Environ Resource Econ (2009) 44:263–285
DOI 10.1007/s10640-009-9285-5

An Experimental Analysis of Optimal Renewable Resource Management: The Fishery

John D. Hey · Tibor Neugebauer · Abdolkarim Sadrieh

Accepted: 26 March 2009 / Published online: 9 May 2009

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 subjects. Compared to the theoretic benchmark, the data reveal that efficiency losses increase as the information on population dynamics and stock size deteriorates. Three common patterns of behaviour are identified. The distribution of these patterns is significantly affected by the informational setting.

Keywords Experimental economics · Renewable resources · Dynamic decision making · Decisions under risk and uncertainty · Misperceptions of feedback

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

J. D. Hey
Department of Economics and Related Studies, University of York, York YO10 5DD, UK
e-mail: jdh1@york.ac.uk

T. Neugebauer
Luxembourg School of Finance, University of Luxembourg, 4, rue Albert Borschette, 1246 Luxembourg, Luxembourg
e-mail: Tibor.Neugebauer@uni.lu

A. Sadrieh (✉)
Faculty of Economics and Management, University of Magdeburg, Postbox 4120, 39016 Magdeburg, Germany
e-mail: sadrieh@ovgu.de

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% 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 an 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 importance of errors in measurement or assessments of the stock level for harvesting policies has been stressed in the recent literature. Moxnes (2003) pointed out that due to the quota management, applied in practice, the implications of measurement errors for harvesting policies are much stronger than, for instance, of natural stochastic variations which have been extensively studied in the theoretical literature (e.g., Clark 1990).

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 in absence of the commons problem 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 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 thus contributes also to the growing experimental evidence on misperceptions in dynamic decision making problems (for a survey see Rouwette et al. 2004).

Within the scope of resource extraction decisions other contributions to the literature have involved problems of optimal-sized fishing fleets (Moxnes 1998a, 2000; Schnier and Anderson 2006), optimal use of reindeer rangelands (Moxnes 1998b, 2000, 2004) and optimal harvesting in a multiple species fishery (Brekke and Moxnes 2003). Most of these studies involve more complex settings than ours and do not allow an exact measurement of the efficiency losses due to errors in measurement or assessment. An exact measurement is usually not possible, because the efficient solutions are ex-ante undetermined.

In contrast, we use a single-species model featuring logistic growth for which the optimal extraction path is uniquely determined. The model is the standard in lecture books and we adapt it to laboratory conditions. Subjects are introduced to a finite-horizon, neutrally framed, deterministic decision problem in discrete time and are motivated by salient rewards to maximise efficiency under risk and ambiguity. We observe that risk regarding the stock size and ambiguity about the growth function have both an extensive and significant effect on the efficiency of the individual extraction decisions. Most of our subjects follow one of three typical extraction patterns. The decisions of about 34% of our subjects seem to aim at control. They either try to achieve a desired constant stock level or a desired constant harvest. The behaviour of 45% displays oscillations that suggest pulse fishing that is characterised by periods of harvest and periods of recovery. Finally, 17% of our subjects extract only minimal amounts or resource, possibly due to a linear mental model of growth.

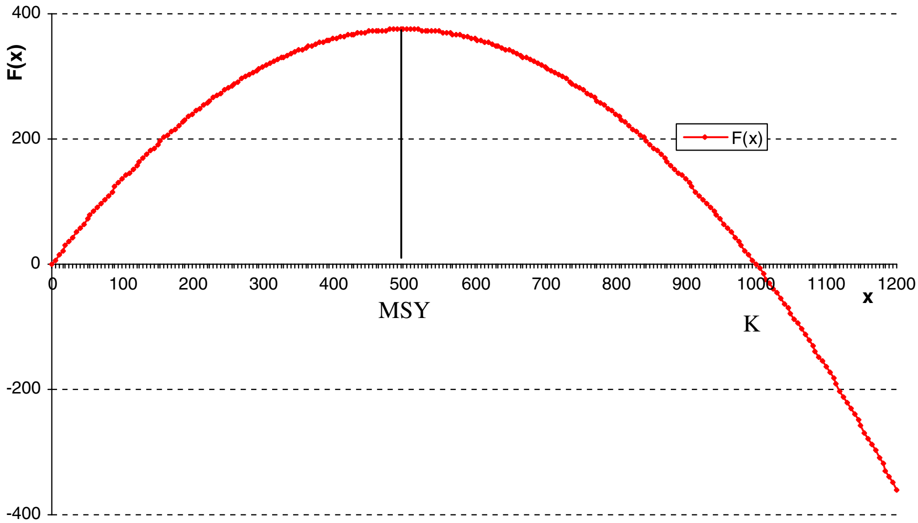

Fig. 1 Logistic growth function for $K = 1000$ and $r = 1.5$. *Note*: MSY denotes the maximal sustainable yield. The graph represents the experimental parameterisation

While some of these behavioural patterns have been reported in the literature,[1] our results show that the informational setting significantly affects the distribution of the behavioural patterns. Patterns of control are mainly observed with full information on the stock and the growth function, but also play a major role in the cases in which only exact stock information is given. Pulse fishing is especially frequent in the case with noisy stock information. Finally, behaviour that is in line with a linear growth model is especially frequent in the ambiguous environment.

The paper is organised 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 Fig. 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.[2] Assume harvesting costs equal zero, normalise the price to one, and let the discount factor be denoted by $\rho = 1/(1+\delta)$, where $r > \delta$.[3] The optimal extraction

[1] Oscillations have been reported by a number of authors, e.g. Moxnes (1998a,b). Schnier and Anderson (2006) report pulse fishing behaviour that is very similar to our observations. Sterman (1994) reports a number of studies that find a misperception of non-linear growth. We discuss further findings in dynamic non-linear decision problems in Sect. 3.

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

[3] If the intrinsic growth-factor r is smaller than the interest rate δ, costless harvesting implies an immediate extinction of the stock. Thus $r > \delta$ is a necessary and sufficient condition for an interior solution of the sole-owner's maximisation problem. In the open-access fishery this condition is insufficient to prevent extinction

policy in the finite-horizon management problem can be determined as the solution to the following program.

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

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 $t \in \{1, 2, \ldots, T\}$. The optimal solution to this problem can be calculated by means of Bellman (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.

$$J_0(x) = \max_{y_T} \left\{ \rho^T y_T \right\} \tag{2}$$

The final extraction y_T that maximises the value function in Eq. 2 is equal to the maximal feasible y_T, which coincides with the stock remaining in period T, x_T. Hence, $J_0(x) = \rho^T x_{T,}$ 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 maximisation procedure, $J_1(x)$.

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

From the first order condition follows the optimality equation $F'(z_{T-1}) = r(1 - 2z_{T-1}/K) = \delta$. Solving this equation, we obtain the end-of-period stock size $z^*_{T-1} = (1 - \delta/r)K/2$, 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, $J_1(x) = \rho^{T-1}(x_{T-1} - z^*_{T-1}) + \rho^T(F(z^*_{T-1}) + z^*_{T-1})$. Proceeding by backward induction, the following general expression is determined,

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

where $z^* = (1 - \delta/r)K/2$ and $y_t = \max\{x_t - z^*; 0\}$. Hence, the first term of the maximisation procedure is $J_T(x) = x_0 - z^* + \sum F(z^*)\rho^{T-j} + \rho^T z^*$. The first extraction is determined by the initial stock size $x_0 = K$ and the optimal end stock size z^*, $y^*_0 = x_0 - z^* = K - z^*$.[4] 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 problem.[5] Hence, the extraction plan in the finite-horizon management problem coincides with the one in the infinite-horizon case

Footnote 3 continued
of the resource stock, as the equilibrium level of the resource stock is determined by the ratio of harvesting cost to the price and thus immediate extinction of the stock follows for any positive price at zero cost.

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

[5] See Clark (1990, Chap. 2) for a derivation of a solution to the infinite-horizon problem and a discussion.

(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

3.1 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; 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*.[6] People are insensitive to non-linearity and violate basic rules of probability. The robustness of the misperception of feedback and the poor performance … 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 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 environment. As pointed out in the

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

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.[7] 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 cannot 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.[8] 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.[9] Finally, there might be an emotional decision bias of subjects—particularly of pity—which might be associated with slaughtering of fish.[10] In order to guarantee salience of the incentive structure the experiment must be neutrally framed. In the experiment we ask subjects to maximise their savings, which correspond to the number of extracted units on each subject's account. The procedures are detailed in the following subsection.

3.2 Experimental Procedures

In the computerised experiment,[11] 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

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

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

[9] In the laboratory, the experimenter usually sets the exchange rate of actual to virtual money in expectation of the subjects' average performance. In a deterministic setting like ours, subjects may use the exchange rate to infer the bioeconomic optimum by guessing the average payoffs the experimenter expects. To avoid any such bias in behaviour, the incentive structure should ensure that laboratory market prices are unrelated to the bioeconomic optimum. In our setting, the market price of one extracted unit is set equal to one and subject payments are, in fact, determined relative to the bioeconomic optimum (i.e. relative to efficiency). Our subjects, however, cannot derive any information on the optimal behaviour from the payment scheme, because they have no clue how many units they must extract for a Euro, even though they know that extracting more increases their current payoff.

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

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

Table 1 Experimental treatments : number of independent observations

	Accurate stock size		Noisy signal[a] about stock size	
Growth function information	*GS*:	25	*G*:	35
No growth function information	*S*:	31	*No*:	30

[a] The noisy signal equals the true stock size multiplied by a random number from the interval [0.75, 1.25]

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.[12] Subjects were instructed accordingly.[13]

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.[14] Efficiency was hence a number between zero and one. The payoff a subject received at the end of the experiment was the product of efficiency and the premium to be paid in a treatment which was known to subjects.[15]

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 period 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 121 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 (12 subjects per treatment) and the other at the CentERlab at Tilburg University (13–23 subjects per treatment). Each subject participated in only one treatment.

12 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 zero 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."

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

14 The maximum is easily calculated by applying the results from Sect. 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.

15 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.

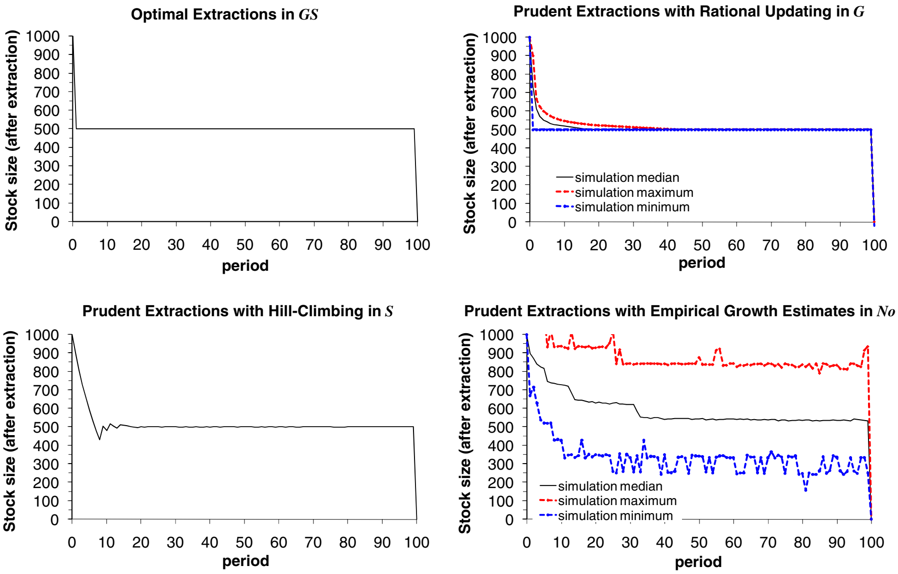

Fig. 2 Theoretical benchmarks

3.3 Theoretical Benchmarks

In Sect. 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 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.' The benchmark 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 Fig. 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 Fig. 2) and *S* (bottom left in Fig. 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 Fig. 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% efficiency. While average efficiencies of 99.7 and 85.4% are achieved in the treatments *G* and *No*, correspondingly. These average efficiencies have been computed on the basis of the theoretical benchmarks and the individual realisation of the stock signals.

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 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 Fig. 2 uses a simple hill-climbing search algorithm with an exponentially decreasing experimentation rate ε_t. In any period t, the decision maker extracts an amount that leaves $1 - \varepsilon_t$ of the last period'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 period τ, extraction is adjusted to restore the previous stock size, before continuing experimentation at an halved rate, i.e. at the rate $\varepsilon_\tau = \varepsilon_{\tau-1}/2$. 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.[16] Using the experimental parameters and an initial experimentation rate of .1, we can show that after only 9 of 100 periods, the process converges to stock levels that are within 10 points around the perfect information benchmark. By period 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 period, the extraction history and stock size signals can be combined with the growth function information to tighten the lower and upper boundaries on the initial stock size. As more and

16 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 Fig. 2 is derived for an initial stock size of 1000 and an initial experimentation rate of $\varepsilon_0 = .1$.

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 period 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 periods), 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 Fig. 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.[17] 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 periods. Even in the worst case observed in the simulation, no more than 40 periods 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%.[18]

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 Fig. 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 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 periods will not entail enough

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

[18] 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.

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

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 organised 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 the penultimate period and extinction of the resource with the last decision. We conclude the section by classifying observed individual behavioural pattern.

4.1 The First Extraction Decisions

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

The deviations from the optimal extraction increased from treatment *GS* through *No* (two-tailed Jonckheere test of ordered alternatives at $\alpha = .01$). Equation 5 represents a pooled dummy regression of the distance of stock after the first extraction from the optimum on *G* and *S*. The variables D*G* and D*S* denote dummy variables that take a value of one if a subject receives growth information and accurate stock information, respectively, and zero otherwise. 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.

$$
\begin{array}{lllll}
|Stock_1 - 500| = & 438^{**} & -223\mathrm{D}G^{**} & -57\mathrm{D}S^{*} & \hat{R}^2 = .35 \\
 & 23.96 & 27.64 & 27.71 & \text{[std. error]} \\
 & 18.30 & -8.08 & -2.07 & [t\text{-ratio}]
\end{array}
$$

** significant at 1%, * significant at 5%, one-tailed (5)

Table 3 Efficiency

Treatment	#	Benchmark prediction	Observed			
			Maximum	Average	Minimum	SD
GS	25	1.000	0.997	0.874	0.613	0.108
G	35	0.997	0.946	0.727	0.091	0.254
S	31	0.986	0.922	0.660	0.057	0.252
No	30	0.854	0.853	0.398	0.018	0.287

4.2 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.[19] The maximum of the observed efficiency levels is close to the efficiency level proposed by the theoretical benchmarks in every treatment. However, the deviation of the average observed efficiency from the benchmark differs substantially across treatments, which is due to the differences in the variance across treatments.

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).[20] In treatment *G* (and only in treatment *G*), three subjects extinguish the resource within the first 19 extractions (see Table 8 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 Eq. 6 indicates that both treatment variables had a significant effect on efficiency.[21] The knowledge of the growth function implied an increase of average efficiency by 27.6%, the accurate stock size information by 21%.

$$\begin{array}{lllll} \text{Efficiency} = & .427^{**} & +.276\text{D}G^{**} & +.210\text{D}S^{**} & \hat{R}^2 = .31 \\ & .038 & .044 & .044 & [\text{std. error}] \\ & 11.17 & 6.26 & 4.66 & [t\text{-ratio}] \end{array}$$

$$^{**}\ \text{significant at 1\%, one-tailed} \tag{6}$$

4.3 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.[22] While the overall average distance of

[19] Table 8 in the appendix records individual efficiency levels.

[20] If we consider only the data from the second half of the experiment the two-tailed Mann–Whitney test rejects the null hypothesis of same efficiency for all pair-wise comparisons at the 5% level of significance. The theoretical solutions suggested that efficiency should be indistinguishable between treatments in later rounds. However, this suggestion is not supported by the data, Since the efficiency levels in *G* are significantly greater than in *S*.

[21] Conducting the regression with two additional dummies, one to determine the cross-effect of growth and stock information and one to determine the subject pool effect, we find no further significance. Obviously, the two types of information have no synergies and the results are robust across subject pools.

[22] An anonymous referee suggested that the gradual decline of the stock sizes towards the end of the experiment could be due to a relaxed 'risk aversion,' as intuition may suggest that there is less to lose in the case of

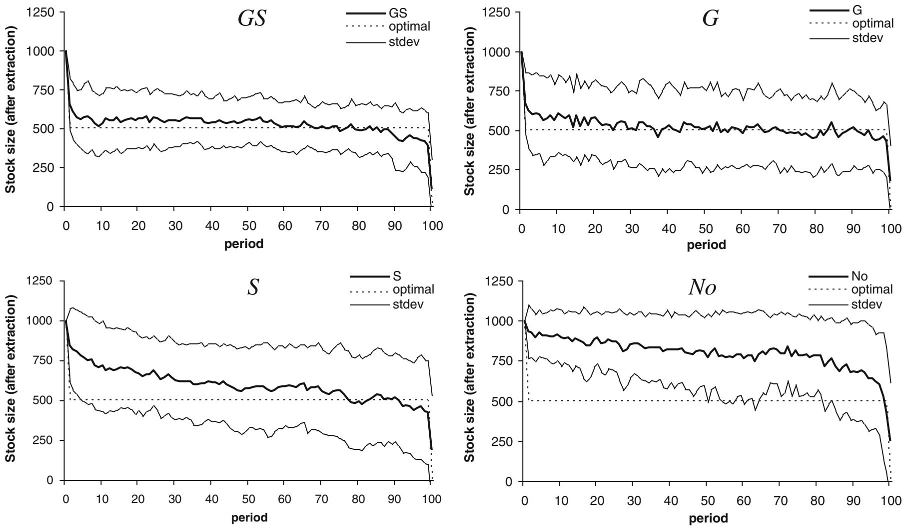

Fig. 3 Evolution of average end-of-period stock (after extraction) compared to the optimal path. *Note*: *horizontal line* indicates efficient extraction path, *oscillating lines* indicate average stock size after extraction and standard deviation bands

observed end-of-period stock from the optimum is rather small, especially towards the end of the experiment, the spread of the outcomes is large enough to substantially impact on the overall efficiency. Yet we find that efficiency increases across periods. To show this increase we compare the observed extraction with the most rapid approach to the optimal path. Payoff maximisation 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 233 units the stock would have to be rebuilt by zero extraction.

To check whether the stock converges to the optimal level over time, i.e. over the sequence of extraction decisions, we run fixed effect regressions to estimate the parameter for time. The regression results recorded in Table 4, particularly the negative coefficients of the time variable, confirm that efficiency increases in the course of the experiment. The size of 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. Since the subjects in the treatments *S* and *No* started further away from optimum than those in the treatments *GS* and *G*, they faced a greater potential for convergence over time. Not counting for the last extraction, the average distance of the observed extraction and optimal extraction did not decrease below 177 and 241 units in any period of the treatments *S* and *No*, respectively. In contrast, in the treatments *GS* and *G*, the average deviation from the optimum never exceeded 175 and 218 units, respectively.

Footnote 22 continued
an irreversible mistake. If so, such 'risk aversion' may have contributed to a certain elevation of the stock prior to the last few rounds. Indeed, unless the resource is wiped out, due to our design, the same mistake usually leads to the same loss whether made at the beginning or at the end of the experiment. But in the treatments where the growth–function was unknown, subjects obviously were not given this information, so they had to find it out by trial and error.

Table 4 Distance of observed and the optimum extraction across time

		Coefficient	Std.error	*t*-ratio
GS	Constant	148**	2.998	49.30
	Period	−0.313**	0.052	−6.08
G	Constant	186**	3.730	49.92
	Period	−0.502**	0.064	−7.83
S	Constant	266**	3.543	75.29
	Period	−0.915**	0.061	−15.03
No	Constant	413**	3.971	104.01
	Period	−1.343**	0.068	−19.70

** Significant at 1%, two-tailed

Table 5 Over- and under-harvesting

Treatment	# Over-harvester	# Under-harvester
GS	10 (40.0%)	15 (60.0%)
G	18 (51.4%)	17 (48.6%)
S	11 (35.5%)	20 (64.5%)
*No***	3 (10.0%)	27 (90.0%)

** Difference is significant at 1%, two-tailed

Greater efficiency also corresponds to smaller stock sizes across treatments. On average, our experimental data suggest a rather equal spread of over- and under-harvesting in all but the *No* treatment. As can be seen in Table 5, using the binomial test on the distribution of over- and under-harvesters, we only observe a significant bias in the *No* treatment, in which 90% of the subjects under-harvest.[23] Over-harvesting was most heavy in treatment *G*, where three subjects extinguished the resource within 19 decisions. In all of these three observations, the last extraction before extinction did not exceed the signalled stock, but in fact it did exceed the actual stock size.

The propensity to harvest more in *GS* and *G* than in *S* and *No* suggests that subjects are more confident with their extraction decisions when they receive exact information on the stock dynamics. In the *No* treatment, subjects have little information and seem to be reluctant to exploit the resource too much. The level of extraction they choose is perhaps related to the fact that initial condition in the experiment is close to the carrying capacity.[24] Brekke and Moxnes (2003) show that subjects that start off above the target capacity tend to have higher levels of stock throughout. Even though we do not vary the initial condition, our finding may be related to that result.

[23] We classified a subject as over-harvester, if the subject's end-of-period stock was below the optimum in more periods than above. Since we observed no tied cases, all the other subjects were classified as under-harvesters.

[24] At the carrying capacity the non-linear growth curve implies that growth first increases with extraction before it decreases. If the initial conditions are close to the maximum sustainable yield, however, growth decreases with extraction. In the former case, constant harvesting can lead to a stable equilibrium while in the latter case it may accelerate depletion.

Table 6 Distribution of final extinction and non-extinction

Treatment	Zero-stock left	Extinction attempted	Minimal stock left	Non-extinction
GS	15	–	1	9 (36%)
G	14	4	–	17 (49%)
S	14	–	4	13 (37%)
No	17	5	–	8 (27%)
Total	60	9	5	47 (39%)

4.4 The Final Extraction Decision & Extinction of the Resource

With the final extraction, subjects were expected to extinguish the resource. However, only about one half of them did so as shown in Table 6. Sixty subjects (50%) ended the experiment with a zero stock, 9 subjects in treatments *G* and *No* did not extinguish the resource, but extracted all units signalled to them in the last decision. Apparently they had forgotten that the signal was most likely incorrect.[25] There were also 12 subjects 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 *No* 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.

4.5 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.[26] 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).[27] 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 Fig. 4, we have plotted the individual stock development of four subjects, who exhibit behaviour that can be identified as "typical" for the control theory hypothesis. The number of observations that follow similar patterns is stated in Table 7. For instance, 15 individual charts or 60% of

[25] In total, 23 subjects in treatments *G* and *No* extracted the signalled stock size in the last decision. In the other 14 cases the signalled stock size exceeded the actual stock size.

[26] 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.

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

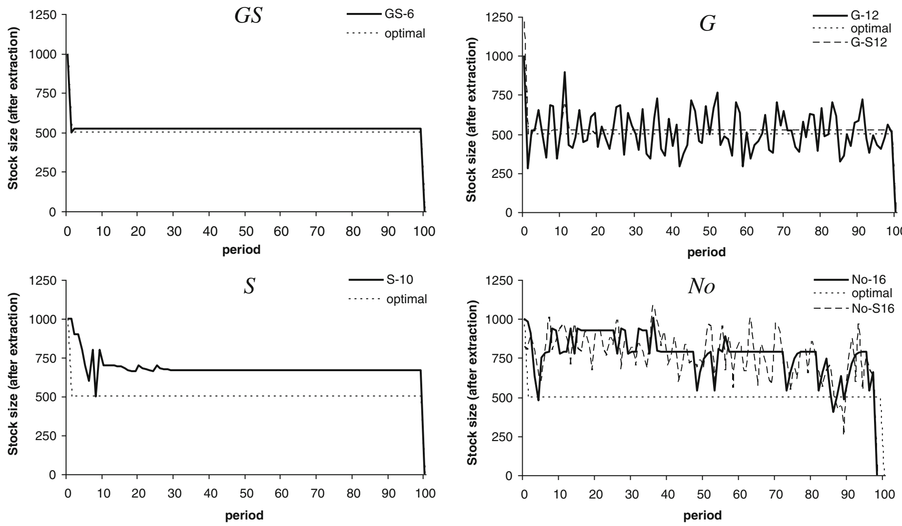

Fig. 4 Behavioural pattern—control theory

Table 7 Individual pattern: summary

Treatment	#	Control theory	Pulse fishing	Linear world	Unexplained
GS	25	15 (60%)	9 (36%)	–	1 (4%)
G	35	12 (34%)	19 + 3 (63%)	–	1 (3%)
S	31	12 (39%)	10 (32%)	5 (16%)	4 (13%)
No	30	2 (7%)	13 (43%)	15 (50%)	–
Total	121	41 (34%)	54 (45%)	20 (17%)	6 (5%)

the observations in *GS* display straight lines similar to that presented in the top left panel of Fig. 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 that is indicated with a dotted line, these plots exhibit two further lines: The unbroken line represents the movement of end-of-period stock (after extraction) and the dashed line represents the noisy end-of-period 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 dashed line is straight indicating that the extractions were meant to maintain a constant stock signal. In contrast to this, the straight segments in treatment *No* exhibit a constant end-of-period stock, which hints at a policy of constant extraction.

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 Fig. 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 rationalised if

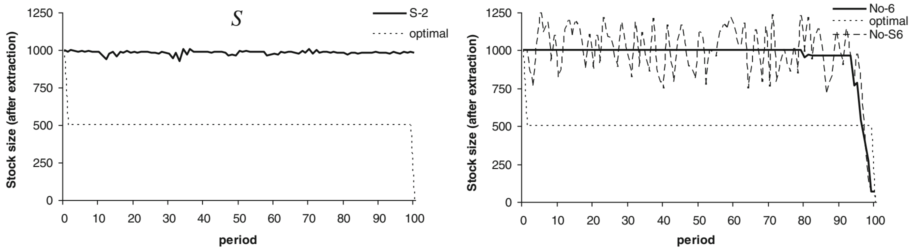

Fig. 5 Behavioural pattern—linear world

not in the light of Brehmer (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.

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 Sect. 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 Fig. 6). Such fluctuations in stock can be due to "pulse fishing," a behaviour that makes sense in some fishery environments (Schnier and Anderson 2006),[28] but not in our experiment. Our subjects were informed that they are facing a deterministic system. In such settings, the optimal harvesting behaviour is non-pulsing. 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 observe comparable patterns, claim that subjects' learning in complex environments is poor. This argument could in fact explain the persistency of these oscillations in the data.

The behavioural patterns that classify our subjects almost perfectly are summarised in Table 7. About 34% of subjects tried to achieve control over the complex system by holding either stock or extraction levels constant. Another 45% of subjects managed their stocks by pulse fishing and 17% misperceived the non-linearity of the environment and extracted almost nothing. None of these behavioural patterns could be used to classify the remaining 5%.

Testing for differences in distributions with pairwise chi-squared tests, we find that the distribution of patterns significantly differs between *No* and any other treatment, as well as between *G* and any other treatment.[29] The distribution of patterns across GS and G, however, are statistically indistinguishable.

5 Summary and Conclusions

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) that is in line with the infinite-horizon solution in all but the last extraction period. Hence, the

[28] Pulse fishing refers to a behaviour that alternates between harvesting and not-harvesting from a stock (Schnier and Anderson 2006). Note that seeing the periods with zero harvesting in our graphs is not possible, because our graphs display the end-of-period stock of each period and not the harvest.

[29] The difference between distributions in *GS* and *G* are significant at the 5% level, two-tailed, all other effects are significant at the 1% level, two-tailed.

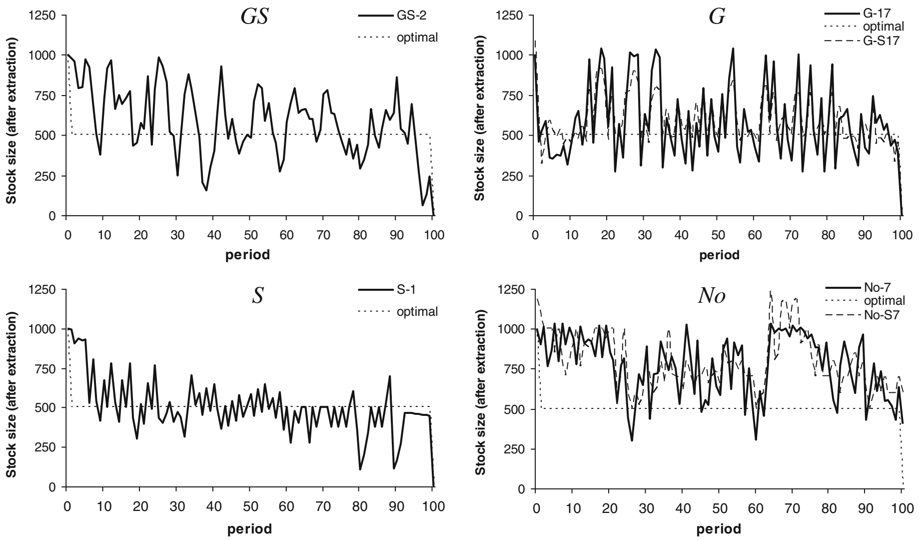

Fig. 6 Behavioural pattern—pulse fishing

theoretical benchmark basically does not differ from that used in other fishery experiments, where "infinite" horizon tasks are simulated by either providing a residual resource value in the final period or unexpectedly cutting sessions short that were ostensibly longer. For laboratory studies, our finite-horizon design seems more compelling and transparent than the residual payment and indefinite length settings.

In line with the experimental literature on resource extraction (Mason and Phillips 1997; Moxnes 1998a,b, 2000, 2004; Brekke and Moxnes 2003; Schnier and Anderson 2006), extinction of the resource before the end of the experiment was generally not a problem in our study. We only observed a few early terminations in the treatment, in which subjects received information on the growth function, but only a noisy signal of the stock.

In the treatments with some information, we find over-harvesting and under-harvesting in almost equal frequencies. We find a significant under-harvesting problem only when subjects neither have exact information on the growth function nor on the stock size. This result diverges from most of the findings in the literature, where over-harvesting is generally more prevalent than under-harvesting. Our result for the no-information treatment may be due to the specific experimental characteristics: in our study, subjects had a direct control over the fishery resource and the initial stock levels were set to the carrying capacity. As Moxnes (2004) shows, direct control may have an effect on behaviour. It, however, seems difficult to compare those findings with our results, due to the numerous other design differences. Whether the initial stock size in our setting has an effect on harvesting behaviour also remains an open issue for future research. There is evidence suggesting that subjects facing a stock that is well above the target will tend to maintain a higher stock level throughout (Brekke and Moxnes 2003). Obviously, this result is in line with our finding. It, however, cannot explain the strong treatment differences that we observe.

Our study also adds to the evidence on the behavioural relevance of pulse fishing, which was experimentally first reported by Schnier and Anderson (2006). It seems that using a total moratorium to help resource recovery is not only a path often chosen by subjects in patchy environments (e.g. Schnier and Anderson 2006), but also in a single resource pool. In fact, the instrument is also often used by political administrations that seem to prefer a total ban on harvesting for a limited time over other recovery methods. Perhaps it is easier to control a moratorium both on an individual self-control level and in a public law-enforcement setting.

The efficiency of extraction decisions in our experiment was an estimated 21% higher if the stock signal was accurate and 27.6% higher if subjects had knowledge on the growth function. Both results seem to suggest that research can significantly help to increase extraction decisions. This corresponds well to the findings by Brekke and Moxnes (2003), who report that decision support systems (i.e. "better" information) can enhance outcomes. Our study adds to these finding, because we can identify different behavioural patterns that are more or less likely to evolve, depending on the informational setting. In cases in which there is ample information on the growth function and the stock size, we can expect to find many subjects using control heuristics to either keep the stock size or the harvest constant. Whenever stock information is noisy, pulse fishing behaviour is especially frequent. If additionally the growth function is not known, we can expect a substantial number of decision makers, whose behaviour indicates a linear misperception of the growth dynamics. Obviously, identifying the behavioural patterns that are most prevalent in a certain resource extraction environment may help design institutions that are especially effective.

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. Yet, all these complications can conveniently be accommodated within the presented framework. The standard logistic growth model which we considered in the experiment seems to be ideally behaved to provide the experimenter with a rich research environment.

Appendix

See Table 8.

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 maximise 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.

Table 8 Individual efficiency

#	*No*	*S*	*G*	*GS*
1	0.018	0.057	0.091[a]	0.613
2	0.036	0.066	0.095[a]	0.618
3	0.043	0.168	0.151[a]	0.706
4	0.056	0.192	0.396	0.741
5	0.059	0.358	0.437	0.803
6	0.084	0.484	0.488	0.810
7	0.087	0.552	0.504	0.822
8	0.157	0.561	0.507	0.851
9	0.161	0.606	0.521	0.877
10	0.178	0.643	0.580	0.879
11	0.192	0.648	0.649	0.882
12	0.218	0.657	0.656	0.891
13	0.221	0.658	0.782	0.903
14	0.267	0.666	0.803	0.904
15	0.276	0.697	0.814	0.917
16	0.388	0.740	0.836	0.920
17	0.487	0.757	0.853	0.939
18	0.557	0.771	0.860	0.942
19	0.611	0.773	0.866	0.950
20	0.613	0.791	0.866	0.964
21	0.623	0.824	0.868	0.970
22	0.639	0.836	0.874	0.974
23	0.646	0.843	0.890	0.980
24	0.648	0.851	0.896	0.993
25	0.686	0.861	0.902	0.997
26	0.695	0.866	0.904	
27	0.755	0.897	0.907	
28	0.836	0.901	0.916	
29	0.846	0.906	0.927	
30	0.853	0.914	0.927	
31		0.922	0.929	
32			0.936	
33			0.936	
34			0.937	
35			0.946	

Note: Subjects' results are arranged according to their performance
[a] 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

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% and 125%. 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

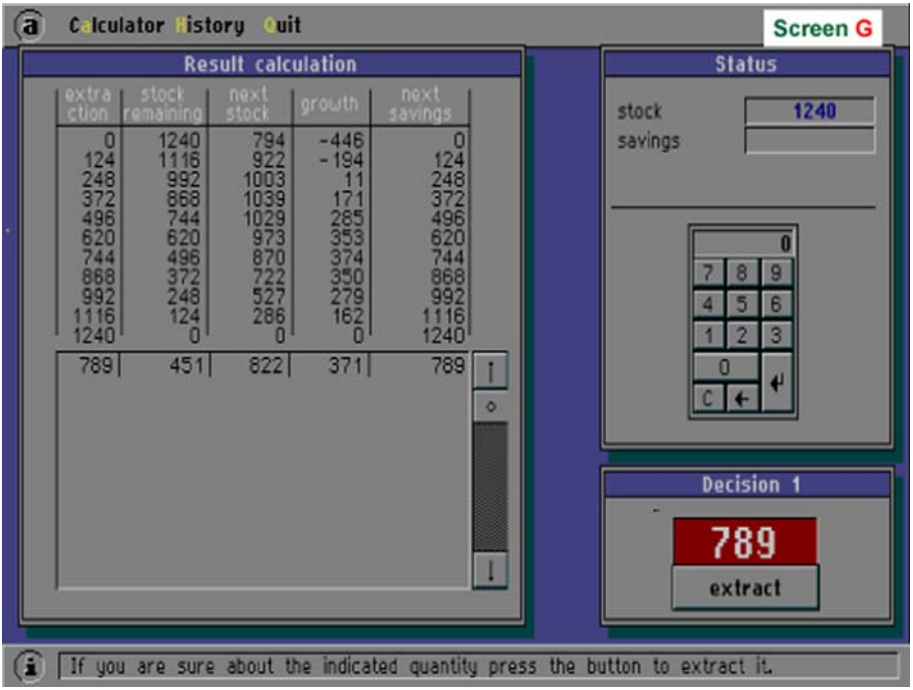

References

Abbink K, Sadrieh A (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? J Econ Behav Organ (in press)

Bellman R (1957) Dynamic programming. Princeton University Press, Princeton

Brehmer B (1980) In one word: not from experience. Acta Psychol 45:223–241

Brehmer B (1992) Dynamic decision making: human control of complex systems. Acta Psychol 81:211–241

Brekke KA, Moxnes E (2003) Do numerical simulations and optimization results improve management? Experimental evidence. J Econ Behav Organ 50:117–131

Clark CW (1990) Mathematical bioeconomics, 2 edn. John Wiley, New York

Diehl E, Sterman JD (1995) Effects of feedback complexity on dynamic decision making. Organ Behav Human Decis Process 62(2):198–215

Edwards W (1962) Dynamic decision theory and probabilistic information processing. Hum Factors 4:59–73

Gordon SH (1954) The economic theory of a common property resource: the fishery. J Polit Econ 62:124–142

Mason CF, Phillips OR (1997) Mitigating the tragedy of the commons through cooperation: an experimental evaluation. J Environ Econ Manage 32:148–172

Moxnes E (1998a) Not only the tragedy of the commons: misperceptions of bioeconomics. Manage Sci 44:1234–1248

Moxnes E (1998b) Overexploitation of renewable resources: the role of misperceptions. J Econ Behav Organ 37:107–127

Moxnes E (2000) Not only the tragedy of the commons: misperceptions of feedback and policies for sustainable development. Syst Dyn Rev 16(4):325–348

Moxnes E (2003) Uncertain measurements of renewable resources: approximations, harvest policies, and value of accuracy. J Environ Econ Manage 45(1):85–108

Moxnes E (2004) Misperceptions of basic dynamics, the case of renewable resource management. Syst Dyn Rev 20(2):136–162

Paich M, Sterman JD (1993) Boom, bust, and failure to learn in experimental markets. Manage Sci 39(12):1439–1458

Rapoport A (1975) Research paradigms for the study of dynamic decision behavior. In: Wendt D, Vlek C (eds) Utility, probability and human decision making. Reidel, Dordrecht pp 349–369

Rouwette EAJA, Größler A, Vennix JAM (2004) Exploring influencing factors on rationality: a literature review of dynamic decision-making studies in system dynamics. Syst Res Behav Sci 21(4):351–370

Schnier KE, Anderson CM (2006) Decision making in patchy resource environments: spatial misperception of bioeconomic models. J Econ Behav Organ 61(2):234–254

Sterman JD (1989a) Deterministic Chaos in an experimental economic system. J Econ Behav Organ 12:1–28

Sterman JD (1989b) Misperception of feedback in dynamic decision making. Organ Behav Human Decis Process 43(3):301–335

Sterman JD (1989c) Modelling managerial behavior: misperceptions of feedback in a dynamic decision making experiment. Manage Sci 35(3):321–339

Sterman JD (1994) Learning in and about complex systems. Syst Dyn Rev 10(2–3):291–330