C.E.S.S. Working Paper Series

Working Paper #2007-05

Imitation and Luck: An Experimental Study on Social Sampling

Theo Offerman and Andrew Schotter

Key Words: Imitation; Social Learning
JEL Classification: D83, C92

19 West Fourth Street, New York, New York 10003, 212-992-8660

The Center for Experimental Social Science

The Center for Experimental Social Science (C.E.S.S.) is an inter-disciplinary research Center at New York University engaged in laboratory experimental work in the social sciences that combines high quality economic theory, psychology and political science. The over arching research tool used in the Center is laboratory experiments in which human subjects engage in decision-making tasks whose purpose is to test both economic theory and/or the properties of proposed or existing economic, political, or social institutions. The working papers in this series reflect these interests.

Imitation and Luck: An Experimental Study on Social Sampling*

Theo Offerman
University of Amsterdam

Andrew Schotter
New York University

February 12, 2007

Abstract

In this paper, we present the results of two experiments on social sampling. In both experiments, 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 are ranked on the basis of their success. In the first experiment, we find that, by and large, subjects sample and imitate lucky risk seekers, while they could have sampled others to retrieve information that is valuable to solve their problem rationally. The simple behavioral rule of imitating the best appears to be robust to the setting of the problem. In the second experiment, we find that subjects tend to imitate successful others in both the winner's curse version and the loser's curse version of the Bazerman-Samuelson takeover game. Because of the way these problems are constructed, imitation exacerbates the winner's curse while it alleviates the loser's curse. In all problems, social sampling makes people look more risk seeking than the people who do not have the opportunity to sample.

JEL: D83, C92, Key Words: Imitation Social Learning

1 Introduction

Imitation may be called the poor man's rationality. What we mean is that if a decision maker were fully rational and capable of costlessly making all necessary calculations, he would not need to imitate anyone before making a decision. Those who feel the need to imitate must, by definition, either not be able to do all the necessary calculations or, if capable, not have the time or inclination to do so. This raises the question of whom to imitate.In many problems imitating the most successful other provides a shortcut to the rational outcome. For instance,

*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.

when a firm invents a new and more efficient production process, it makes a lot of sense for the competetion to try and copy this process. It is therefore not surprising that imitating the best is ingrained in all of us.

In this paper, we ask the question whether people rely on imitating the best in situations where doing so involves certain risks. In particular, we focus on situations where idiosyncratic luck plays a role in determining previous agents' performace. In these situations, those who have been successful may be the foolhardy but lucky ones. For example, if we were to rank people by their degree of success (i.e. their *ex post* payoff) on a task they have performed once, we would expect that on the top of the list would be those people who took big gambles and were lucky. In fact, we know who would be on the top. If the sample of people we are looking at is large, it would be that person who chose that action which, when coupled with the most luck (i.e. the highest realization of the random variable defined in the problem) would determine the highest payoff. This would not necessarily be the person who chose the optimal *ex ante* action or that which would determine the highest expected payoff. On the bottom of the list might be others who made the same choices but were unlucky. In many problems, 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. What we demonstrate is that the desire to imitate the best is so tempting that people consistently fail to distinguish correctly between these situations.

In the experiments we present, the role of imitation varies greatly. In the first problem, a problem of optimal production, the task of the decision maker is to discover the price of his product on the market and then choose an output appropriately. In this problem the decision maker can sample the actions of those who have gone before him and copy them if he likes. The problem is constructed so that sampling the actions of the most successful past participants is not optimal since it is least informative about the price he faces. In short, the subjects must decide whether he or she wants to "sample for information" and sample those people whose decisions ar emost informative, or "sample for imitation" and copy the best. Most subjects tend to not only sample the people on the top of the payoff list but also copy their actions.

The other problem we investigate is the Bazerman-Samuelson (1983) takeover game. Here, subjects are given the opportunity to sample from a list of the takeover bids made by predecessors ranked on profit. In one version of the problem, the Losers' Curse version, imitating the best is optimal in that it leads subjects to make offers that mitigate the loser's curse. In the other version, the Winners' Curse version, imitating the best exacerbates the winner's curse. What is significant is the inability of our subjects to distinguish those situations where imitating the best is beneficial from those where it is not. In all our experiments, subjects tend to imitate the best no matter where that leads them.

We include this game and its two versions since in both versions the cognitive task is almost identical yet in one imitation is beneficial while in the other it is not. Hence, these games help support our belief that what we observe in the production game is not an artifact of the specifics of that problem. In other words, we include these two games to buttress our belief that blind imitation is an ingrained instinct and not an artifact of the context of our production problem.

The consequences of this 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.[1] 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.[2]

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 experimentally. 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.

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,

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

[2] 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.

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. The setup of 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.

Gilboa and Schmeidler (1995; 2001) 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.[3] In such circumstances, decision makers may search their memory for past cases that are similar to theirs.[4] 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.

This paper also 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

[3]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).

[4]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.

access to either all the decisions that have gone before her or at least a subset (see Anderson and Holt (1997), Celen and Kariviv (2004a, 2004b, 2005) and Celen, Kariv and Schotter (2007) where advice is added to the conventional social learning problem. 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 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 three problems presented to 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 concludes.

2 The Problems

In this Section we will present the problems faced by our experimental subejcts in what we call the Production Problem and the winners' and losers' curse versions of the Takeover Game.

2.1 The Production 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 $\underline{q}$ $\leq$ q" $\leq$ $\bar{q}$ where $\underline{q}$ and $\bar{q}$ are the lower and upper limits on production. (In the experiment q" = 40 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),

[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).

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,

$E(\pi) = 2 \cdot E(p) \cdot q - q^2((1 + 0.01 \cdot E(\varepsilon))$. 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.

2.2 The Take-Over Game Problem

We next investigate the role of imitation in the Bazerman-Samuelson (1983) takeover game. This game is played by two players, the target firm and the bidder. The bidder is interested in acquiring the target firm but only wants to do so if the value of the firm is sufficiently high. The value of the target firm is only revealed to the target firm. The bidder knows the distribution from which the true value is drawn but not the true value itself. The value of the firm is worth 3/2 more in the hands of the bidder – the bidder is the better manager. The bidder submits a take-it-or-leave-it bid and the target firm accepts or rejects the bid. Payoffs are determined in accordance with the decisions of the players. That is, if the target firm with a value V accepts the bid B, the target firm will earn a profit equal to B and the bidder will earn a profit equal to (3/2)*V-B. If the target firm rejects the bid, the target firm earns V while the bidder earns 0.

The theoretical predictions and experimental results depend on the distribution that is used to draw the true value (Holt and Sherman, 1994). For some distributions, like the uniform [0,1000] distribution, the market is predicted to fail. The only viable bid that does not make an expected loss is a bid of 0. The risk-neutral prediction of the theory fails in the lab. Subjects submit bids dispersed all over the support and thus fall prey to the winner's curse, winning the firm but paying on average more than the firm is worth to them (see also Samuelson and Bazerman, 1985; Ball, Bazerman and Caroll, 1991; Selten, Abbink and Cox, 2005 and Charness and Levin, 2006).

In contrast, for other distributions the market is not predicted to fail. For the uniform [1000,2000] distribution for instance, the risk-neutral prediction is that the bidder will bid the maximum value in the support, i.e., 2000. This bid will be accepted by any target firm and both firms make a positive expected profit. Holt and Sherman note that in such situations subjects tend to underbid. They often regret ex post that their bid was below the ex ante optimal bid of 2000 and that it was not accepted. Thus, subjects experience what Holt and Sherman call the loser's curse.

Now consider what may happen when players have access to information on how well previous bidders did in similar situations. In the winner's curse version of the game (U[0,1000]), bidders who submit the optimal bid of 0 will be above the middle but below the top of the ranked list of bidders. At the top of the list are the bidders who submitted positive bids and were lucky. In fact, with sufficiently many bidders in the seed, a person with a very high bid and a lot of luck will appear at the top of the list. Most bidders with positive bids make a loss and appear at the bottom of the list, however. If people sample from the top of the list and imitate what these bidders did, they will submit higher and more risk-seeking bids than they otherwise would. Thus, imitation may exacerbate the winner's curse.

In the loser's curse version of the game, the list of ranked bidders will have the following features when it becomes sufficiently long. At the top of the list will be the ones who submitted very high bids and were lucky, below the middle will be the ones with very low bids and at the bottom will be those with very high bids but who were unlucky. Again, imitation will encourage higher bids, but in this setting it is beneficial to bid higher. Therefore, imitation may alleviate the loser's curse. In the next section, we describe the details of the experiment on the question how sampling affects the loser's curse and the winner's curse in the lab.

3 The Experiments, Experimental Design, and Hypotheses

3.1 The Production Problem

The experiment performed on the Production Problem 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, 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).

The production problem was presented to the subjects on computer terminals. The procedures were then reviewed and questions answered. In total there were three production problem treatments called Prod-Seed, Prod-Sample-One and Prod-Sample-Three. 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 negative earnings.

The Prod-Seed was a treatment run to provide the ranked list of subjects that later subjects in the Prod-Sample-One and Prod-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 Prod-Sample-One treatment subjects would come into the lab and read the instructions describing the problem but would be told that 60 in the Prod-Seed treatment 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 Prod-Seed treatment 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 Prod-Sample-Three treatment was identical to the Prod-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 production treatments except Prod-Seed, subjects made the production decision only once. In Prod-Seed we added one extra task which we will describe below.

The above treatments were embellished in several ways. First of all in Prod-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 Prod-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 production treatments the realized price was 38**). This important feature was explicitly mentioned in the instructions: "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 the Prod-Sample-One and Prod-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

[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.

intervals 10-19, 20-29, 30-39, 40-49, 50-59, 60-69, 70-79, and 80-90. They had 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]

3.1.1 The Production Problem and 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 Prod-Seed treatment. To begin, since in stage 1 of the Prod-Seed treatment prices are assumed to be distributed uniformly over the interval [10,90] and since subjects in Prod-Seed cannot sample for information, risk neutral subjects should 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, they should choose 38. These expectations furnish us with the following hypotheses.

Hypothesis 1: Prod-Seed Behavior

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

Hypothesis 2: Prod-Seed with Price Behavior

Subjects who receive the price in stage 2 of the Prod-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 Prod-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 Prod-Sample-One treatment should be 38. Note that sampling three times in our Prod-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 Prod-Sample-One and Prod-Sample-Three experiments.

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

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

Hypothesis 4: Prod-Sample-Three Behavior;

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

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 Prod-Sample-One behavior with Prod-Sample-Three behavior, because such a comparison reveals whether potential biases observed in Prod-Sample-One are robust. It is possible that with multiple searches subjects find out that imitating the top is risky and not necessarily optimal.

The next hypothesis is a very important one. In stage 1 of the Prod-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 Prod-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 Prod-Sample treatments than we do in Prod-Seed, then we know that subjects sample for imitation and not for information.

Hypothesis 5: Seed-Sample Comparisons

Subjects in both the Prod-Sample-One or the Prod-Sample-Three treatments set lower production levels than those in stage 1 of the Prod-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 Prod-Sample-One and Prod-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).

3.2 The Take-Over Problems

The experiment on the takeover game was run at the University of Amsterdam. This experiment was not computerized. After the instructions were handed out, subjects could read at their own pace before they made their decisions. In total, 110 subjects participated. These subjects were assigned to one of three treatments: 49 participated in the TO-Seed treatment, 31 in the Sample-WC treatment that implemented the winner's curse game and 30 in the Sample-LC treatment that implemented the loser's curse game. The experiment lasted for about 40 minutes in which subjects earned on average 11 euros.

In all takeover treatments, subjects received a starting capital of 10 euros. Subsequent earnings and losses were added to or substracted from the starting capital. All subjects played the role of bidder. It was explained to them that the experimenter would play the role of the target firm and that he would accept all bids that were at least as high as the value of the firm. Note that the target firm has a simple weakly dominating strategy to reject a bid if and only if the bid is smaller than the value of the firm. Therefore, it makes sense to simplify the experiment and simulate the role of the target firm. This simplifying procedure has been used before by Selten, Abbink and Cox (2005) and Charness and Levin(2006).

In the TO-Seed treatment, subjects were informed that they would play two rounds, one of which would be randomly selected for actual payment at the end of the experiment. In the first round, 24 subjects faced the winner's curse problem described above. Only after they had submitted their bids, they received the instructions for the second round, in which they faced the loser's curse problem. The other 25 TO-Seed subjects played the two versions of the game in the reverse order.

Subjects knew that only after the second round the experimenter would determine the value of the firm in each of the two versions of the game. For each subject and each version of the game, the value was drawn independently. A throw with two ten-sided dies determined the value of the firm in the following way. One die had sides labelled 00, 10, 20, ..., 90 and the other had sides labelled 0, 2, ..., 9. Adding the two outcomes gives a draw from a discrete U[0,99] distribution. This number was multiplied by 10 to determine the value of the firm in the winner's curse problem. Here, the value was equally likely 0 cent, 10 cents, ... , 980 cents or 990 cents. A new throw with the two dies determined the value of the firm in the loser's curse problem. Again, the sum of the outcomes of the two dies was multiplied by 10, and to this numer 1010 cents were added. So in the loser's curse version the value of the firm was equally likely 1010 cents, 1020 cents, ..., 1990 cents or 2000 cents.[8]

Before we carried out the other treatments that allowed for sampling, we ranked the subjects on the basis of their profits. We constructed two separate rankings, one for the winner's curse game and the other for the loser's curse treatment. We refer to the former list as Seed-WC and to the latter list as Seed-LC.

In the treatments Sample-WC and Sample-LC, we allowed subjects to sample from the corresponding list of ranked TO-Seed persons. They knew that the subjects of the TO-Seed had faced exactly the same problem as they did. They

[8] Notice that in the experiment we had to approximate the continuous U[0,1000] and U[1000,2000] distributions mentioned in the previous section because the experiment was run by hand.

were also exactly informed about the procedure used to determine the value of each firm in the TO-Seed sessions, and they knew that the value for their own firm was going to be determined independently with the same procedure.

We informed subjects that they could observe the bids of 1, 2 or 3 participants of the ranked list. Subjects got the first observation for free. The experimenter informed the subject (privately) about the bid of the selected ranked person, but not abouth this person's profit. Subjects made the decision to sample an additional person after they had observed what the previous person did. To observe a second person's bid, they had to pay a cost of 10 cents. To observe a third person's bid, they had to pay an additional cost of 50 cents.

Notice that there are some small differences between the sampling procedure of the takeover game compared to the sampling procedure of the production decision. Unlike in the production decision experiment, we did not inform subjects about the profit of the selected TO-Seed person. This information was necessary for a person who wanted to choose an optimal production level, but here the information would be redundant. (Note that in the takeover game a rational person ignores the information from the sample anyway). Another difference was that we allowed subjects to decide how many searches they made at increasing marginal cost. This allows us to infer to what extent subjects search for information in the realistic situation where searching is costly.

3.2.1 The Takeover Problem and Hypotheses

We designed this experiment in particular to test the following hypotheses.

Hypothesis 7: Sample-WC Treatment

Our conjecture is that subjects will sample and imitate people at the top of the list of ranked subjects and that the winner's curse will be exacerbated. The null-hypothesis is that the subjects of the seed and the subjects who sample choose equal bids.

Hypothesis 8: Sample-LC Treatment

Again, we expect that subjects will sample and imitate top ranked people. This will soften the loser's curse. The null-hypothesis is that the subjects of the seed and the subjects who sample choose equal bids.

4 Results

In general the results of our experiments demonstrate a clear proclivity for subjects to imitate the behavior of that subject who performed best in the past on the task at hand. In the Production Problem this leads to extreme production choices and obvious sampling at the top of the Prod-Seed treatment despite the fact that optimal behavior suggests sampling in the middle and more moderate production choices. In the Takeover Problem we again see extreme choices made and sampling from the top. Here, however, while this is payoff decreasing in the Winners' Curse version of the problem, it is actually welfare increasing in the Losers' Curse version.

We will proceed by first discussing the results of our Production Problem and then move on to our two versions of the takeover game.

4.1 Results in the Production Problem

We will first present the results of our Production Problem Experiment by looking at the behavior of subjects in Prod-Seed before we discuss the more important question of how people sample.

4.1.1 The Prod-Seed treatment

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

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

There are several interesting things about subject behavior in Prod-Seed. First, as can be seen in the first four columns of Tables 2A and 2B, which show the rank of each person in Prod-Seed, the market he or she produced in, the production level chosen and the profit made, after 60 subjects made their choices in the Prod-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 Prod-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 Prod-Seed treatment, after they were told the price was 38, production levels actually went up to a mean of 47.0 (compared to 42.6 in stage 1). The difference in stage 1 and stage 2 production levels was not significant according to 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 production level in the Prod-Seed treatment was higher than 38.

Figure 2 shows that in stage 2 of Prod-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 had learned that the price was 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 Prod-Seed is the natural benchmark for the Prod-Sampling treatments and in the remaining part of the paper we will compare the Prod-Sampling treatments with stage 1 of Prod-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 Prod-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 Prod-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 Prod-Sample-One and Prod-Sample-Three treatments sample. We now turn our attention to those treatments.

4.1.2 The Prod-Sample-One and Prod-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 Prod-Sample-One and Prod-Sample-Three treatments we have 32 and 25 subjects respectively. Figures 3 and 4 present the histograms of production levels for subjects in the Prod-Sample-One and Prod-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 Prod-Sample-One and Prod-Sample-Three treatments was equal to 38 since the median choice in both the Prod-Sample-One and Prod-Sample-Three treatments is 60 (compared to a median of 40 in stage 1 of Prod-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 Prod-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 Prod-Seed rose from 42.6 to 54.5 and 55.3 for the Prod-Sample-One and Prod-Sample-Three treatments, respectively. Mann-Whitney rank tests reveal that the Prod-Sample-One as well as the Prod-Sample-Three treatment produce significantly higher production levels than the Prod-Seed treatment. In accordance with Hypothesis 4, the difference in the production level of Prod-Sample-One and Prod-Sample-Three is far from significant.[9] The effect of sampling is also

[9] 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

apparent in the market that subjects use for their production. While only 38.3% of subjects invested in the national market in the Prod-Seed treatment, for the Prod-Sample-One and Prod-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 Prod-Sample-One and Prod-Sample-Three treatments to be below those of Prod-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 did not sample for information but rather imitated the successful.

4.1.3 The Prod-Sample-One and Prod-Sample-Three Treatments: Sampling Behavior

Here, we focus on how people sampled in the production experiment. 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 Prod-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 Prod-Seed. Overall, 63.2% of subjects sampled the top ranked person with an additional 5.3% sampling the subject ranked 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 Prod-Sample-Three treatment. Here 76.0% of all first samples were on people who chose nationally in Prod-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 Prod-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 Prod-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

are significantly different from those of the Seed at the 1% level, while they are not significantly different from each other.

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 Prod-Sample-One and Prod-Sample-Three treatments, in the Prod-Sample-One (Prod-Sample-Three) treatments 75% (76%) of subjects sampled a subject who produced for the national market. (For the Prod-Sample-Three treatment we are looking only at the first sample). Of those who sampled in the national market in the Prod-Sample-One (Prod-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 Prod-Sample-One (or first sample in the Prod-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 = Prod-Sample-Three) and d_1 = 0 if (treatment = Prod-Sample-One)

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

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

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

d_5 = 1 if (treatment = Prod-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 Prod-Sample-One and Prod-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.[10]

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.

4.1.4 Risk Taking

In the production 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 is an upperbound on the concavity of the utility function, $r<1$, to ensure that utility is increasing.[11]

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.[12] 3 of these subjects switched back 3 or more times and we drop these people from the analysis be-

[10] 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 production levels between these two groups is not significant (Mann-Whitney test: p=0.49).

[11] 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. Below, we will present evidence showing that in Sample-Three subjects' reported beliefs by and large coincide with the uniform distribution.

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

cause 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 production 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 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 Prod-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 Prod-Seed treatment 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 Prod-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 Prod-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 Prod-Sample-One and Prod-Sample-Three treatments. Here, the distributions of r_{prod} implied by production choices reveal that engaging in these treatments 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 Prod-Seed treatment we see that over 61% of them were in the local market, representing a risk averse choice, while in the Prod-Sample-One and Prod-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.

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 Prod-Seed. Figure 10 presents the cumulative frequency distribution of payoffs of our subjects in Prod-Seed and the treatments where they were allowed to sample (i.e., Prod-Sample-One and Prod-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 Prod-Seed. In fact, payoffs decrease when sampling is allowed. While the mean payoff for subjects in Prod-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 Prod-Sample-One and Prod-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.

The most direct evidence on the question whether people learn is provided by the beliefs that subjects reported in the Prod-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 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 Prod-Sample-One and the Prod-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).

4.2 Results of the Takeover Game

We first discuss the choices made by the subjects in TO-Seed. Then we discuss how subjects sampled and what bids subjects submitted after sampling.

In Seed-WC, subjects submitted on average a bid of 476.6 with a standard deviation of 216.9. The mode of the distribution is at 500: 11 subjects chose a bid close to 500. Remarkably, 6 out of 49 subjects chose the risk-neutral optimal bid of 0. As expected, though, 22 out of 49 subjects chose bids higher than or equal to 550. Many of these subjects are either found at the top or at the bottom of the ranked profit list. Table 8a presents the ranking and choices of the seed subjects in the winner's curse game together with the sampling decisions of those who sampled to which we will come back below. The person with the highest earnings was the one who chose a bid of 850. Notice that the top-3 choices were well above the average. In the winner's curse game, many subjects in the seed made a profit of 0 because their bid fell short of the value of the firm. These subjects were ranked in a random order. The bids that made 0 profits were sampled only rarely, however, which means that the order of these bids is of little importance anyway.

Tables 8a and 8b here

The results of Seed-LC are presented in Table 8b. Here, subjects chose on average a bid of 1597.1 with a standard deviation of 169.1. A total of 14 out of 49 subjects chose a bid close to 1500, which is the highest peak of the distribution. A clear majority of 29 subjects chose bids higher than or equal to 1550. The person who submitted a bid of 1900 made the highest profit. The top-3 bids were all above the average. Although most subjects chose bids quite far from the risk-neutral optimal levels of the two versions of the takeover game, they did react to the shift of the distribution in the expected direction. That is, if we substract an amount of 1010 from the bids of the TO-Seed-subjects in the loser's curse game, a distribution results that dominates the distribution of bids of the TO-Seed-subjects in the winner's curse game. A Wilcoxon rank test shows that the bids in Seed-WC are significantly lower than the shifted bids of Seed-LC ($p=0.020$).

We now turn to the question how people used the possibility of sampling in the takeover games. The person on top of the ranked list is the one that is most frequently sampled. A majority of 61.3% of the subjects in Sample-WC and 60.0% of the subjects in Sample-LC chose to observe the bid of the highest ranked person on their first or second search. Next, the person with the second highest earnings was sampled most frequently: 19.4% of the subjects in Sample-WC and 20.0% of the subjects in Sample-LC observed what this person did on one of their searches. The other searches were spread out more or less evenly over the persons with ranks 3, 5, 10, 15, 20, 25 and 49. Subjects used a very similar sampling strategy in the two versions of the takeover game, and a lot of their attention was focussed on the top-ranked persons.

In the takeover games, subjects decided how many times they wanted to sample. Virtually all subjects made use of the first sampling possibility that was provided for free. Only 2 subjects in Sample-WC indicated that they did not want to see anybody of the list. More remarkable is that a clear majority of the subjects, 58.1% in treatment Sample-WC and 70.0% in Sample-LC, paid a cost of 10 cents to make use of the possibility of observing a second person of the list. A fraction of 22.6% of the subjects in Sample-WC and 6.7% in Sample-LC were even willing to pay 50 cents to observe a third person. Thus, many subjects voluntarily pay costs to observe what other people did, which expresses a genuine preference for social learning.

Finally, we assess how sampling affected subjects' choices by comparing the bids of the TO-Seed with the bids of the people who had the opportunity to sample. First we deal with the winner's curse game. Here, subjects chose on average a bid of 559.8 after sampling, at a standard deviation of 250.7. This exceeds the average bid of the subjects in the TO-Seed of 476.6 (s.d. 216.9) by a fair margin of almost 20%. Likewise, the median of the bids in Seed-WC of 500 lies below the median of bids of Sample-WC of 600. Figure 11 displays the situation graphically and shows that the empirical density of the bids in Sample-WC has more mass on the right side than the one for Seed-WC. The difference in bids between Seed-WC and Sample-WC is significant according to a (two-sided) Mann-Whitney rank test at p=0.057. The evidence rejects hypothesis 7, in favor of the hypothesis that sampling encourages higher bids in the winner's curse game.

Figure 11 here

The effect of sampling looks more pronounced in the loser's curse game. Here, subjects submitted bids with an average of 1727.5 after sampling, at a standard deviation of 118.7. This is clearly higher than the average bid of Seed-LC of 1597.1 (at a s.d. of 118.7). The median bid of the subjects who sampled equals 1750, exceeding the median bid of the seed at 1580. Figure 12 shows a picture of the empirical densities of the bids. Notice that the mode of the distribution of bids in the seed is located at 1500, far below the mode of the distribution of bids in Sample-LC that is found at 1800. A Mann-Whitney test shows that bids in Sample-LC are significantly higher than those in Seed-LC at p=0.0001. This means that hypothesis 8 fails. Bids in the loser's curse game become substantially higher when sampling is allowed. Thus, in accordance with what can be expected if people sample for imitation, sampling reduces the loser's curse while it aggravates the winner's curse.

Figure 12 here

5 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 Production Problem 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 (as seen again in our Winner's Curse version of the Bazerman-Samuelson game) 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] Anderson, L. and C. Holt (1997) "Information Cascades in the Laboratory." *American Economic Review,* 87(5), pp. 847-62.

[3] 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.

[4] Bazerman, Max and William Samuelson, 1983. "I Won the Auction But Don't Want the Prize". *Journal of Conflict Resolution* 27, 618-634.

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

[6] Çelen, B. and S. Kariv (2004a) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review,* 94(3), pp. 484-497.

[7] Çelen, B. and S. Kariv (2004b) "Observational Learning Under Imperfect Information." *Games and Economic Behavior,* 47(1), pp. 72-86.

[8] Çelen, B. and S. Kariv (2005) "An Experimental Test of Observational Learning under Imperfect Information." *Economic Theory*, 26(3), pp. 677-699.

[9] Çelen, B. and S. Kariv, and Schotter A. (2007) "An Experimental Test of Advice and Social Learning". Mimeo, 2007.

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

[11] Charness. Gary, and Dan Levin (2006). "The Origin of the Winner's Curse: A Laboratory Study". University of California, Ohio State University. Mimeo.

[12] 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

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

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

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

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

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

[18] Holt, Charles, and Roger Sherman (1994). "The Loser's Curse". *American Economic Review* 84, 642-652.

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

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

[21] 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.

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

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

[24] 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.

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

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

[27] Selten, Reinhard, Klaus Abbink and Ricarda Cox (2005). "Learning Direction Theory and the Winner's Curse". *Experimental Economics* 8, 5-20.

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

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

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

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