPERIMENT (Treatment: SEED)

Welcome to this experiment on decision-making! You can make money in this experiment. Read the instructions carefully. There is paper and a pen on your table. You can use these during the experiment.

THE EXPERIMENT

You will earn points in the experiment by making an investment decision. At the end of the experiment your points will be exchanged to dollars. Each 500 points will yield 1 dollar. The experiment consists of 2 rounds. At the end of the experiment the computer will randomly select one of the two rounds. Your earnings will be equal to the earnings of this randomly selected round. In addition we will give you 5000 points. There is a possibility that you will lose some points in a round. In case you lose points in the round that is actually paid, we will deduct this loss from the amount of 5000 points that we have given you. After the two rounds you will be able to make some more money in an additional experiment that we will describe later. Your earnings for this additional experiment will be added to the earnings that you made in the round that was actually paid.

INVESTING IN LOCAL OR NATIONAL MARKET

In this experiment you will act as a firm who has to decide on how much of his or her product to produce. You may sell your good in one of two markets, the local market or the national market. The price of the good in both markets will be identical but determined randomly. The costs of production of the good, however, will vary depending on which market you choose to produce in.

If you opt for the local market, your production level must be a lower number than when you opt for the national market. If you produce in the local market you must choose a production level which is an integer between 10 and 40, while if you choose to produce in the national market the production level must equal an integer number between 41 and 90. In the first round you will not be informed about the price that you will receive for each unit produced before you make your production decision. This will be determined by a random draw of the computer. The price will be a number between 10 and 90 points, and each number between 10 and 90 is equally likely.

Each product that you will produce will cost you some points. In the local market your costs are described by the attached cost table which lists the costs for all production levels that you may choose. Please look at this table now. Note that on this table the cost of producing more output is not only increasing but increasing at an increasing rate so that, for example, the cost of producing 30 is more than twice the cost of producing 15 etc.
In the national market your costs will be determined in a similar way as in the local market, but now the resulting cost will be multiplied by a factor that equals (1+0.01*random shock). In other words, the national market is more risky than the local market because your costs of production are random whereas they were not random in the local market. The random shock will be equal to an integer number between -60 and 60, and each number between -60 and 60 is equally likely. Thus, for example, say you choose a production level of 60 in the national market. If such a production level were feasible in the local market it would cost you 3600 (please refer to your cost table) while your costs in the national market will depend on what random shock you receive. If the random shock is +30 then your costs will be (3600(1+ 0.30) or 4680. If the random shock were -30 then the same production level would cost only 2520.

Your profit in a round will be determined as follows. The price for the product will be multiplied by your production level. The resulting number is multiplied by two. Your profit equals this number minus your costs. Or,

Profit = 2* price * your production - your costs

Round 2

After you have made your choice for round 1 you will proceed to round 2. Round 2 is identical to round 1 except that here you will learn the price of the product you produce before you make your production decision. That is, at the beginning of round 2 you will receive the information about what price you will receive for each unit that you produce. Then you may again decide whether you want to produce for the local or the national market and how much you want to produce. When everybody has made their decisions for round 2 we will hand out a new set of instructions describing another experiment that we would like you to participate in. When you finish this experiment, you will be informed about the results of the first two rounds and be paid. Your payment will be equal to your earnings in either round 1 or round 2 (chosen randomly) plus your earnings in the additional experiment we will perform after round 2.

END

You have reached the end of the instructions. If you want to read some parts of the instructions again, push the button BACK TO THE START or the button PREVIOUS. When you are ready, push the button READY. When all participants have pushed READY, the experiment will start. When the experiment has started, you will NOT be able to return to these instructions.

If you still have questions, please raise your hand!

INSTRUCTIONS ROUND 2

In this round the random draw by the computer is revealed to you. The random draw for the PRICE equals 38. You are again asked to make a decision whether you want to produce for the local market or the national market and how much you want to produce. The production circumstances remain the same. In case that you decide to produce for the national market, a new independent random shock will determine your costs.

Cost Table (distributed in all treatments)

production	costs	production	costs	production	costs
10	100	37	1369	64	4096
11	121	38	1444	65	4225
12	144	39	1521	66	4356
13	169	40	1600	67	4489
14	196	41	1681	68	4624
15	225	42	1764	69	4761
16	256	43	1849	70	4900
17	289	44	1936	71	5041
18	324	45	2025	72	5184
19	361	46	2116	73	5329
20	400	47	2209	74	5476
21	441	48	2304	75	5625
22	484	49	2401	76	5776
23	529	50	2500	77	5929
24	576	51	2601	78	6084
25	625	52	2704	79	6241
26	676	53	2809	80	6400
27	729	54	2916	81	6561
28	784	55	3025	82	6724
29	841	56	3136	83	6889
30	900	57	3249	84	7056
31	961	58	3364	85	7225
32	1024	59	3481	86	7396
33	1089	60	3600	87	7569
34	1156	61	3721	88	7744
35	1225	62	3844	89	7921
36	1296	63	3969	90	8100

INSTRUCTIONS EXPERIMENT (Treatment: Sample -Three)

Welcome to this experiment on decision-making! You can make money in this experiment. Read the instructions carefully. There is paper and a pen on your table. You can use these during the experiment.

THE EXPERIMENT

You will earn points in the experiment by making one investment decision. At the end of the experiment your points will be exchanged to dollars. Each 500 points will yield 1 dollar. Your earnings will be equal to what you earn for your one decision choice. In addition we will give you 5000 points. There is a possibility that you will lose some points in the experiment. In case you lose points, we will deduct this loss from the amount of 5000 points that we have given you. After this part of the experiment is over you will be able to make some more money in two other experiments that will be described to you later. Your earnings for these experiments will be added to the earnings that you made in the first experiment.

After you have made your investment decision but before you learn your profit we will ask you to participate in two other experiments which will be described later. Your earnings in these experiments will be added to those you have made already.

INVESTING IN LOCAL OR NATIONAL MARKET

In this experiment you will act as a firm who has to decide on how much of his or her product to produce. You may sell your good in one of two markets, the local market or the national market. The price of the good in both markets will be identical but determined randomly. The costs of production of the good, however, will vary depending on which market you choose to produce in.

If you opt for the local market, your production level must be a lower number than when you opt for the national market. If you produce in the local market you must choose a production level which is an integer between 10 and 40, while if you choose to produce in the national market the production level must equal an integer number between 41 and 90. You will not be informed about the price that you will receive for each unit produced before you make your production decision. This price has been determined by a random draw of the computer. The price will be a number between 10 and 90 points, and each number between 10 and 90 is equally likely.

Each product that you will produce will cost you some points. In the local market your costs are described by the attached cost table which lists the costs for all production levels that you may choose. Please look at this table now. Note that on this table the cost of producing more output is not only increasing but increasing at an increasing rate so that, for example, the cost of producing 30 is more than twice the cost of producing 15 etc.

In the national market your costs will be determined in a similar way as in the local market, but now the resulting cost will be multiplied by a factor that equals (1+0.01*random shock). In other words, the national market is more risky than the local market because your costs of production are random whereas they were not random in the local market. The random shock will be equal to an integer number between -60 and 60, and each number between -60 and 60 is equally likely. Thus, for example, say you choose a production level of 60 in the national market. If such a production level were feasible in the local market it would cost you 3600 (please refer to your cost table) while your costs in the national market will depend on what random shock you receive. If the random shock is +30 then your costs will be (3600(1+ 0.30) or 4680. If the random shock were -30 then the same production level would cost only 2520.

Your profit in a round will be determined as follows. The price for the product will be multiplied by your production level. The resulting number is multiplied by two. Your profit equals this number minus your costs. Or,

Profit = 2* price * your production - your costs

Before you make your investment decision we will give you a chance to observe what choice others who have done this exact same experiment before you have chosen. More precisely, before you make your investment decision you will see on the bottom of the screen a list of subjects ranked by their success in previous experiments. That person who received the highest payoff from his or her decision choice will be denoted as subject number 1 while that person who did second best will be denoted by the number 2, etc. In other words those who did best will be denoted by lower numbers. In addition, next to each person will be the letters "loc" or "nat" indicating whether that person invested in the local or national market. The experiment performed by these subjects was identical to the one you will be performing here except that the ranked subjects did not have the chance to observe other subjects' actions and profits as you do. Before any subject did the experiment one price was drawn from the distribution of prices described above and that price was used to calculate profits for all subjects. Thus, you will face the same price as the previous subjects did but the level of that price will not be shown to you. In addition, the way their costs and profits was determined is identical to yours. By clicking on the button associated with any previous subject, you will see their production level as well as their profit. You will be able to click on at least one such subject. After you have clicked a person, you will be informed whether you are allowed to click again on another person. You will not know how many persons you are allowed to click on. At some point we will inform you that you can not click on other persons anymore, and then you will have to enter your investment decision.

END

You have reached the end of the instructions. If you want to read some parts of the instructions again, push the button BACK TO THE START or the button PREVIOUS. When you are ready, push the button READY. When all participants have pushed READY, the experiment will start. When the experiment has started, you will NOT be able to return to these instructions.

If you still have questions, please raise your hand!

Beliefs

In the next part of the experiment we ask you to reveal to us what you believe the price for the product was in the investment decision experiment. On the bottom of the computer screen will be 8 boxes with numbers above them. The numbers are indicated as 10-19, 20-29, 30-39, 40-49, 50-59, 60-69, 70-79, and 80-90. What we ask you to do is to indicate the probability that you feel the price in the investment decision experiment was in each of these intervals by allocating percentages to each box that add up to 100%. For example, if you think that the probability was 50% that the price was between 10 and 19 and 50% that it was between 80 and 90 then place the numbers 50 in the first box and 50 in the last box. We will reward you for the accuracy of your predictions as follows: Suppose your beliefs are as shown in the following table:

price	prediction
10-19	15%
20-29	15%
30-39	5%
40-49	15%
50-59	5 %
60-69	5%
70-79	15%
80-90	25%

If these are your beliefs about the price, then enter them in the boxes at the bottom of the computer screen. Note that the total of the percentages you type in must sum to 100. In addition say that the price was actually 27. Since 27 is in the interval 20-29 that is the true price interval. We will then determine your prediction payoff as follows: Prediction Payoff = $3{,}000 - 0.15(100-15)^2 - 0.15(0-15)^2 - 0.15(0-5)^2 - 0.15(0-15)^2 - 0.15(0-5)^2 - 0.15(0-5)^2 - 0.15(0-15)^2 - 0.15(0-25)^2 = 3000 - 0.15(7225 + 225 + 25 + 225 + 25 + 25 + 225 + 625) = 3{,}000-0.15(8600) =3000-1290 = 1710$. In other words, we will give you a fixed amount of 3,000 points from which we will subtract an amount that depends on how inaccurate your prediction was. To do this when we find out what the true price is (i.e.in this example it was 27 which is between 20 and 29), take the number you assigned to that choice, in this example 15, subtract it from 100 (this is the extent to which you made a mistake in guessing the price) and square it and multiply it by 0.15. Then take the numbers you assigned to the offer ranges that did not contain the actual price, (these are also mistakes) square them, add them up and multiply the sum by 0.15 as well. These squared numbers will then be subtracted from the 3,000 points we initially gave you to determine your final point payoff since they represent the sum total of the squared mistakes you made in your predictions. Your point payoff will again be converted into dollars at the rate of 500 points = 1 dollar. Note that the worst you can do under this payoff scheme is to state that you believe that a certain price interval is the true one with a 100% chance and assign 100 to that choice when in fact this is not true. Here your payoff from prediction would be 0. Similarly, the best you can do is to guess correctly and assign 100 to that price interval which turns out to be the actual one. Here your payoff will be 3,000. However since your prediction is made before you know what the true price is, the best thing you can do to maximize the expected size of your prediction payoff is to simply state your true beliefs about the price. Any other prediction will decrease the amount you can expect to earn as a prediction payoff.

If you still have questions, please raise your hand!

Instructions Holt-Laury procedure (all treatments; this part was not computerized)

Your decision sheet shows ten decisions listed on the left. Each decision is a paired choice between "Option A" and "Option B." You will make ten choices and record these in the final column, but only one of them will be used in the end to determine your earnings. Before you start making your ten choices, please let me explain how these choices will affect your earnings for this part of the experiment.

Here is a ten-sided die that will be used to determine payoffs; the faces are numbered from 1 to 10 (the "0" face of the die will serve as 10.) After you have made all your choices, we will throw this die twice, once to select one of the ten decisions to be used, and a second time to determine what your payoff is for the option you chose, A or B, for the particular decision selected. Even though you will make ten decisions, only of these will end up affecting your earnings, but you will not know in advance which decision will be used. Obviously, each decision has an equal chance of being used in the end.

Now, please look at Decision 1 at the top. Option A pays 200 pennies if the throw of the ten-sided die is 1, and it pays 160 pennies if the throw is 2-10. Option B yields 385 pennies if the throw of the die is 1, and it pays 10 pennies if the throw is 2-10. The other Decisions are similar, except that as you move down the table, the chances of the higher payoff for each option increase. In fact, for Decision 10 in the bottom row, the die will not be needed since each option pays the highest payoff for sure, so your choice here is between 200 pennies or 385 pennies.

To summarize, you will make ten choices: for each decision row you will have to choose between Option A and Option B. You may choose A for some decision rows and B for other rows, and you may change your decisions and make them in any order. When you are finished, we will come to your desk and throw the ten-sided die to select which of the ten Decisions will be used. Then we will throw the die again to determine your money earnings for the Option you chose for that Decision. Earnings (in pennies) for this choice will be added to your previous earnings, and you will be paid all earnings in cash when we finish.

So now please look at the empty boxes on the right side of the record sheet. You will have to write a decision, A or B in each of these boxes, and then the die throw will determine which one is going to count. We will look at the decision that you made for the decision that counts, and circle it, before throwing the die again to determine your earnings for this part. Then you will write your earnings in the blank at the bottom of the page.

Are there any questions? Now you may begin making your choices. Please do not talk with anyone while we are doing this; raise your hand if you have a question.

Decision sheet

	Option A				Option B				
Decision	Throw	payoff	throw	payoff	throw	payoff	throw	payoff	choice
1	1	200	2-10	160	1	385	2-10	10	
2	1-2	200	3-10	160	1-2	385	3-10	10	
3	1-3	200	4-10	160	1-3	385	4-10	10	
4	1-4	200	5-10	160	1-4	385	5-10	10	
5	1-5	200	6-10	160	1-5	385	6-10	10	
6	1-6	200	7-10	160	1-6	385	7-10	10	
7	1-7	200	8-10	160	1-7	385	8-10	10	
8	1-8	200	9-10	160	1-8	385	9-10	10	
9	1-9	200	10	160	1-9	385	10	10	
10	1-10	200			1-10	385			

Table 1
Experimental Design

treatment	# subjects	# rounds	# samples	belief elicitation	Holt-Laury risk aversion test
Seed	60	1 without price; 1 with price	0	none	at the end
Sample-One	32	1	1	after sampling	at the end
Sample-Three	25	1	3	after sampling	at the end

Table 2a
Seed and Sampling

Seed: publicly observable		Seed: observable after sampling		Sample-One	Sample-Three		
rnk	mrk	prod	profit	only	first	second	third
1	nat	63	2922.6	63.2%	72.0%	4%	4%
2	nat	90	2790.0	3.5%	4%	12%	28%
3	nat	60	2760.0	1.8%	0%	4%	12%
4	nat	50	2525.0	0%	0%	0%	4%
5	nat	49	2523.5	0%	0%	8%	4%
6	nat	45	2306.3	0%	0%	0%	0%
7	nat	50	2300.0	0%	0%	0%	0%
8	nat	60	2184.0	0%	0%	0%	0%
9	nat	60	2040.0	0%	0%	0%	0%
10	nat	70	2037.0	1.8%	0%	0%	4%
11	nat	60	1752.0	0%	0%	0%	0%
12	loc	40	1440.0	8.8%	4%	36%	4%
13	loc	40	1440.0	1.8%	0%	0%	12%
14	loc	40	1440.0	1.8%	0%	0%	0%
15	loc	40	1440.0	0%	0%	0%	0%
16	loc	40	1440.0	0%	0%	0%	0%
17	loc	40	1440.0	0%	0%	0%	0%
18	loc	40	1440.0	1.8%	0%	0%	0%
19	loc	40	1440.0	0%	0%	0%	0%
20	loc	40	1440.0	1.8%	4%	0%	0%
21	loc	40	1440.0	0%	0%	0%	0%
22	loc	40	1440.0	0%	0%	0%	0%
23	loc	40	1440.0	0%	0%	0%	0%
24	loc	40	1440.0	0%	0%	0%	0%
25	loc	35	1435.0	0%	0%	0%	0%
26	loc	35	1435.0	1.8%	4%	0%	0%
27	loc	35	1435.0	1.8%	0%	0%	0%
28	loc	33	1419.0	0%	0%	4%	0%
29	loc	32	1408.0	0%	0%	0%	0%
30	loc	32	1408.0	0%	0%	0%	4%

Table 2b
Seed and Sampling (continued)

Seed: publicly observable		Seed: observable after sample		Sample-One	Sample-Three		
rnk	mrk	prod	profit	only	first	second	third
31	loc	30	1380.0	0%	0%	0%	0%
32	loc	30	1380.0	0%	0%	0%	0%
33	loc	30	1380.0	1.8%	4%	0%	0%
34	loc	30	1380.0	0%	0%	0%	0%
35	loc	30	1380.0	0%	0%	0%	0%
36	loc	30	1380.0	0%	0%	0%	0%
37	loc	30	1380.0	1.8%	4%	0%	0%
38	nat	63	1374.7	0%	0%	0%	0%
39	loc	29	1363.0	0%	0%	0%	0%
40	loc	29	1363.0	1.8%	4%	4%	0%
41	loc	28	1344.0	0%	0%	0%	0%
42	loc	27	1323.0	0%	0%	0%	0%
43	loc	27	1323.0	0%	0%	0%	0%
44	loc	25	1275.0	0%	0%	0%	0%
45	loc	25	1275.0	0%	0%	0%	0%
46	loc	20	1120.0	0%	0%	0%	0%
47	loc	20	1120.0	0%	0%	0%	4%
48	nat	65	1095.3	0%	0%	0%	0%
49	loc	18	1044.0	0%	0%	0%	0%
50	nat	45	1010.2	0%	0%	0%	0%
51	loc	15	915.0	0%	0%	0%	8%
52	nat	41	813.0	0%	0%	0%	0%
53	nat	63	541.2	0%	0%	0%	0%
54	nat	47	501.5	0%	0%	0%	0%
55	nat	50	-75.0	0%	0%	0%	4%
56	nat	50	-150.0	0%	0%	0%	0%
57	nat	70	-756.0	0%	0%	0%	0%
58	nat	67	-968.2	0%	0%	0%	0%
59	nat	70	-1246.0	0%	0%	0%	0%
60	nat	75	-1275.0	5.3%	0%	28%	8%

Table 3
Production Levels across Treatments

	% national	production
Seed without info	38.3	42.6 (16.0)
Seed with info	40.0	47.0 (19.8)
Sample-One	71.9	54.5 (18.5)
Sample-Three	76.0	55.3 (17.7)
Hypotheses: comparisons of production levels		
Seed with versus without info	Wilcoxon rank test	p=0.25
Seed (without info) versus Sample-One	Mann-Whitney rank test	p=0.00
Seed (without info) versus Sample-Three	Mann-Whitney rank test	p=0.00
Sample-One versus Sample-Three	Mann-Whitney rank test	p=0.93

Notes: standard deviations are listed in parentheses.

Table 4
Observed Market First Sample and Subsequent Production

	observed market of first sample	frequency	subsequent production	produce in national market
Sample-One	local	25%	33.4 (12.7)	25%
	national	75%	61.6 (14.4)	87.5%
Sample-Three	local	24%	43.7 (15.0)	50%
	national	76%	59.0 (17.6)	84.2%

Notes: standard deviations appear in parentheses.

Table 5
Regression Results

Regression		
constant	-0.93 (11.34)	p=0.94
$ß_1$	0.93 (0.19)	p=0.00
$ß_2$	0.00 (0.07)	p=0.96
$ß_3$	0.34 (0.27)	p=0.20
$ß_4$	0.31 (0.25)	p=0.23
$ß_5$	0.04 (0.16)	p=0.83
$ß_6$	0.09 (0.14)	p=0.52
R^2=0.35; n=57		

Table 6
Two Risk Attitudes Measures

	r_{HL}		r_{prod}		Wilcoxon test
Treatment	Mean (std. dev.)	Median	mean (std. dev.)	median	H_0: $r_{HL}=r_{prod}$
Seed	0.31 (0.48)	0.28	0.02 (1.17)	0.37	p=0.16
Sample-One	0.28 (0.36)	0.28	-0.79 (1.94)	-0.57	p=0.00
Sample-Three	0.48 (0.48)	0.55	-0.90 (2.07)	-0.57	p=0.00

Notes: r_{HL} represents the risk coefficient implied by the Holt-Laury procedure; r_{prod} denotes the risk coefficient inferred from the production decision. Standard deviations appear in parentheses.

Table 7
Beliefs

	probability intervals								
	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80-90	mean
Sample-1	9.8	12.7	11.1	13.7	17.5	14.4	12.1	8.9	50.3
Sample-3	10.6	9.0	15.5	15.2	13.4	15.9	11.8	8.6	50.0
Loc	6.9	9.5	12.0	15.6	17.7	15.1	14.4	9.1	52.7
Nat	11.2	11.6	13.4	13.9	15.1	15.0	11.2	8.7	49.3

Notes: Each cell reports the probability (in %) that an average subject in the row assigns to the probability interval of the column. Loc [nat] reports the beliefs of all subjects whose first (or only) sample was in the local [national] market. For the final two rows the data of Sample-One and Sample-Three are pooled.

Figure 1
Histogram Production Levels Seed without Information Price

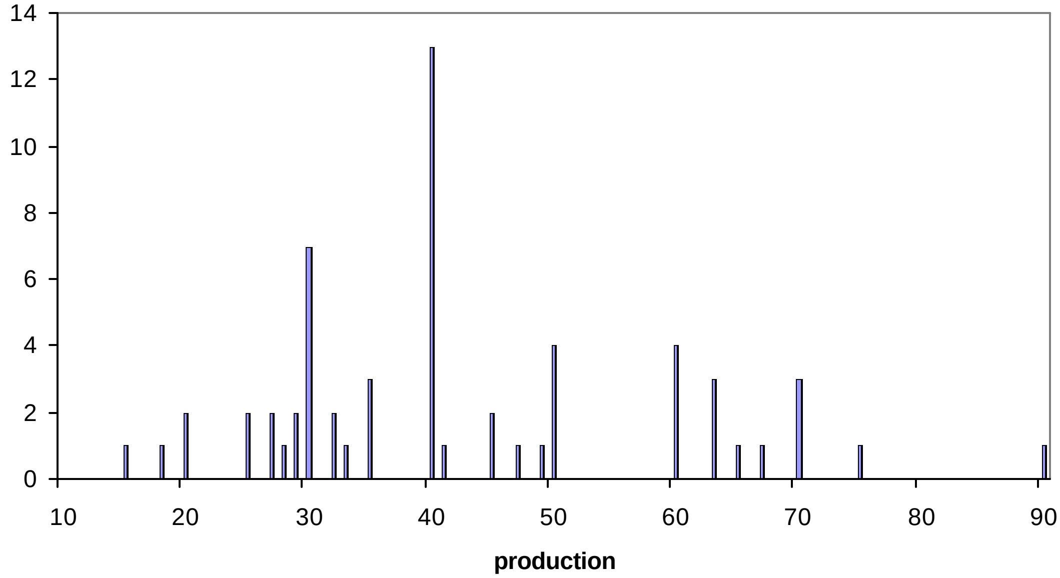

Figure 2
Histogram Production Levels Seed with Information Price

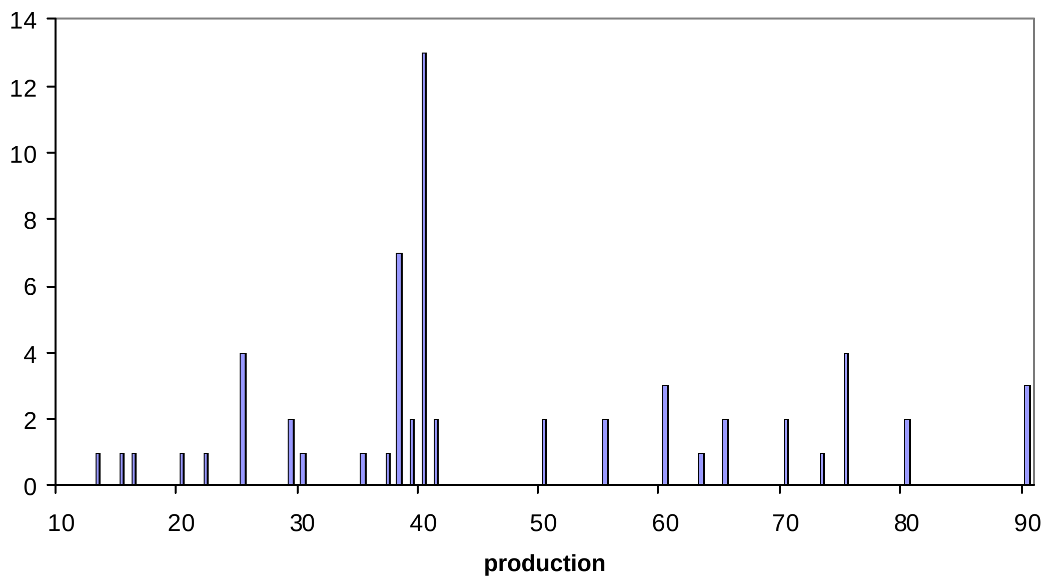

Figure 3
Histogram Production Levels Sample-One

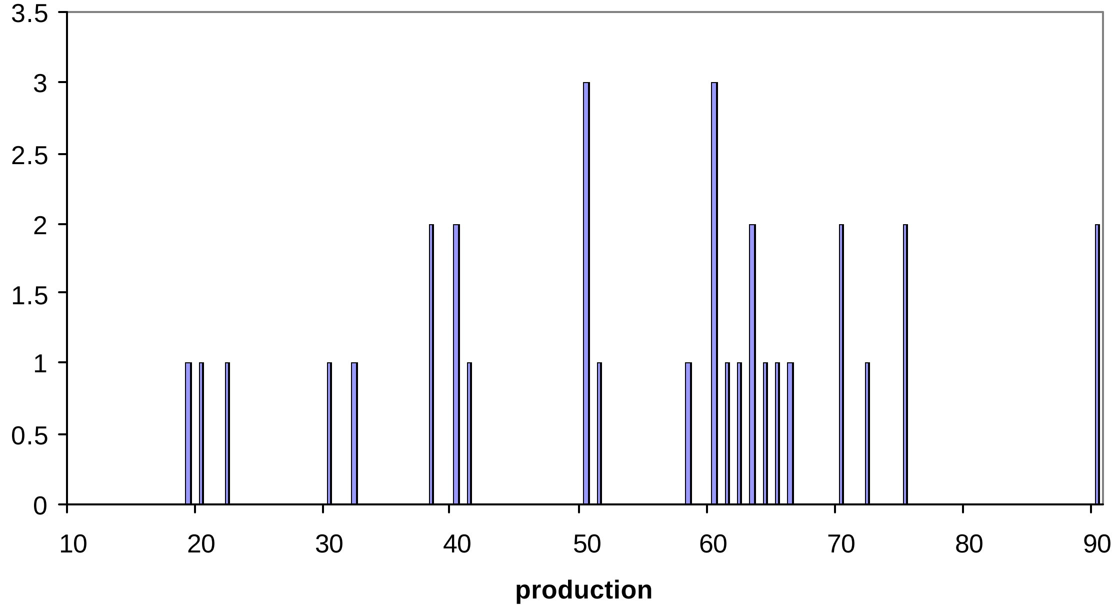

Figure 4
Histogram Production Levels Sample-Three

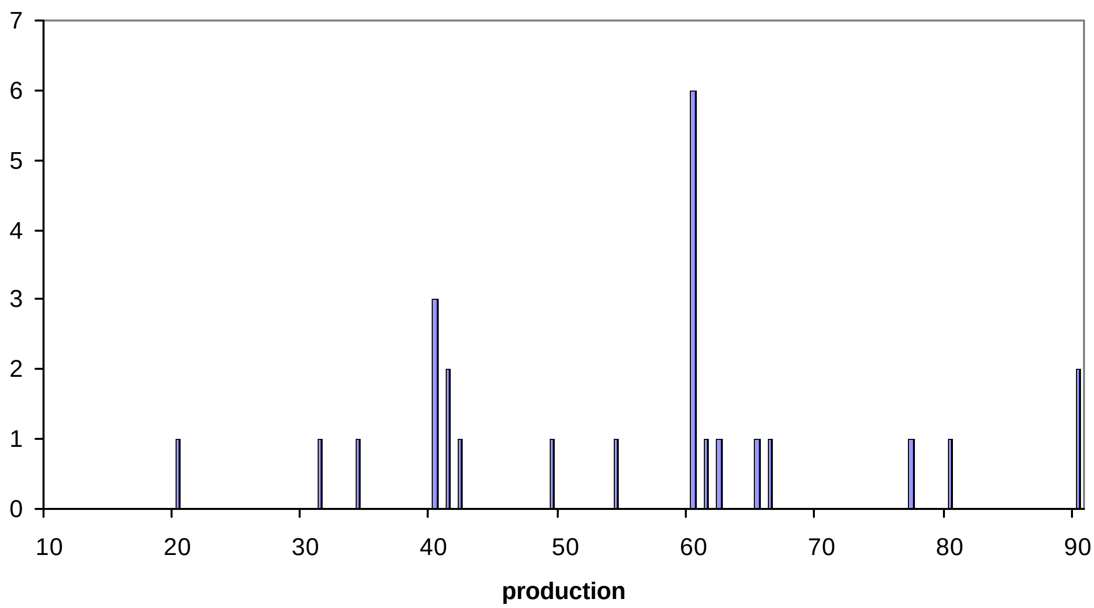

Figure 5
Production Levels Seed versus Sampling Treatments

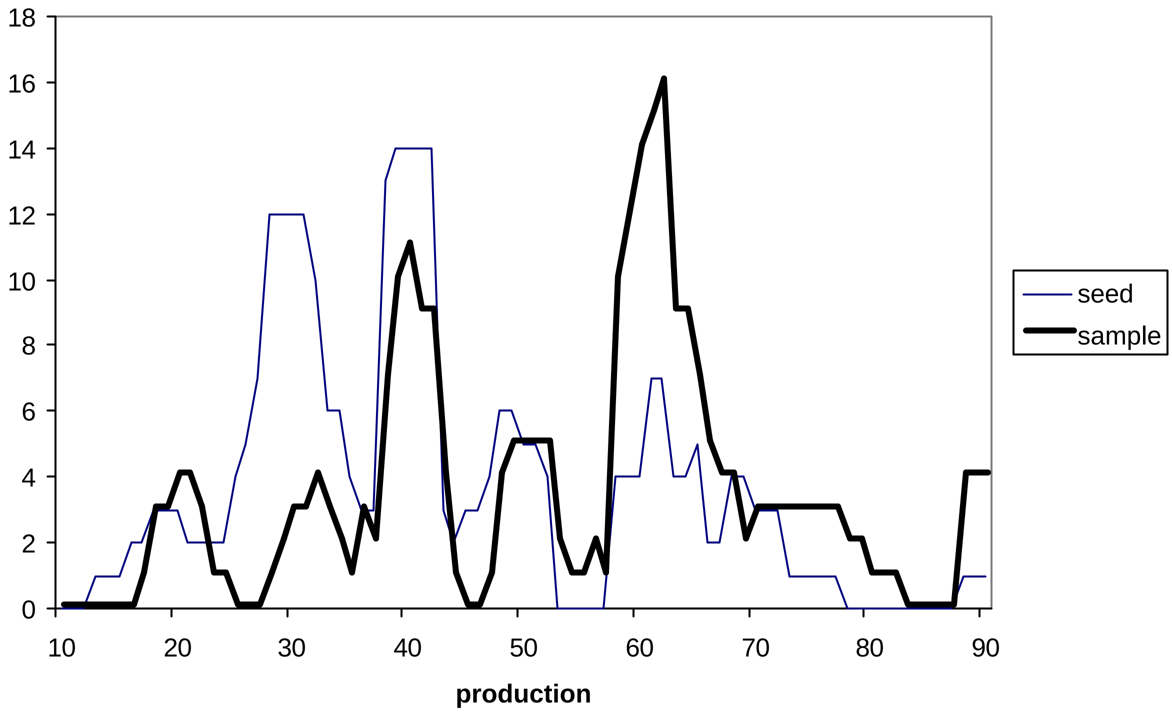

Notes: for each production level on the horizontal axis the vertical axis reports the number of cases that a production level within the interval [production-2, production+2] was observed.

Figure 6
Comparison Cumulative Density Function Risk Holt-Laury across treatments

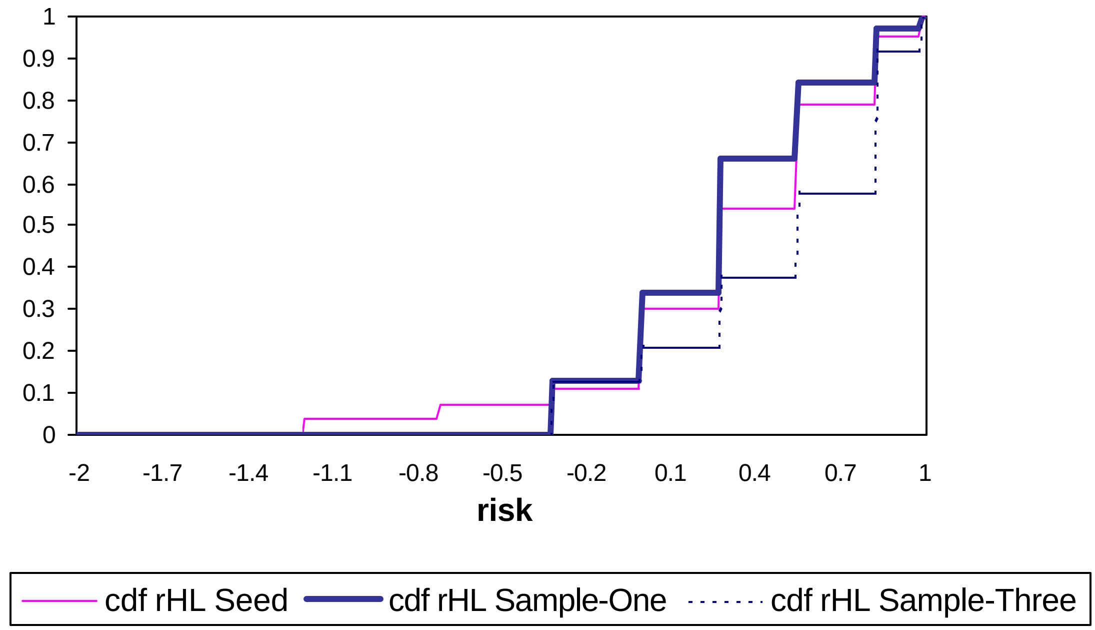

Figure 7
Comparison Cumulative Density Function Risk Coefficients Seed

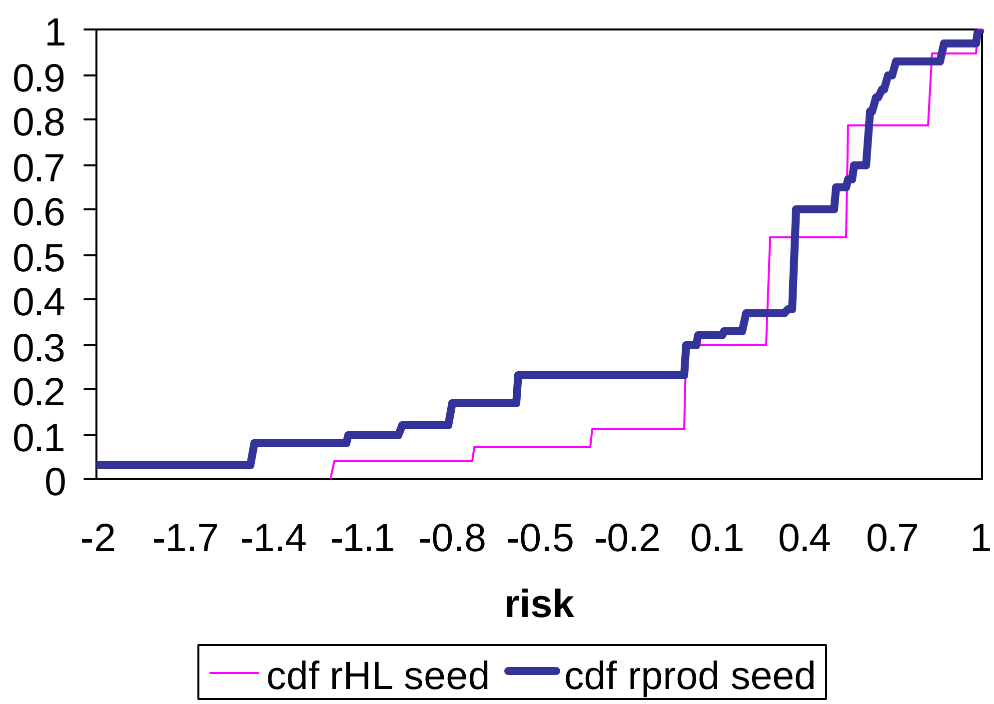

Figure 8
Comparison Cumulative Density Function Risk Coefficients Sample-One

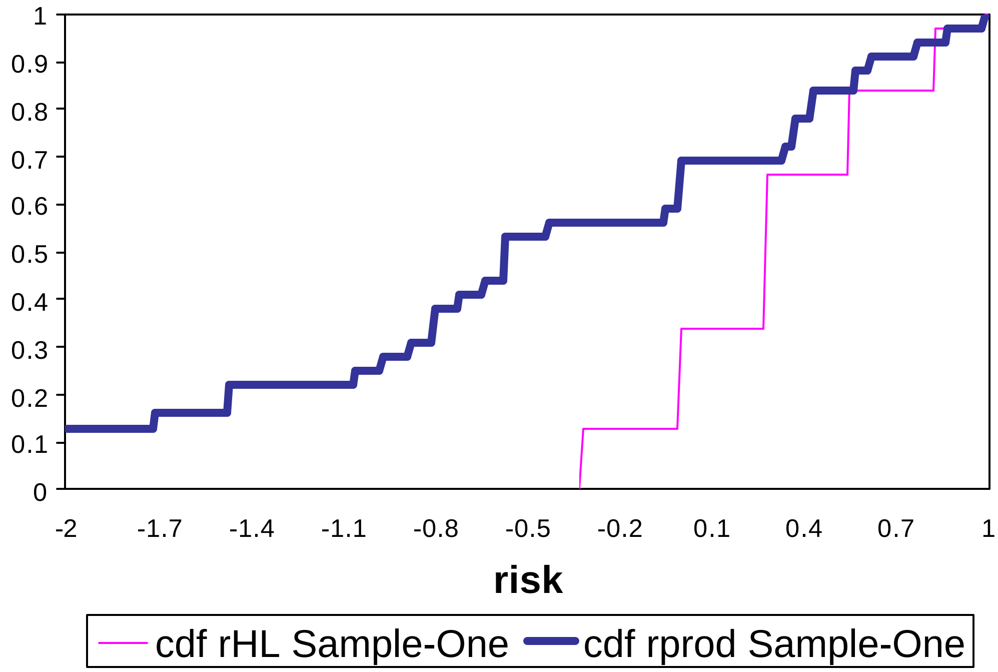

Figure 9
Comparison Cumulative Density Function Risk Coefficients Sample-Three

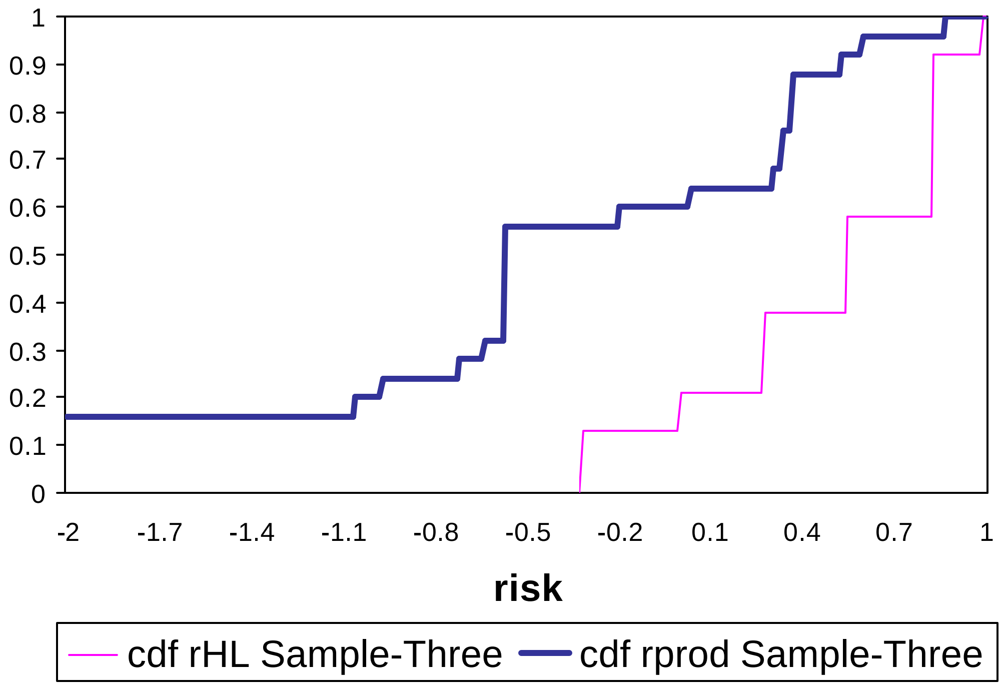

Figure 10
Cumulative Density Function Profits

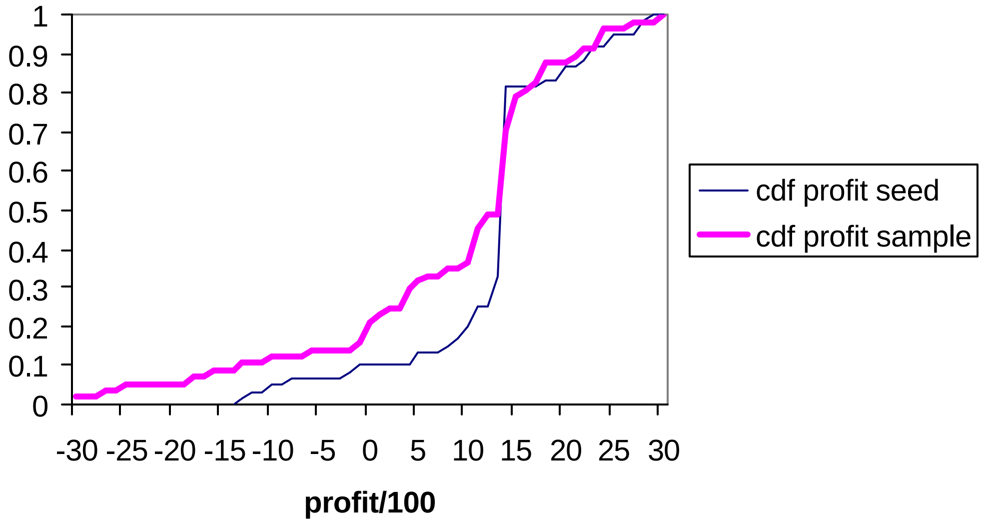

Figure 11
Histogram Probabilities True Price