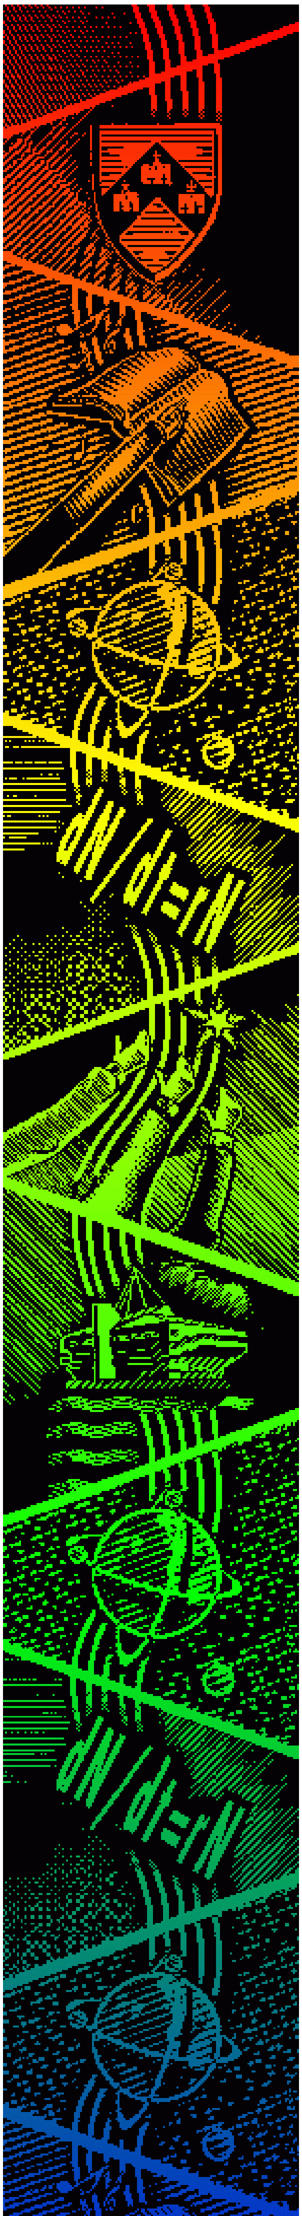

THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 2005/37

An Experimental Analysis of Optimal Renewable Resource Management: The Fishery

by

John D Hey, Tibor Neugebauer and Abdolkarim Sadrieh

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

AN EXPERIMENTAL ANALYSIS OF OPTIMAL RENEWABLE RESOURCE MANAGEMENT: THE FISHERY[‡]

John D Hey, Tibor Neugebauer and Abdolkarim Sadrieh

Universities of York and Bari, University of Hannover, and University of Magdeburg

First version: 2001-12-19

This version: 2005-05-25

Abstract

This paper experimentally studies the extraction decisions of a sole-owner in a fishery, the population dynamics of which behave according to the standard deterministic logistic growth model. Four treatments were implemented which differed in the level of information supplied to the experimental subjects. The theoretical solution was used to evaluate the behaviour of experimental subjects. The data reveal high efficiency losses due to the lack of information on population dynamics and stock size. Efficiency varied between treatments according to the information conditions.

JEL classifications

C91, D81, Q22

Keywords

Experimental economics, renewable resources, dynamic decision making, decisions under risk and uncertainty, misperceptions of feedback

Contact

John D. Hey, Department of Economics, University of York, York YO10 5DD, UK.

Tel: 0044 1904 433786. Fax: 0044 1904 433759. E-mail: jdh1@york.ac.uk

---

[‡] This paper is part of the EU-TMR Research Network ENDEAR (FMRX-CT98-0238). The authors are grateful to Kurt Schnier, seminar participants at Magdeburg, York, Kiel, Salamanca, Barcelona, Amsterdam and Akureyri for valuable comments.

1 Introduction

Renewable resources are those for which the stock can be continually replenished. Fishery resources are renewable. However, if (through human activities or otherwise) the population of some species is drawn down beyond a critical threshold, the species can become extinct. A recent concern has been with the dramatic decline in the populations of several valuable fish species such as cod, halibut and haddock. Since the seminal article of Gordon (1954), difficulties in effective management of fisheries have been attributed to the resource's peculiarity of being a common property. However, due to the new law of the sea (established in 1982) more than 90 percent of fish resources are now under the exclusive jurisdiction of coastal states and can, in principle, be protected. Distant water fishing fleets are restricted to cooperative arrangements. The coastal state has to establish a total allowable catch (henceforth TAC) for each fishery resource in its extended economic zone. The TAC is allocated among the fishermen; the individual quotas are transferable and can be reallocated through a market for certificates. In theory, established property rights and the individual transferable quota system warrant optimal resource management. In practice, however, errors might occur when the decision-maker determines the TAC, because the size, growth and population dynamics of the fishery are not exactly known. Since fish is only observable upon landings, the estimated stock size of the species is likely to be different from the actual one.

The primary research question addressed in the present study is to which extent the accuracy of stock surveys and the knowledge of the population dynamics may alter the decisions of the planner and affect efficiency of resource management. We study the resource extraction decisions of a sole owner under different information conditions in a deterministic laboratory setting. Our experimental results indicate that the knowledge of both the species' growth model and to a smaller extent the accuracy of the stock estimate may produce significant efficiency enhancements in the dynamic decision task. In fact, these significant effects are not only a consequence of the different information conditions of experimental treatments but arise also from subjects' deficiencies in learning non-linear dynamics.

The paper is organized as follows. Departing from the classical logistic growth model, we derive the finite-horizon optimal extraction plan in the subsequent (second) section. In the third section we highlight research issues and present our experimental design. In the fourth section we report the results of our study and relate them to the received literature. Finally, the fifth section concludes.

2 Theoretical Considerations

Consider the standard logistic growth function (as depicted in Figure 1), $F(x_t) = rx_t(1 - x_t/K)$, where x_t denotes stock, $r > 0$ denotes the species' intrinsic growth factor and $K > 0$ denotes the carrying capacity[1]. In the open-access fishery, the equilibrium level of the resource stock is determined by the ratio of harvesting cost to the price of the resource. If harvesting in a commercial fishery is costless (as in the experiment), the species will be extinguished for any positive price. If the intrinsic growth-factor r is smaller than the interest rate d and costs are equal to zero, extinction may be the only solution even in the optimal harvesting policy.

Figure 1. Logistic growth function for K=1000 and r=1.5

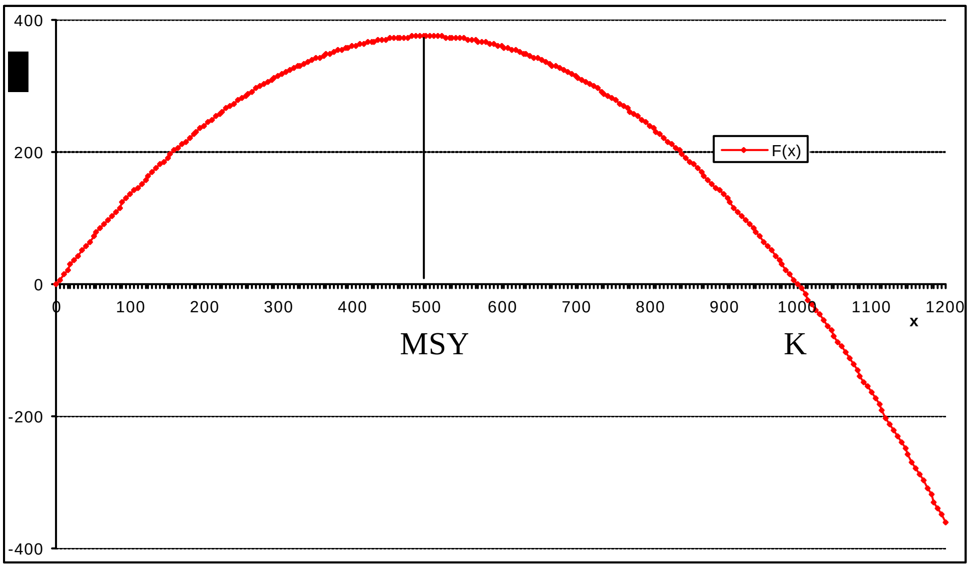

Note: MSY denotes the maximal sustainable yield.
The graph represents the experimental parameterisation.

Let the discount factor be denoted by $r = 1/(1 + \delta)$, the price normalised to one and assume harvesting costs to be equal to zero, the optimal extraction policy in the finite-horizon management problem can be determined as the solution to the following program.

$$\begin{aligned} \max \quad & V = \sum_{t=0}^{T} r^t y_t \\ s.t. \quad & x_{t+1} = z_t + F(z_t) \\ & z_t = x_t - y_t; \quad x_0 = K \end{aligned} \tag{1}$$

[1] This is the maximum viable (long-run) stock size.

Here, x_t denotes the stock before extraction, z_t denotes the stock size after extraction and y_t (the control variable) denotes the extraction in period $t \in \{1, 2,.., T\}$. The optimal solution to this problem can be calculated by means of Bellman's (1957) maximum principle. Define $J_n(x)$ as the maximum total value when only n periods remain, and the state variable at the outset of these n periods is x. Thus, beginning with the last period, the decision-maker faces the following problem.

$$J_0(x) = \max_{y_T} \{ r^T y_T \} \tag{2}$$

The final extraction y_T that maximizes the value function in equation (2) is equal to the maximal feasible y_T, which coincides with the stock remaining in period T, x_T. Hence, $J_0(x) = r^T x_T$, which, according to (1), is a function of the extraction in period T-1, $x_T = x(y_{T-1})$. Given $J_0(x)$, we can calculate the next term of the maximization procedure, $J_1(x)$.

$$J_1(x) = \max_{y_{T-1}} \{ r^{T-1} y_{T-1} + J_0(x(y_{T-1})) \} \tag{3}$$

From the first order condition follows the optimality equation $F'(z_{T-1}) = r(1 - 2\, z_{T-1}/K) = d$. Solving this equation, we obtain the end stock size $z^*_{T-1} = K/2(1 - d/r)$, which is constant as it does not depend on time. Given initial stock size x_{T-1}, the optimal extraction in period T-1 is determined by the optimal end stock size, $y^*_{T-1} = x_{T-1} - z^*_{T-1}$. Thus, $J_1(x) = r^{T-1}(x_{T-1} - z^*_{T-1}) + r^T (F(z^*_{T-1}) + z^*_{T-1})$. Proceeding by backward induction, the following general expression is determined,

$$\begin{aligned} J_n(x) &= \max_{y_{T-n}} \{ r^{T-n} y_{T-n} + J_{n-1}(x(y_{T-n})) \} \\ &= r^{T-n}(x_{T-n} - z^*) + \sum_{j=0}^{n-1} F(z^*) r^{T-j} + r^T z^* \end{aligned} \tag{4}$$

where $z^* = (1 - d/r)\, K/2$ and $y_t = \max\{x_t - z^*; 0\}$. Hence, the first term of the maximization procedure is $J_T(x) = x_0 - z^* + \sum F(z^*)\, r^{T-j} + r^T z^*$. The first extraction is determined by the initial stock size $x_0 = K$ and the optimal end stock size z^*, $y^*_0 = x_0 - z^* = K - z^*$.[2] Since the end stock size is constant for all $t < T$ and growth is deterministic, the initial stock size x_t is constant for $t > 1$ and, consequently, the extraction y_t is constant for all periods $1 < t < T$. This result holds for any finite time horizon $T < \infty$, and also in the infinite horizon management

[2] Note the optimal harvest policy is a "most rapid approach" policy, driving the population toward the optimal level z^* as rapidly as possible.

problem.[3] Hence, the extraction plan in the finite-horizon management problem coincides with the one in the infinite-horizon case (exclusive of the last period when the resource has to be extinguished) because at the maximum the marginal productivity of the resource after extraction $F'(z^*)$ must equal the interest rate δ.

3 Laboratory fisheries: Design issues and experimental procedures

Design issues

The model of the previous section (as much as any other theoretical model we can handle) is a vastly simplified representation of the fishery. The perfect description of the population dynamics and the knowledge of the exact stock-size in every instance of time are only two of the idealistic assumptions. If we relax these, the solution to the harvesting-problem -as long as we can find one at all- becomes more involved. Another serious simplification of the model is the assumption of unbounded rationality which implies that a decision-maker is able to determine the optimal catch quota within a system of non-linear dynamics. The literature has shown that subjects experience significant difficulties in non-linear environments (Sterman (1989a, b, c), Brehmer (1992), Paich and Sterman (1993), Sterman (1994), Diehl and Sterman (1995), and Moxnes (1998a, b)). As Sterman (1994) pointed out

> … human performances in dynamic (complex) systems is poor … even compared to simple decision rules. … The observed dysfunction in dynamically complex settings arises from *misperceptions of feedback*.[4] People are insensitive to non-linearity and violate basic rules of probability. The robustness of the misperception of feedback and the poor performance that lead us to create across many domains result from two basic and related deficiencies in our mental models of complexity. First, our cognitive maps of the causal structure of systems are vastly simplified compared to the complexity of the systems themselves. Second, we are unable to infer correctly the dynamics of all but the simplest causal maps.

This paper addresses the efficiency losses that might accrue in fishery management due to the decision-maker's shortcoming in dealing with complexity and due to missing information. For

[3] See Clark (1976, Ch. 2) for a derivation of a solution to the infinite-horizon problem and a discussion.

[4] Moxnes (1998a) referred to *misperceptions of bioeconomics* when he reported from a fishery management experiment.

this study, we have designed and conducted experimental treatments which vary two information conditions involving the knowledge of the species' growth model and the accuracy of the stock estimate. The complexity in the task arises through the non-linearity of the growth function. We measure the efficiency of subjects' extraction decisions by comparing the observed extractions to the maximal possible outcome. In fact, the scope of this study, in which one sole-owner of the fishery decides on the TAC, is limited to the examination of the deterministic *microworld* we described in the previous section. Therefore, many complications authorities actually face when they set the TAC are missing. Though this environment is overly simplistic it still captures essential ingredients of a fishery resource's population dynamics. Given the (relative) easy tractability of this environment, we put up with the drawbacks. More realistic settings may be studied in the future.

Still, there are at least three features with respect to the experimental implementation of the dynamic decision task that should be stressed: First, in the literature the fishery management problem is typically set in the infinite-horizon. As pointed out in the previous section, the optimal harvesting policy in the theoretical model is the same whether we consider the finite or the infinite-horizon setting. Since the infinite-horizon cannot be implemented in the laboratory, we tackle the fishery management problem as a finite-horizon dynamic decision task.[5] Second, the decision-maker's presumed objective should be to maximise the present value of the fishery in every instance. In a world without interest and costless harvest, this objective involves the most rapid approach to the maximum sustainable yield with every extraction decision, including a rebuilding of an eventually depleted resource as rapidly as possible. Naturally, the authorities can not know how well their harvesting-decisions approach the maximal economic rent. In the experiment, we implement this ignorance by a lack of information feedback into the decision-maker's payoff space.[6] Clearly, subjects must be rewarded according to their extractions. However, the exchange rate between the experimental currency and the subject's home currency must not be given before the end of the experiment. Finally, there might be an emotional decision-bias of subjects -particularly of pity- which might be associated with slaughtering of fish.[7] In order to guarantee salience of the incentive structure the experiment must be neutrally framed. In the experiment we ask a

---

[5] Given the earth does not exist indefinitely this approach does not seem less plausible, either.

[6] Apesteguia (2005) finds no behavioural differences in a common pool experiment if payoff information is not revealed.

[7] See Moxnes (1998b) for a discussion.

subject to maximize savings, which are identical to the number of extracted units on the subject's account. The procedures are detailed in the following subsection.

Experimental Procedures

In the computerized experiment,[8] a subject had to decide one hundred times on the TAC, i.e., how much to extract from a privately owned resource stock. The extracted units were saved on the subject's account and the logistic growth function was applied to the units that remained after an extraction. The initial stock size coincided with the carrying capacity $x_0=K=1000$ units, the intrinsic growth parameter was $r=1.5$, and the discount rate was $\delta=0$.

The experiment involved four treatments which differed in the level of on-screen information. In *Table 1* an overview is given: the letter *G* denotes growth, the latter *S* denotes stock and the letters *No* indicates no information on growth or stock. Before every extraction, the subject received a stock signal revealing information about the number of existing resource units. This signal was accurate in the treatments *GS* and *S* – i.e., equal to the resource stock x_t – and noisy in the other two treatments *G* and *No* – i.e., the signal was equal to the resource stock multiplied by a random draw from the uniform distribution over the interval [0.75,1.25] and rounded to the next integer. In the treatments *GS* and *G*, an on-screen facility (in *Table 1* referred to as information about the growth function) was provided by means of which a subject could anticipate the consequences of any possible extraction for the nearest future before she/he confirmed an extraction.[9] Subjects were instructed accordingly.[10]

Efficiency in the experiment was defined as the quotient of extracted units and 38125, which was the maximum number of possible extractions unknown to experimental subjects.[11] Efficiency was hence a number between zero and one. The payoff a subject received at the

---

[8] The software was programmed by means of Abbink and Sadrieh's (1995) *RatImage*.

[9] Before making an extraction decision, the subject was given an on-screen record of 11 possible extractions in 10 percentiles of the signalled stock in the first column. In the second column the corresponding after-extraction stock sizes were displayed, in the third column the resulting next stock sizes, in the fourth column the growth of the resource (i.e., the difference between the third and the second column) was displayed, and finally the savings were recorded in the fifth column. Additionally, the subject could explore the effects of every possible extraction at any point in time and before making an extraction decision -between nothing and the maximal available number of units (i.e., in *G* the maximum extraction was 4/3 * stock signal). The results of any such enquiries were displayed in a *scroll-box* appended to the standard record of possible extractions. Finally, if the subject was satisfied with the consequences of her/his latest inquiry (displayed at the end of the table) she/he confirmed it as the harvesting-decision by pressing the "extraction button."

[10] Instructions and the computer-screen (for *G*) are depicted in the Appendix.

[11] The maximum is easily calculated by applying the results from Section 2: First, extracting 500 units to reach the steady state (the maximal sustainable yield since the interest rate is zero); then, extracting 375 units (equal to the growth in the steady state); and finally, extinguishing the resource in the last decision.

end of the experiment was the product of efficiency and the premium to be paid in a treatment which was known to subjects.[12]

Table 1. Experimental treatments

	Accurate stock size		Noisy signal[a] about stock size	
Growth Function information	**GS**	#25	**G**	#35
No growth Function information	**S**	#27	**No**	#30

a) The noisy signal equals the true stock size multiplied by a random number from the interval [.75, 1.25]

If a subject extinguished the resource before having made 100 extraction decisions the experiment ended instantaneously, regardless of the number of decisions made to that point. In order to limit erroneous extractions from the stock, subjects were warned if the extracted number of units exceeded the stock signal. At the other extreme, an extraction decision of zero units also triggered a warning. In addition, before the last decision (in round 100) the subject was informed that no further extraction would be possible thereafter. The preceding extractions and the on-screen information, including the stock signal before and after extraction as well as the resulting savings, were recorded in a history-window that subjects could access at any time during the experiment.

In total 117 subjects participated in the experiment. The set of decisions made by each subject represents an independent observation for our statistical analyses. The number of subjects participating in each treatment is displayed in *Table 1*. The experimental sessions were conducted on two occasions, one at the ESSE laboratory at the University of Bari and the other at the CentERlab at Tilburg University. Each subject participated in only one treatment.

Theoretical Benchmarks

In section 2 above, we derived the optimal extraction strategy in the full information *GS* treatment. This strategy is clearly not applicable in the other treatments. Moreover, it is not clear in these other treatments what the optimal strategy is. However, in this section we

[12] The premium (i.e., the maximal payoff) in *GS* was €15, in *G* and *S* €17.50, and in *C* €20 (1€ ≈ 1$). The average payoff was €11; the experiment took about an hour.

propose 'reasonable' strategies in these other treatments, and justify their 'reasonableness' by showing that the implications of following these strategies are close to the implications of following the optimal strategy (if it were known). Of course, the subjects in our experiment could not know what was the optimal strategy, but we, the experimenters, know, and can use that knowledge to justify these 'reasonable' strategies. In what follows, we refer to these as our 'theoretical benchmarks.' That for the *GS* treatment coincides with the optimal strategy; those for the other treatments are justified in what follows.

The development of the stock on the optimal extraction path for the treatment *GS* is shown in the top panel of *Figure 2*. On first sight, it may seem inadequate to use this perfect information benchmark for calculating efficiencies in the imperfect information settings. However, the benchmarks for the two treatments *G* (top right in Figure 2) and *S* (bottom left in Figure 2), are so close to this simple benchmark that using more elaborate comparisons would not yield any substantially different results. Even the benchmark for the treatment *No* that shows substantial stochastic variation centres with a median path around the perfect information benchmark (bottom right panel in Figure 2). In other words, given subjects follow a path of action that uses the information input consistently, they are likely to come close to the perfect information optimum rather quickly in an early phase of the experiment.

In order to avoid the problem of extinction, the suggested theoretical benchmarks for the treatments without perfect information are all "prudent", i.e. extraction choices that do not lead to the extinction of the resource with certainty are preferred to those that risk extinction. While this requirement may be too conservative in general, it seems useful, because it defines the most cautious benchmarks, below which no reasonable extraction path should fall, no matter how risk-averse the decision-taker is. Interestingly, prudence does not really pose a major source of inefficiency in any of the settings. In treatment *S*, the prudent benchmark behaviour achieves 98.6 percent efficiency. While an average efficiencies of 99.7 and 85.4 percent are achieved in the treatments *G* and *No*, correspondingly.

In the treatment *S*, in which the stock size is known, but not the growth function, the decision-maker must use the information on stock size change (i.e. growth) to infer the best possible plan of action. Given that the range of possible actions is finite and given that the growth function is unknown, but single-peaked and fixed, the task can be reduced to a parameter search problem. The goal of the search algorithm is to identify the stock level inducing the

Figure 2. Theoretical benchmarks

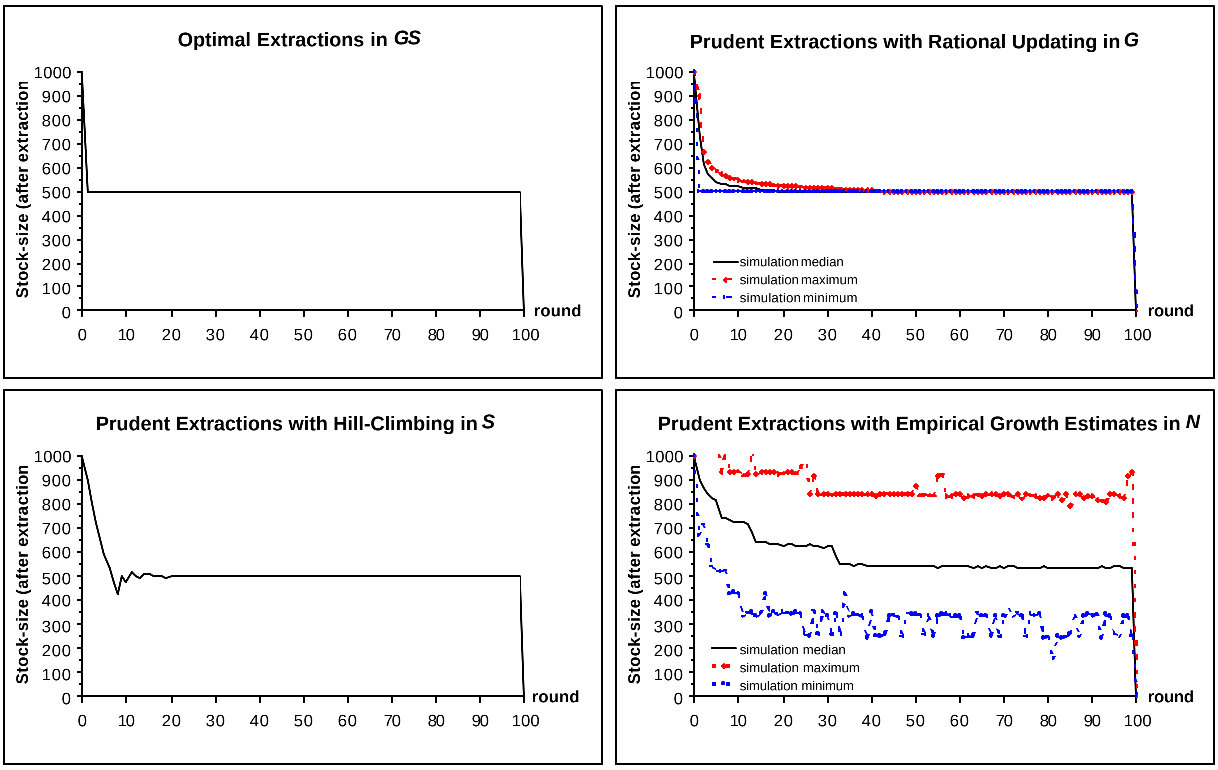

maximum growth. Since the stock size information is perfect and the growth function well-behaved, a *hill-climbing* algorithm can be used that searches the parameter space employing systematic experimentation and consistent adaptation of the choice variable to achieve higher and higher values of the goal variable. The only major difficulty that the process must deal with is the *lock-in* hazard that is due to the missing information on the growth function. A lock-in situation arises when experimentation entails the risk of "being stuck" at such a low level of growth that a return to the optimal stock size is no longer feasible within the decision horizon. In the case of an extremely skewed growth function, for example, even relatively cautious experimentation may lead to lock-in situations, on the one hand, while perfectly conservative stock preservation will obstruct the optimisation process, on the other. Hence, the decision-maker will have to trade-off the efficacy of the search process against the risk of being locked in at a sub-optimal stock level.

The benchmark we present in the lower left panel of figure 2 uses a simple hill-climbing search algorithm with an exponentially decreasing experimentation rate ε_t. In any round t, the decision-maker extracts an amount that leaves $1 - \varepsilon_t$ of the last round's post-extraction stock level. As long as the observed absolute growth in t is greater than in $t - 1$, the process is continued with the same experimentation rate ε_t, i.e. $\varepsilon_{t+1} = \varepsilon_t$. As soon as, a decline of the resource growth is observed in a round t, extraction is adjusted to restore the previous stock size, before continuing experimentation at an halved rate, i.e. at the rate $\varepsilon_t = \varepsilon_{t-1}/2$. At what speed the process converges and how difficult it is to recover from local lock-ins, does not only depend on the growth function, but also on the initial conditions, i.e. the initial size of the stock and the initial experimentation rate ε_0.[13] Using the experimental parameters and an initial experimentation rate of .1, we can show that after only 9 of 100 rounds, the process converges to stock levels that are within 10 points around the perfect information benchmark. By round 20 experimentation ends and the processes rests exactly at the optimal stock level.

Things are a bit more complicated in the *G* treatment, in which the growth function is known, but not the exact stock size. In this case, using the growth function information, the decision-maker can calculate the optimal target stock size, which is identical to the target stock size in the perfect information treatment. However, due to the stochastic nature of stock size feedback, extraction decisions that perfectly hit the targeted optimal stock size cannot be made. Instead, to improve the quality of the extraction decisions, the information arriving after each decision must be used to increase the precision of the stock size estimate. In any round, the extraction history and stock size signals can be combined with the growth function information to tighten the lower und upper boundaries on the initial stock size. As more and more observations are made, the range of possible initial stock sizes is reduced, ultimately making a perfect estimate possible. Once the initial stock size can be pinpointed, the current stock size can be calculated by reconstructing the history of extractions and applying the growth function, correspondingly.

While the inference logic described above is unique, neither the realised path of information disclosure nor the level of extractions up to the point of perfect inference are unambiguous. The path of inference is not unique, because of the stochastic nature of the stock size feedback. Obviously, certain sequences of random draws will enable a quicker perfect inference than other sequences. Furthermore, the optimal extraction behaviour *before* the perfect inference is achieved depends on how the threat of pre-mature extinction is treated.

We have chosen a benchmark that deals with the extinction issue by assuming "prudence" in the sense that any pre-mature extinction of the stock is excluded. The prudent extraction x_t in

[13] If ε_0 is small the risk of overshooting the optimal growth stock level is small, but the speed of convergence is low. If ε_0 is large then the contrary is true. The path displayed in Figure 2 is derived for an initial stock size of 1000 and an initial experimentation rate of $\varepsilon_0 = .1$.

round t is limited to being no greater than the minimum estimated stock size $\underline{s}_t$ at time t (i.e. given all information collected in the previous rounds), hence $x_t = \max(0, \underline{s}_t - 500)$, where 500 is the optimal stock size (derived from the growth function information).

The top right panel in *Figure 2* displays the development of the stock in a small Monte-Carlo sample of runs with *prudent extraction and rational updating*. The "median path" shown in the panel, describes which development of the stock size we should be expecting, if subjects are prudently extracting and rationally updating. The "minimum" and "maximum" paths show the extremes of the simulated distribution[14]. As can be seen, the stock size quickly converges to the optimum of 500 (i.e. the exact initial stock size is quickly inferred from the history), even though the assumed behaviour is very cautious concerning the threat of pre-mature extinction. On the median path, the maximum sustainable yield at a stock size of 500 is reached after only 20 of 100 rounds. Even in the worst case observed in the simulation, no more than 40 rounds were needed for full convergence.

What is perhaps even more important than the point of full convergence is the fact that the path comes close to optimum very quickly and hence induces only minor losses due to the imperfect information. On the median path the total extraction is just slightly below 38000 compared to the 38125 in the optimum of the perfect information setting. Hence, the median efficiency loss due to application of the prudent extraction rule would be less than 0.4% and even in the worst case only 1%.[15]

Finally, defining a convergent benchmark in the treatment *No* involves using blending the two methods used for the benchmarks in *S* and *G*, because both growth function information and perfect stock size information are missing. The bottom right panel in *Figure 2*, displays the median, the minimum, and the maximum path that were observed in a small Monte-Carlo simulation using such a combined procedure. In this procedure, the hill-climbing process (i.e. the search for the stock size that induces maximum growth) cannot be controlled by simply comparing resource growth at different levels of stock, because the stock size information is imperfect. Instead, the success measurement has to be based on the distribution of empirically estimated growth numbers. Since the growth function information is also unavailable, achieving the same precision of the empirical estimation of the growth numbers as in

[14] It should be noted that these do not represent actual paths, but just the upper and lower envelopes of the various possible paths.

[15] Since subject payments in the experiment were rounded up to multiples of 50 Cents, even in the worst case simulation subject payments would have coincided with the maximal payoff.

treatment *G* requires having many more observations. Hence, the low level of information in *No* leads to a high dispersion in the speed and path of convergence. As the minimum and maximum paths we observed in our simulation show, often 100 rounds will not entail enough empirical observations as to allow a convergence of the process to the maximum growth point. Nevertheless, the median simulation path converges well within the first half of the experiment, indicating that the distribution of experimentally observed post-extraction stock sizes in the second half of the experiment should be located around 500, the stock level that induces maximum growth.

4 Experimental Results

This section is organized corresponding to the optimal extraction plan. First, we survey the efficiency of initial extraction decisions; second, we consider the overall efficiency and the evolution of extraction decisions; and, third we report on the efficiency of subjects' last extractions. As a benchmark we refer to the decisions on the optimal path. These imply a stock-size after extraction at the maximum sustainable yield (i.e., 500 units) until pen-ultimate decision and extinction of the resource with the last decision. We conclude the section by classifying observed individual behavioural pattern.

The First Extraction Decisions

The first extraction induced significant under-harvesting in all treatments (two-tailed Wilcoxon signed ranks test at α=.01): subjects extracted less than the optimal 500 units. Table 2 records the statistics on stock after the first extraction.

Table 2. Stock size after first extraction

treatment	#	minimum	Maximum	average	std. error	Wilcoxon-test H_0: average=500
GS	25	300	975	653	165	-3.40^{**}
G	35	280	999	670	197	-3.98^{**}
S	31	100	1000	842	232	-4.34^{**}
No	30	200	1000	932	166	-4.69^{**}

**significant at 1%, one-tailed; *significant at 5%, one-tailed

The deviations from the optimal extraction increased from treatment *GS* through *No* (two-tailed Jonckheere test of ordered alternatives at α=.01). Equation (5) represents a pooled dummy regression of the distance of stock after the first extraction from the optimum on *G* and *S*. In accordance with the Jonckheere test, the regression results reveal that growth and accurate stock information both had a significant influence on the efficiency of subjects' first extraction decisions. D*G* and D*S* denote dummy variables which take a value of one if a subject receives growth information and accurate stock information, respectively, and zero otherwise.

$\lvert Stock_1 - 500 \rvert$	=	438^{**}	$- 223\ DG^{**}$	$- 57\ DS^{*}$	$\hat{R}^2 = .35$	(5)
		23.96	27.64	27.71	[std. error]	
		18.30	-8.08	-2.07	[t-ratio]	

**significant at 1%, one-tailed; *significant at 5%, one-tailed

Average Efficiency

The subjects' presumed objective in the experiment was to maximise efficiency, which we define as the ratio of the actual extraction to the maximal possible one. *Table 3* records the minimum, maximum, and average efficiency attained in the experiment. Standard deviations are reported in the last column.[16] Efficiency increased across treatments from *No* to *GS*. Differences between treatments are significant at 1% for all pair-wise comparisons, except for the comparison of *G* to *S* (Mann-Whitney test, two-tailed). In treatment *G* (and only in treatment *G*), three subjects extinguish the resource within the first 19 extractions (see *Table A* in the appendix). If we do not take the extinction observations into account, efficiency in treatment *G* is significantly greater than in *S*. The dummy regression of efficiency on the treatment dummies growth and accurate stock reported in equation (6) indicates that both treatment variables had a significant effect on efficiency. The knowledge of the growth function implied an increase of average efficiency by 27.6%, the accurate stock size information by 21%.

Efficiency	=	$.427^{**}$	$+ .276\ DG^{**}$	$+ .210\ DS^{*}$	$\hat{R}^2 = .31$	(6)
		.038	.044	.044	[std. error]	
		11.17	6.26	4.66	[t-ratio]	

**significant at 1%, one-tailed; *significant at 5%, one-tailed

[16] *Table A* in the appendix records individual efficiency levels.

The maximum of the observed efficiency levels does not deviate too much from the efficiency levels proposed by the theoretical benchmarks in any treatment. However, actual behavioural patterns differed much from the benchmark; they will be discussed below.

Table 3. Efficiency

Treatment	#	minimum	Maximum	average	std. error
GS	25	0.613	0.999	0.874	0.108
G	35	0.091	0.949	0.727	0.254
S	31	0.057	0.922	0.661	0.252
No	30	0.018	0.853	0.398	0.287

The Evolution of Extractions and Stock

Figure 3 contrasts the evolution of average stock levels after extraction in all treatments with the optimal path represented by the 500-units line. Payoff maximization involved in each but the 100th decision an extraction of the maximum of zero units and the difference of the actual stock size and 500 units. In case the resource was depleted below 500 units the stock would have to be rebuilt. Equation (7) represents the pooled regression of distance from optimum on time, by which we checked whether efficiency increased over the sequence of extraction decisions.

$$|Stock_{it} - Optimum| = \beta_0 + \beta_1{}' t + \varepsilon_{it},\ 1 \leq t \leq 99 \qquad (7)$$

The regression results recorded in *Table 4* confirm that efficiency increased in the experiment. The coefficients indicate that the effect of time on efficiency was greater in treatments *S* and *No* than in the treatments *GS* and *G*, in which subjects received information about the growth function. This result does not say that subjects in the treatments without growth information were more efficient in the end than those who had this information. Convergence of the average stock after extraction to the optimum was limited at 205 and 258 units in the treatments *S* and *No*, respectively. In the treatments *GS* and *G*, in contrast, the average deviation from the optimum did never exceed 199 and 252 units, respectively.

Figure 3. Evolution of average stock after extraction compared to optimal path

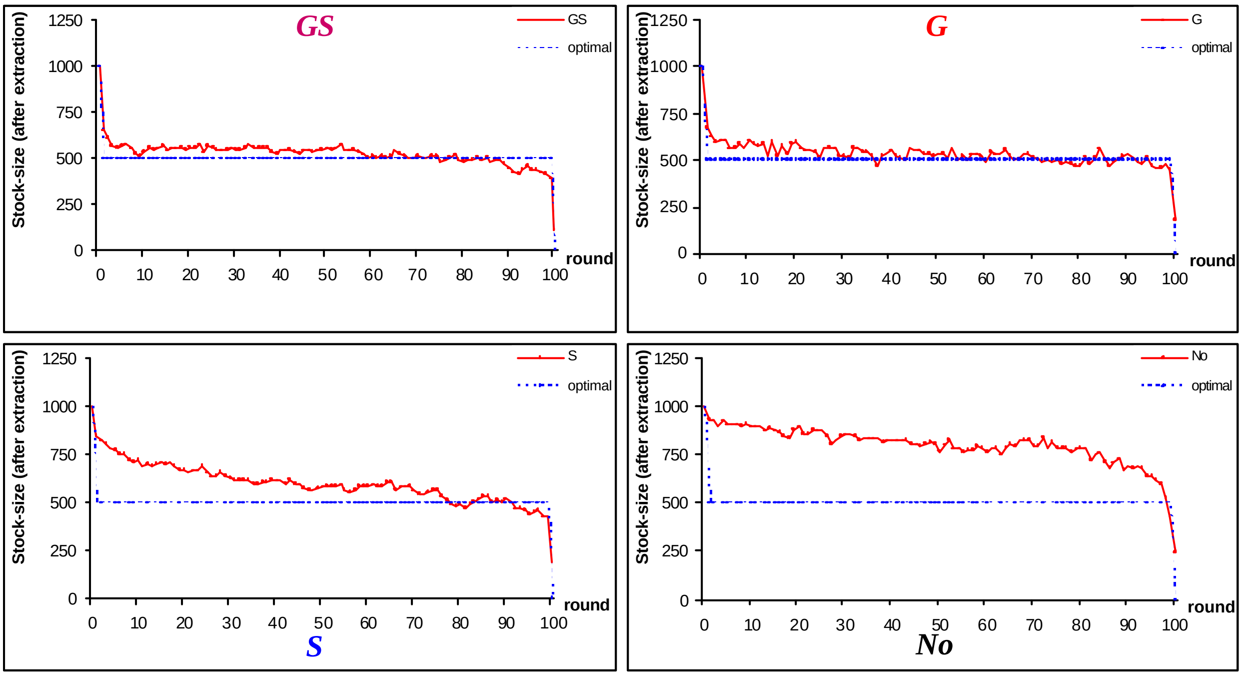

Table 4. Distance from optimum on time (eq. 7)

		coefficient	std. err.	t-ratio	$\hat{R}^2$
GS	constant	182**	6.220	29.33	0.03
	period	-0.67**	0.108	-6.22	
G	constant	197**	6.063	32.55	0.01
	period	-0.53**	0.105	-5.07	
S	constant	248**	5.708	43.53	0.08
	period	-1.22**	0.099	-12.36	
No	constant	417**	7.516	55.47	0.05
	period	-1.31**	0.130	-10.10	

**significant at 1%, one-tailed; *significant at 5%, one-tailed

Indeed, greater efficiency also corresponds to smaller stock sizes across treatments, as, on average, our experimental data suggest rather under-harvesting than over-harvesting. This is confirmed by the two-tailed Wilcoxon signed ranks test reported in *Table 5*. The first column of Table 5 exhibits the percentage of subjects whose stock size after extraction was more frequently below than above the optimum; the second column records the ratio of observations that involved lower stock than optimum to the number of observations in which stock was not equal to optimum. Over-harvesting was most heavy in treatment *G*, where three subjects extinguished the resource within 19 decisions. In all three observations, the last extraction did not exceed the signalled stock, but in fact it did exceed the actual stock size. The two-tailed Wilcoxon signed ranks test reported in the third column of *Table 5* reveals that over-harvesting was not significant in any treatment.[17] The only significant result that comes out from the test indicates under-harvesting in treatment *No*. This reveals important differences to the results of Moxnes (1998a, 1998b). While Moxnes finds a strong tendency to over-harvesting in early rounds, we rather find a tendency to under-harvesting. This seems to indicate that the behavioural biases in such dynamic settings are not very robust concerning changes to the experimental set-up which is also reflected in the treatment differences we observe. The propensity to harvest more in *GS* and *G* than in *S* and *No* suggests furthermore that subjects are more confident with their extraction decisions when they receive information of the stock dynamics.

Table 5. Over-harvesting

	stock-size < optimum		Wilcoxon test
Treatment	#subjects	#observations	H_0: #subjects=50%
GS	40.0%	45.8%	-0.700
G	51.4%	53.1%	-0.532
S	35.5%	36.9%	-1.925
No	10.0%	15.4%	-4.502**

**significant at 1%, one-tailed; *significant at 5%, two-tailed

The Final Extraction Decision & Extinction of the Resource

With the final extraction, subjects were expected to extinguish the resource. However, only one half of them did so as surveyed in *Table 6*. 58 subjects (50%) ended the experiment with a zero stock, 10 subjects in treatments *G* and *No* did not extinguish the resource but extracted all units signalled to them before the last decision. Apparently they had forgotten that the signal was most likely incorrect. In *Table 6*, the numbers in brackets indicate the subjects who did not even extract all the units signalled to them. On the other hand, there were 12 subjects

[17] We computed for each subject the number of times the stock was greater than or equal to the optimum. The rest, the difference between 99 and that sum, revealed the number of occasions in which the stock was smaller than the optimum. For each subject we computed the surplus of number of times the stock was greater than the optimum over the number of times it was smaller than the optimum. The test was run on these surpluses.

who extinguished the resource too early, i.e. before they reached the 100th decision: eight subjects did so in the 98th and 99th decision in treatment *S* and one subject in the 93rd decision of *GS*. As already pointed out above, we observed three cases of apparently unintentional extinction in treatment *G* within the first 19 decisions.

Table 6. Non-extinction [positive end-signal]

Treatment	#	%
GS	10	40%
G	22 [17]	63%
S	16	59%
No	15 [10]	50%
Total	59 [49]	50%

Behavioural pattern: Control Theory, Linear World and Misperceptions of Feedback

In agreement with Edwards' (1962) classical description, the present work contributes to the laboratory studies on dynamic decision making.[18] Brehmer (1992) suggests that experiments on dynamic decision making are particularly valuable since real world problems such as company management or even everyday life involve many dynamic tasks, and field data is difficult to obtain. As a general framework for the study of dynamic decision making, Brehmer (1992) proposed control theory (although not the mathematical term).[19] He pointed out, subjects' *overall goal* in a dynamic decision task should be one of "… *achieving control*: that is, that decisions are made to achieve some desired state of affairs, or to keep a system in some desired state."

As we observe literally no incidence of individual decision making in support of our above outlined theoretical benchmarks, we establish the alternative research hypothesis that subjects either try to hold the stock signal constant or the extraction level (through the extractions 2-99). This hypothesis is based on the idea that subjects try to take control over the dynamic system. Actually, we can find support for both extraction policies. In *Figure 4*, we have plotted the individual stock development of four subjects who exhibit behaviour that can be

[18] A dynamic decision problem implies that 1) a series of decisions is required to reach the goal, 2) the decisions are not independent, and 3) the state of the decision problem changes. See Brehmer (1992) for a discussion.

[19] This was noted before; see for instance Rapoport (1975).

identified as “typical” for the control theory hypothesis. The number of observations that follow similar patterns is stated. For instance, 15 individual charts or 60% of the observations in *GS* display straight lines as presented in *Figure 4*.

In treatments *GS* and *S*, in which subjects received accurate stock information, it is difficult to tell whether subjects were maintaining stock or extraction as both variables depend on each other. However, these questions can be addressed by examining the plots of treatments *G* and *No*, where such information was not supplied. Next to the optimal path indication, these plots exhibit two further lines: The unbroken line represents the movement of stock after extraction and the dashed line represents the noisy stock signal after extraction. The displayed plots represent 34% and 7% of similar patterns in the treatments *G* and *No*, respectively. In the representative plot of treatment *G* the dotted line is straight indicating extractions to maintain a constant stock signal. In contrast to this, the straight segments in treatment *No* exhibit a constant stock after extraction which hints at a policy of constant extraction.

Figure 4. Behavioural pattern – control theory

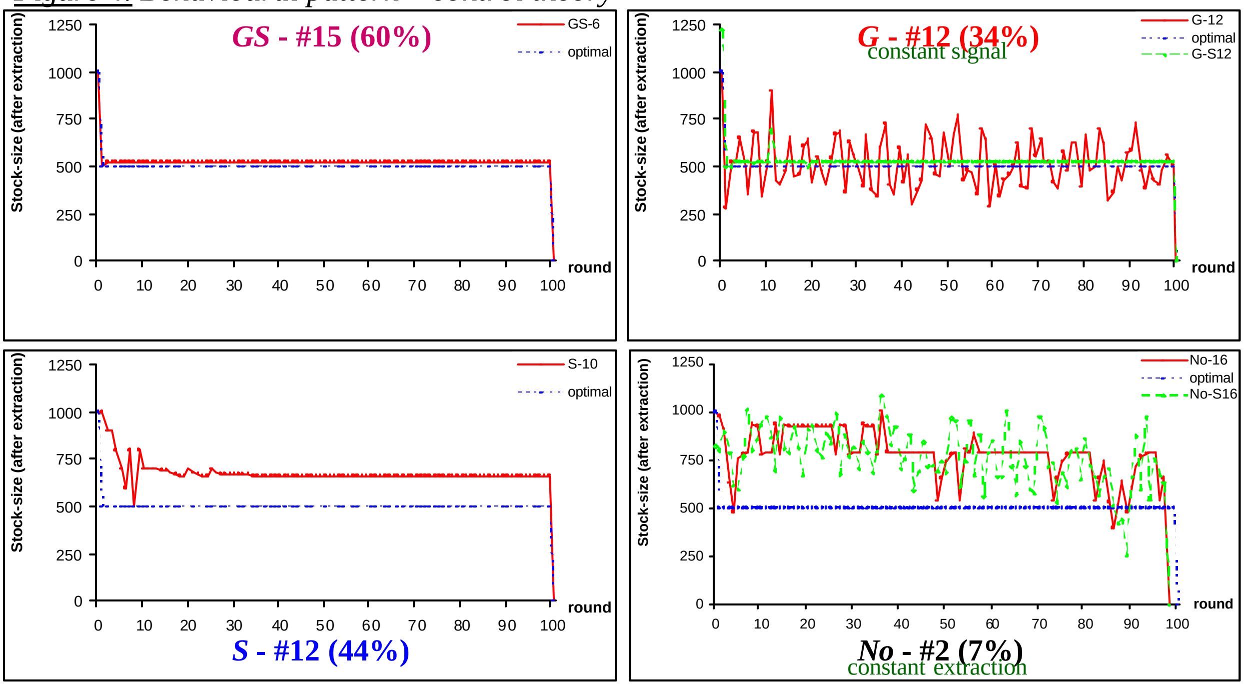

In fact, more plots of individual extraction decisions indicate a constant stock size in treatments *S* and *No*, but not in support of the control theory story. Samples of these are displayed in *Figure 5*. The striking pattern is that half of the subjects in treatment *No* and 19% in *S* extracted almost nothing. They held their stocks near the biological equilibrium size of 1000 units where growth was very close to zero. This odd behaviour can hardly be rationalized if not in the light of Brehmer’s (1980) observation that people tend to believe in a linear model rather than in other models. If subjects actually believed in a linear relationship between stock size and growth they might have taken for granted that growth increases with stock. From this perspective it would make sense to let stock grow and extract at the end the profit maximizing stock size.

Figure 5. Behavioural pattern – linear world

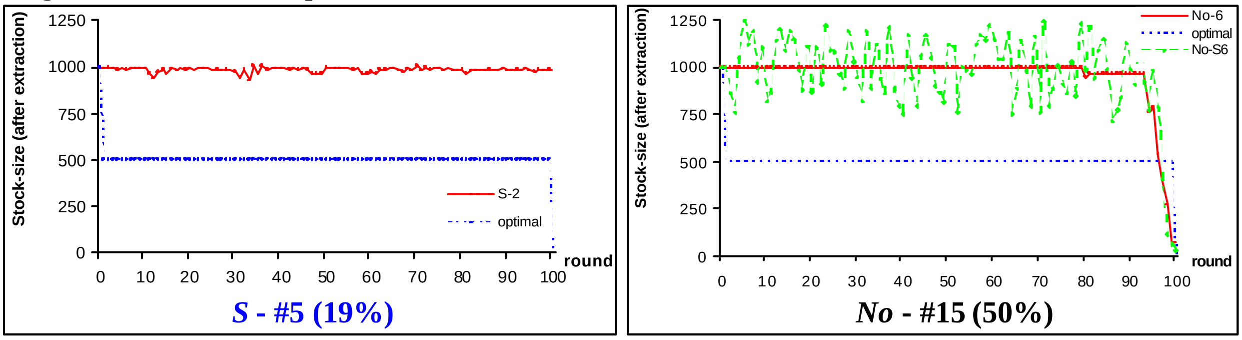

Such misperception of linearity in non-linear dynamic systems has been reported in earlier experimental research (see Sterman (1994) for a survey)), as we pointed out in section 3. Another behavioural pattern, which Sterman (1989a, b, 1994) called the *misperceptions of feedback*, must be seen in the fluctuations of the stock-sizes after extractions (see *Figure 6*). Such "pulse fishing" makes sense in some fishery environments (Schnier 2005), but not in our experiment since subjects were informed about facing a deterministic system. It seems particularly surprising that even in the transparent setting of treatment GS (in which subjects experienced feedforward information) cycles and oscillations of stock occurred. Paich and Sterman (1993) who observed comparable pattern claimed that subjects' learning in complex environments is poor. This argument could in fact explain the persistency of these oscillations in the data.

5 Summary

In this paper we have considered the fishery management problem under a finite-horizon condition. We established the benchmark solution (in the full information treatment) which disagrees with the infinite-horizon solution only in the last extraction. This approach differs from other fishery experiments as in Moxnes (1998a), Mason and Phillips (1997) or Schnier

Figure 6. Behavioural pattern – pulse fishing

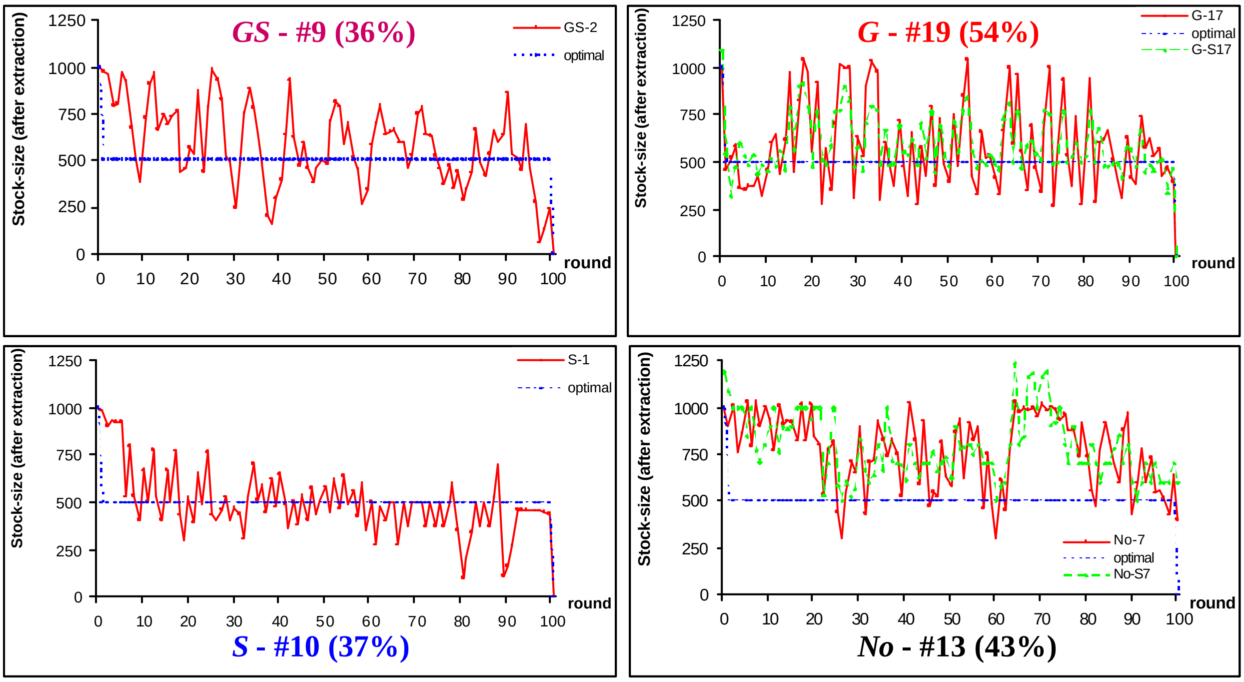

(2005) where "infinite" horizon tasks were intended but sessions lasted only 20 or 35 extraction periods.[20] In line with this literature, extinction of the resource before the end of the experiment was generally not a problem in our study, although it eventually happened when subjects received information on population dynamics but a noisy signal of stock. Over-harvesting was not a problem either in our study. This result diverges from those of Moxnes (1998a, b) and may lead back to the differences in the experimental structure: in the present study, subjects had a direct control over the fishery resource while settings in Moxnes' studies were much more complex. The behavioural patterns that classify our data almost completely can be summarised as follows. About 35% of subjects tried to achieve control over the complex system by holding either stock or extraction levels constant. Another 44% of subjects managed their stocks by pulse fishing and 17% misperceived the non-linearity of the environment and extracted almost nothing.

[20] Mason and Phillips (1996) considered a dynamic extraction game in which they varied the number of extractors from the common pool. Moxnes (1998) reported an experiment in which subjects acted as sole-owners of the fishery. The optimal solution to the management problem could only be determined numerically. The task was as of infinite horizon: Subjects had to maximise over the horizon of 20 periods the sum of extractions and remaining fish units. The control variables were the orders of vessels and utilization level of the fleet. Every vessel had an average lifetime of 25 periods, thus that all vessels should be ordered at the beginning of the experiment. Only the subject with the high score received any price. Moxnes (1998) reported over-capacities of the fleet and so over-harvesting.

Efficiency of extraction decisions was an estimated 21% higher if the stock signal was accurate, and 27.7% higher if the growth function was revealed to the subjects. It suggests that research helps to increase extraction decisions significantly. However, it should be noted that we considered here a highly simplified, deterministic model in which the precise growth function is given or not. In a real world resource management problem the decision maker faces an inaccurate growth model, non-deterministic stocks and positive market parameters as interest rates, costs, and prices. Furthermore, we are aware that political influences may affect the decisions of the authority as well (e.g., lobbyism) but we left these imperfections in the decision making process out of focus. However, all these complications might be introduced to the laboratory in the future. The logistic growth function which we considered in the experiment seems to be ideally behaved to provide the experimenter with a rich research environment: the unique optimal solution to the maximization problem can be calculated, although it is not easily recognized by experimental subjects.

6 References

Abbink, K. and A. Sadrieh, 1995, "Ratimage, research assistance toolbox for computer-aided human behavior experiments," *Discussion paper* B-325, University of Bonn.

Apesteguia, J., 2005, "Does Information Matter in the Commons?" *Journal of Economic Behavior and Organization* (in press).

Bellman, R., 1957, *Dynamic Programming*, Princeton University Press, Princeton.

Brehmer, B., 1989, "In One Word: Not From Experience," *Acta Psychologica* 45: 223-41.

Brehmer, B., 1992, "Dynamic Decision Making: Human Control of Complex Systems," *Acta Psychologica* 81: 211-41.

Clark, C. W., 1976, *Mathematical Bioeconomics*, John Wiley & Sons, New York.

Diehl, E., and J. D. Sterman, 1995, "Effects of Feedback Complexity on Dynamic Decision Making," *Organizational Behavior and Human Decision Processes* 62 (2): 198-215.

Gordon, S. H., 1954, "The Economic Theory of a Common Property Resource: The Fishery," *Journal of Political Economy* 62: 124-142.

Mason, C. F. and O. R. Phillips, 1997, "Mitigating the tragedy of the commons through cooperation: an experimental evaluation," *Journal of Environmental Economics and Management* 32: 148-172.

Moxnes, E., 1998a, "Not Only the Tragedy of the Commons: Misperceptions of Bioeconomics," *Management Science* 44: 1234-1248.

Moxnes, E., 1998b, "Overexploitation of Renewable Resources: The Role of Misperceptions," *Journal of Economic Behavior and Organization* 37: 107-127.

Paich, M. and J. D. Sterman, 1993, "Boom, bust, and failure to learn in experimental markets," *Management Science* 39 (12): 1439-1458.

Rapoport, A., 1975, "Research Paradigms for the Study of Dynamic Decision Behavior," in Wendt, D. and C. Vlek (eds.), *Utility, Probability and Human Decision Making*: 349-369, Reidel, Dordrecht.

Schnier, Kurt, 2005, "Decision Making in Patchy Resource Environments: An Experimental

Analysis" *Journal of Economic Behavior and Organization* (in press).

Sterman, J. D., 1989a, "Deterministic Chaos in an experimental economic system," *Journal of Economic Behavior and Organization* 12: 1-28.

Sterman, J. D., 1989b, "Misperception of feedback in dynamic decision making," *Organizational Behavior and Human Decision Processes* 43 (3): 301-335.

Sterman, J. D., 1989c, "Modelling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment," *Management Science* 35 (3): 321-339.

Sterman, J. D., 1994, "Learning In and About Complex Systems," *System Dynamics Review* 10 (2-3): 291-330.

Table A. individual efficiency

#	No	S	G	GS
1	0.018	0.057	0.091†	0.613
2	0.036	0.066	0.095†	0.620
3	0.045	0.168	0.149†	0.706
4	0.056	0.192	0.394	0.740
5	0.060	0.358	0.434	0.803
6	0.084	0.484	0.488	0.810
7	0.085	0.554	0.503	0.820
8	0.155	0.560	0.507	0.853
9	0.160	0.606	0.520	0.877
10	0.178	0.640	0.583	0.879
11	0.190	0.646	0.651	0.880
12	0.220	0.657	0.656	0.893
13	0.220	0.658	0.783	0.900
14	0.267	0.669	0.800	0.907
15	0.275	0.697	0.814	0.917
16	0.390	0.743	0.836	0.920
17	0.485	0.760	0.853	0.939
18	0.557	0.771	0.860	0.942
19	0.611	0.771	0.866	0.950
20	0.613	0.791	0.869	0.967
21	0.625	0.823	0.869	0.967
22	0.640	0.834	0.874	0.974
23	0.645	0.846	0.890	0.980
24	0.650	0.851	0.897	0.993
25	0.686	0.863	0.903	1.000
26	0.695	0.869	0.904	
27	0.755	0.897	0.907	
28	0.835	0.903	0.914	
29	0.845	0.909	0.926	
30	0.853	0.914	0.926	
31		0.922	0.931	
32			0.937	
33			0.937	
34			0.937	
35			0.949	

Note: Subjects' results are arranged according to their performance. †In G, three subjects extinguished the resource within the first 19 extractions. The last extraction was two or five units smaller than the signalled stock but exceeded the actual stock size.

Instructions

In the experiment you are asked to make saving decisions. With every decision you determine how many units you extract from a fictitious resource stock. Every extracted unit is credited to your savings account, which is displayed on your screen (in a window labelled “status”). Your objective in the experiment is to maximize savings. You begin with zero savings.

With each extraction you transfer units from your stock to your savings account. After the decision, the stock will be subject to deterministic growth. That is, the resource stock grows by an amount that is unequivocally determined by the number of units that remain after extraction. If the stock is zero, growth is zero. Unless you extract the entire stock you are asked to make 100 extraction decisions.

The stock size information

At every time before you make an extraction decision, the stock, i.e., the number of units from which you can extract, will be revealed to you on the screen. [subjects in **G** and **No** read: Yet, this information is biased. Your information reflects the product of a random number in the range 0.75-1.25 and the actual stock. In other words, the number of units you have in the stock is multiplied by a randomly determined number between 75 percent and 125 percent. The computer determines a new random number after each of your decisions. Consequently, you never know whether the actual stock is greater, smaller or equal to the revealed one.]

[subjects in **GS** and **G** read: The growth function

You are given information about the relation of stock size and growth through an onscreen tool in a window titled “result calculation”. It is easy to handle: Insert a potential number of units to be extracted (how to do it is detailed below). The corresponding stock after extraction and the resulting stock from which you can extract at your next decision will be stated in the second and the third column. The fourth and the fifth column record the corresponding growth and the savings after extraction, respectively. By default, this information is recorded for all potential extractions involving 10 percentiles (i.e., 10%, 20%,…, 100%) of the [subjects of **G** read: revealed] stock, as recorded in the first column of the result calculation window.]

Your payoff

There is an optimal extraction plan, though you will not be told any details about it. However, your payoff relates to the maximum possible amount of savings as follows. At the end of the experiment your payoff will depend on the quotient of your actual savings and the maximum possible savings (i.e., the quotient corresponds to the result of dividing your savings by the maximum possible ones). This quotient will be taken times {(subjects in **GS** read 15), (subjects in **G** and **S** read 17.50), (subjects in **No** read 20)} Euro to determine your payoff, which will be paid to you in private as soon as you have taken your last decision in the experiment.

The software

To make your decision you proceed in 2 steps: First, insert a potential number of units to be extracted with the keyboard or the mouse, and confirm it with the enter key. The number will be highlighted in the display of the "decision" window on the bottom right of your screen. Second, to make your extraction decision final you press the button labelled "extract". Note, unless you press the extraction button with the mouse you can insert other numbers as often as you like without any consequences.

The history

From the menu bar at the top left of your screen you can retrieve the "history" window. The history records all information you have received and the decisions you have taken in the experiment.

The screen