Journal of Bioeconomics (2006) 8:147–165
DOI 10.1007/s10818-006-9001-2

Information Structure and the Tragedy of the Commons in Resource Extraction

RABAH AMIR
CORE, University of Louvain, 34 Voie du Roman Pays, 1348 Louvain-la-Neuve, Belgium (amir@core.ucl.ac.be)

NIELS NANNERUP
Department of Economics, University of Southern Denmark, 5230 Odense M, Denmark (nna@sam.sdu.dk)

Synopsis: This paper considers the well-known Levhari-Mirman discrete-time model of resource extraction, and investigates the effects of the information structure of the dynamic game – open-loop, Markovian or history-dependent – on the equilibrium consumption path and the overall utility of the agents. Due to the special structure of the model, the open-loop regime yields a Pareto-optimal outcome. The Markovian regime leads to the most pronounced version of the tragedy of the commons. History-dependent behavior yields an outcome set that is intermediate between the other two cases, and that may include the Pareto-optimal outcome in some cases. The level of efficiency of equilibrium behavior is thus U-shaped as a function of the level of information the agents' extraction strategies are based on. The analysis suggests that in environments characterized by a dynamic (and no market) externality, forcing agents to commit to open-loop behavior would constitute welfare-improving regulation.

Key words: dynamic resource games, open and closed loop strategies, trigger strategies, Pareto optimality, regulation

JEL classification: Q20, C73, H41, Q22

1. Introduction

The tragedy of the commons is one of the most readily accepted conclusions in economic analysis. In the absence of clear-cut property rights assignment or in the presence of public goods, the emergence of the tragedy of the commons is almost always a foregone conclusion. In other words, a first-best outcome is typically not expected to prevail in such environments in the presence of multiple agents whenever their behavior is assumed noncooperative. Among the many diverse economic settings characterized by this outcome, common-property resource extraction is one of the most natural and most intensively investigated problem.

Of the several different competing models of noncooperative resource extraction[1], the work of Levhari & Mirman (1980) remains the leading discrete-time model, and one of the most influential overall. Their simple model postulates two agents as joint owners of a renewable resource, with each of them maximizing the infinite-horizon discounted sum of utilities depending on own consumption only. An important feature of their model is that it reflects no market externality. They refer to the only interdependence in their model, the intertemporal common-property feature, as a dynamic externality. Adopting the specific framework of log utility and isoelastic growth function for the resource, they provided a simple analysis of the Markovian (also known as feedback or closed-loop no-memory[2]) equilibrium of the dynamic game based on a closed-form solution with linear consumption strategies. Their results confirm that, relative to the first-best or cooperative solution, the Markovian equilibrium leads to overconsumption (at every stock level[3]) and to a lower overall utility level for each agent. Their model has recently been generalized by Fisher & Mirman (1996) to cover the more general case where there are two biologically interacting species of fish, and further extended by Datta & Mirman (1999).

In a novel attempt to model history-dependent behavior in such an environment, Cave (1987) subsequently considered the symmetric version of this model and analyzed 'cooperative' equilibria secured by trigger strategies. In his setting, the two agents agree on extraction paths that mutually improve on the Markovian outcome, with reversion to the latter for the indefinite future constituting the punishment mode in case a deviation from the 'cooperative path' is detected.[4] At any point in the course of the game, the players recall all the previous history of play: states as well as actions. Some recall of history is obviously necessary for the players to be able to monitor compliance with the cooperative extraction path. The equilibria thus derived have the desirable property of subgame-perfection, due to the equilibrium nature of the punishment phase. Ingeniously exploiting the rich structure of this framework, Cave (1987) provides a complete characterization of the associated large equilibrium set. In particular, whether or not this set includes a Pareto-optimal extraction path depends on a derived simple condition on the parameters on the problem.

The present paper is an attempt to understand the effects of the information structure of the dynamic game of Levhari & Mirman (1980) on the characteristics of the resulting equilibria. The main underlying question is whether a monotonic relationship exists between the level of information available to the agents and the efficiency properties of the resulting equilibria. We consider three different information structures, listed in order of increasing information for the players as open-loop, closed-loop no-memory (or Markovian), and history-dependent strategies. For the latter two cases, we rely completely on the results of Levhari-Mirman and Cave, respectively, as described above and in detail in the body of the paper. Our first task then is to investigate the structure of equilibrium under open-loop behavior by the agents: Consumption at any period depends only on the initial stock level and on the date. The agents are thus committed at the very beginning of the game to a fully specified course of play that cannot be altered along the way, as no dynamic information becomes available in the course of the game. The game may thus be viewed as a one-shot game with infinite sequences (of consumption levels) as strategies.

Our first main result is that the open-loop equilibrium coincides with the symmetric Pareto-optimal solution. While at first very surprising, we argue that this result is not so unnatural when one considers that the only externality present in the game is the dynamic externality. In other words, this result would survive an extension to more general functional forms for the utility and growth functions, but would not extend to a more general setting that includes a market externality.[5]

This finding, together with the results of the two previous studies, leads to the following assessment: Equilibrium efficiency does not depend in a monotonic way on the level of information available to the agents. Rather, efficiency as a function of information is U-shaped, with a maximum at the lowest information level (i.e. with open-loop strategies) and possibly at the maximal information level (i.e. with full recall of history), with a local maximum being guaranteed there, and a minimum at the intermediate level of information (i.e. with Markovian strategies). An intuitive account of this interesting conclusion is provided below, together with a detailed description of the general properties of each of the information structures at hand.

Besides shedding some light on the relationship between information structure and equilibrium efficiency in the specific dynamic game at hand, our results offer one policy implication worth exploring: If regulatory action could induce the agents to behave according to open-loop strategies in environments characterized by the sole presence of the dynamic externality, the resulting equilibrium would be a first-best outcome. A possible way of accomplishing this is to force the agents to submit a vector of specified consumption levels, over some fixed horizon, to which they would remain committed.[6] If the players are then forced to stick to their announcements, the outcome would be Pareto-optimal over the given horizon. An interesting property of such a scheme is that it is likely to require much less monitoring to ensure compliance than current regulation via quota assignment. The reason for this important feature follows from an interesting property of open-loop equilibria: If all other players are using open-loop strategies, a given player cannot (unilaterally) improve on his payoff by using more complex strategies.[7] In other words, a unilateral violation of this scheme would not be worthwhile to the perpetrating agent.

By contrast, the standard quota regulation is subject to unilateral violations since, given the quotas assigned to the other players, the best response of a player is often to consume more than his own quota. Indeed, in practice, it has been observed in various settings that this widespread regulation scheme has enjoyed rather limited success. For instance, the EU has on several occasions experienced violations by member states of the quota levels agreed upon within the Community. Another well-known example is OPEC, an international oil cartel that has frequently encountered major difficulties in ensuring compliance of the member states with their assigned extraction quotas, the explanation being again the unilateral incentive for over-extraction, given the partners' compliance with their quotas.[8]

The rest of the paper is organized as follows. Section 2 presents the basic model. In Sections 3–6, the solutions under the three different information structures and the Pareto-optimal case are discussed. Section 7 discuss the main comparative results and their derivation. Section 8 discuss policy implications of the results. Section 9 explores the relationship between the sustainability of the natural resource and the information structure. Section 10 provides a short conclusion.

2. The model

Two identical agents share the rights to exploit a renewable resource over an indefinite future. The resource stock develops over time according to a biological growth rule given (upon a normalization of units) by

$$x_{t+1} = x_t^{\alpha}, \quad 0 < \alpha \leq 1, \tag{1}$$

and this rule is common knowledge to both agents. Following Levhari & Mirman (1980) and Cave (1987), the utility of agent i is assumed to be:

$$u_i(c_i^t) = \log c_i^t, \;\; i = 1, 2 \tag{2}$$

where c_i^t is his consumption at time t.[9] Let $0 < \delta < 1$ be the common discount factor of the agents. Agent's objective is to maximize the sum of the discounted utility of consumption

$$Max \sum_{t=0}^{\infty} \delta^t \log c_t^i \tag{3}$$

subject to

$$x_{t+1} = \left(x_t - c_t^1 - c_t^2\right)^{\alpha} \;\; t = 0, 1, \ldots \tag{4}$$

Upon specification of the strategies of the two players, the formulation of an infinite horizon dynamic game between the agents will be complete. Over the next four sections of the paper we compare the equilibria that arise under different information structures or, equivalently, under different sets of strategies used by the agents, in terms of their efficiency properties and their consequences on the resource stock.

This model is specific on two important counts. First, the utility and growth functions have special (though commonly used) functional forms. Second, the agents face no market externality here, the only externality being what Levhari & Mirman referred to as a dynamic (common-property) externality.

In what follows, we can consider each of the four solutions separately. For each case, we provide a definition and summary of the salient features of the strategies at hand, and then derive the associated equilibrium. Further details may be found in Basar & Olsder (1999), Fudenberg & Tirole (1986) or Amir (2003).

3. The open-loop equilibrium

3.1. *Definition and general facts*

This subsection provides a definition of open-loop strategies in deterministic Markovian dynamic games[10] and a summary of their general properties and limitations.

In the present setting, an open loop (henceforth OL) strategy σ for an agent is an infinite sequence $\sigma(x_0) = (c_0, c_1, \dots) \in R^{\infty}$ specifying the resource consumption level at every time period t over an infinite horizon as a function of the initial stock x_0 and calendar time t only. Open-loop behavior thus rests on the premise that the players simultaneously commit at the beginning of the game to a completely specified list of actions, without any possibility of revision during the entire game. The players receive no new information, not even about the value of the current state. Hence, no contingency planning of any sort is possible.

Several important observations concerning open-loop equilibria should be noted. Interestingly, while some of these properties may appear complex, they actually all follow directly from the well-known properties of Markovian dynamic programming. To begin with, in deterministic Markov dynamic optimization, there always exists an optimal open-loop strategy (under minor regularity conditions). In other words, in one-person problems, restricting oneself to open-loop policies results in no loss of value compared to using more sophisticated behavior such as Markov or history-dependent policies. This fact is certainly intuitive, as is its failure in the presence of chance moves or stochastic transitions.[11]

The game-theoretic partial analog[12] of the above fact is perhaps less intuitive, though no less important. In deterministic dynamic games, an open-loop equilibrium remains an equilibrium when the strategy spaces are expanded to include Markovian or history-dependent strategies. The reason is that if the rival is using open-loop strategies, a player cannot achieve a higher payoff by using more sophisticated strategies than open-loop. This follows directly by invoking the above fact for the player's best-response problem which, given the open-loop strategies of the rivals, is a deterministic (Markovian) dynamic optimization problem.

Open-loop equilibria are generally not subgame-perfect. This is often viewed as a major drawback of this type of equilibrium. On the other hand, open-loop equilibria are typically much simpler to analyze than Markovian equilibria. This relative simplicity is at the heart of the widespread use of open-loop strategies in the early stages of the adoption of the tools of strategic dynamics in economics.[13] In recent years, economists have generally agreed that the commitment to a completely specified course of action over the indefinite future is not a realistic behavioral postulate in most cases of interest in economic dynamics. Furthermore, subgame-perfection is broadly viewed as a desirable property of dynamic equilibria. Consequently, focus has markedly shifted to Markovian behavior.

3.2. The open loop solution

In order to characterize the OL equilibrium, we begin with an analysis of the best response problem of agent 1 (say). Fix an open loop strategy $\tau^2 = \left(c_0^2, c_1^2, c_2^2 \dots\right) \in R^\infty$ for agent 2 that is feasible, i.e., that is such that $0 \le c_0^2 \le x_0$ and $x_{t+1} = \left(x_t - c_t^2\right)^\alpha \ge 0$, for $t = 0, 1 \dots$ Given the OL strategy τ^2, agent 1 clearly faces a dynamic optimization problem, and his best-response may be solved for via the discrete-time Maximum Principle as follows. With $\{\lambda_t\}$ being a vector of costate variables, the Hamiltonian is (the superscript is dropped for simplicity):

$$H\left(x_t, c_t^1, \lambda_{t+1}\right) = \delta^t \log c_t^1 + \lambda_{t+1}\left(x_t - c_t^1 - c_t^2\right)^\alpha. \tag{5}$$

The first-order condition for the maximization of the Hamiltonian is

$$\frac{\delta^t}{c_t^1} = \alpha\lambda_{t+1}\left(x_t - c_t^1 - c_t^2\right)^{\alpha-1}. \tag{6}$$

The associated second-order conditions hold, so the first-order conditions are sufficient. The costate equation is $\lambda_t = \partial H / \partial c_t^1$, or

$$\lambda_t = \alpha\lambda_{t+1}\left(x_t - c_t^1 - c_t^2\right)^{\alpha-1}. \tag{7}$$

Since the right-hand sides of (6) and (7) are the same, we have

$$\frac{\delta^t}{c_t^1} = \lambda_t. \tag{8}$$

Rewriting (8) with a time shift as $\lambda_{t+1} = \frac{\delta^{t+1}}{c_{t+1}^1}$ and substituting this for λ_{t+1} in (6),

$$\frac{1}{c_t^1} = \frac{\delta\alpha}{c_{t+1}^1}\left(x_t - c_t^1 - c_t^2\right)^{\alpha-1}. \tag{9}$$

An analogous equation can be derived for agent 2 given a feasible consumption path by agent 1. Then invoking symmetry, we can postulate that $c_t^1 = c_t^2$ for all t, at an open-loop equilibrium. Furthermore, we can postulate – and later confirm – that the (symmetric) solution of (9) has a closed-loop or feedback representation g that is linear in x. Then pluging $g(x) = \lambda(x)$ in (9) and dropping time subscripts leads to

$$\frac{1}{\lambda x} = \frac{\delta\alpha\,[(1-2\lambda)\,x]^{\alpha-1}}{\lambda\,[(1-2\lambda)\,x]^{\alpha}}. \tag{10}$$

After simplification, we have $\lambda = (1-\delta\alpha)/2$ or

$$g_{ol}(x) = \left(\frac{1-\delta\alpha}{2}\right)x. \tag{11}$$

Viewed as $g(x_t) = \left(\frac{1-\delta\alpha}{2}\right)x_t,\ t = 0, 1, \ldots,$ this is the closed-loop (or feedback) representation of the symmetric open-loop equilibrium consumption strategies.[14]

A simple way to calculate the corresponding equilibrium value function $V_{ol}(x)$ for each agent (i.e the equilibrium total discounted utility given an initial state x) is to substitute (11) in the functional equation of dynamic programming (see Corollary 4 in Section 7 for a justification), yielding

$$V_{ol}(x) = \log\left(\frac{1-\delta\alpha}{2}x\right) + \delta V_{ol}(x)\,[(\delta\alpha x)^{\alpha}], \tag{12}$$

and then postulating the 'guess' $\mathrm{V_{ol}}(x) = A\log x + B$ into (12), we get

$$A\log x + B = \log\left(\frac{1-\delta\alpha}{2}x\right) + \delta\left\{A\log[(\delta\alpha x)^{\alpha}] + B\right\}. \tag{13}$$

Identifying terms leads to

$$V_{ol}(x) = \frac{\log x}{1-\delta\alpha} + \frac{\log(1-\delta\alpha) + \delta\alpha\log\frac{\delta\alpha}{2}}{1-\delta} \tag{14}$$

where x is to be thought of as the initial state here. Under the open-loop equilibrium, the resource stock evolves according to the following difference equation

$$x_{t+1} = \left[x_t - 2\left(\frac{1-\delta\alpha}{2}x_t\right)\right]^{\alpha} = (\delta\alpha x_t)^{\alpha}. \tag{15}$$

At a steady-state equilibrium of the ressource stock, $\overline{x}$, we have $\overline{x} = (\delta\alpha\overline{x})^{\alpha}$ or

$$\overline{x}_{ol} = (\delta\alpha)^{\frac{\alpha}{1-\alpha}}. \tag{16}$$

4. The closed-loop (no memory) equilibrium

4.1. Definition and general facts

A feedback or closed-loop memoryless strategy for an agent is a function γ from the set of all possible stock levels to the set of possible consumption levels.[15] Since with such strategies, players are allowed to condition their extraction only on the value of the current stock, they necessarily consume the same amount of resource every time the same stock is observed, regardless of calendar time.

It is well-known that a feedback or Markov-stationary equilibrium of a Markov-stationary infinite-horizon dynamic game remains an equilibrium when history-dependent strategies are allowed. This also follows from the fact that with the rival playing a Markov-stationary strategy, a player's best-response problem is a Markov-stationary dynamic program, which is known to admit a Markov-stationary optimal policy. Note in this case that this argument is equally valid in the presence of chance moves (i.e. stochastic transitions).[16] Furthermore, an equilibrium in Markovian (or feedback) strategies is always subgame-perfect in a strong sense: Uniformly in the initial state.

4.2. The closed-loop or feedback solution

The results of this subsection are due to Levhari & Mirman (1980). Their solution proceeds via backward induction to derive the closed-loop equilibrium as the length of the horizon is increased, obtaining the infinite-horizon equilibrium as a limit. Here, we analyze instead the infinite-horizon problem directly to maintain comparability with the other cases analyzed in the present paper.

A feedback strategy γ is feasible if $0 \leq \gamma(x) \leq x, \forall x \geq 0$. We consider only the set of strategy pairs (γ^1, γ^2) that are jointly feasible here, in the sense that $0 \leq \gamma_1(x) + \gamma_2(2) \leq x, \forall x \geq 0$.[17]

Fixing a feasible strategy $\gamma(\cdot)$ by agent 2 (say), the best-response problem of agent 1 can be analyzed as follows. Let $V_\gamma(x)$ denote the optimal value agent 1 can obtain when agent 2 follows the consumption function $\gamma(\cdot)$ and the initial state is x. Then agent 1's best response strategy is the solution to the standard optimality equation (note that subscripts are dropped below for the sake of lighter notation, as the identity of the agents is clear):

$$V_\gamma(x) = \max_{0 \leq c \leq x - \gamma(x)} \left\{ \log c + \delta V_\gamma [(x - c - \gamma(x))^\alpha] \right\}. \tag{17}$$

The first order condition is (g denotes the best response consumption policy):

$$\frac{1}{g(x)} = \delta \alpha V'_\gamma [(x - g(x) - \gamma(x))^\alpha] (x - g(x) - \gamma(x))^{\alpha - 1}. \tag{18}$$

We can clearly write

$$V_\gamma(x) = \log[g(x)] + \delta V_\gamma [(x - g(x) - \gamma(x))^\alpha], \; for\, all\, x. \tag{19}$$

Differentiating with respect to x gives

$$V'_\gamma(x) = \frac{g'(x)}{g(x)} + \delta \alpha V'_\gamma [(x - g(x) - \gamma(x))^\alpha] (x - g(x) - \gamma(x))^{\alpha - 1} \times \left(1 - g'(x) - \gamma'(x)\right). \tag{20}$$

Substituting (18) into (20) yields the envelope relation

$$V_{\gamma}'(x) = \frac{1-\gamma'(x)}{g(x)}. \tag{21}$$

Substituting (21) into (18), the Euler equation follows

$$\frac{1}{g(x)} = \frac{\delta\alpha\,(x-g(x)-\gamma(x))^{\alpha-1}}{g\left[(x-g(x)-\gamma(x))^{\alpha}\right]}\left\{1-\gamma'\left[(x-g(x)-\gamma(x))^{\alpha}\right]\right\}. \tag{22}$$

Postulating a symmetric equilibrium with strategies linear in the stock, i.e. $\gamma(x) = g(x) = \lambda x$,(22) becomes

$$\frac{1}{\lambda(x)} = \frac{\delta\alpha\left[(1-2\lambda)x\right]^{\alpha-1}}{\lambda\left[(1-2\lambda)x\right]^{\alpha}}(1-\lambda). \tag{23}$$

This simplifies to $\lambda = \frac{1-\delta\alpha}{2-\delta\alpha}$, so the symmetric CL equilibrium consumption is

$$g_{cl}(x) = \frac{1-\delta\alpha}{2-\delta\alpha}x. \tag{24}$$

This is clearly jointly feasible, i.e. $0 \leq 2g_{cl}(x) < x$, for all $x \geq 0$. The corresponding value function may by computed, as before, via the optimality equation (17), upon substituting (24) for both agents:

$$V_{cl}(x) = \log\left(\frac{1-\delta\alpha}{2-\delta\alpha}x\right) + \delta V_{cl}\left[\left(\frac{\delta\alpha x}{2-\delta\alpha}\right)^{\alpha}\right]. \tag{25}$$

Substituting the guess $V_{cl}(x) = A\log x + B$ into (25) and identifying terms yields

$$V_{cl}(x) = \frac{\log x}{1-\delta\alpha} + \frac{\log\frac{1-\delta\alpha}{2-\delta\alpha} + \frac{\delta\alpha}{1-\delta\alpha}\log\frac{\delta\alpha}{2-\delta\alpha}}{1-\delta}. \tag{26}$$

The steady-state equilibrium level of ressource stock satisfies $\overline{x} = \left(\overline{x} - 2\frac{1-\delta\alpha}{2-\delta\alpha}\overline{x}\right)^{\alpha}$. Solving this equation for $\overline{x}$ yields

$$\overline{x}_{cl} = \left(\frac{\delta\alpha}{2-\delta\alpha}\right)^{\frac{\alpha}{1-\alpha}}. \tag{27}$$

5. The symmetric Pareto-optimal solution

In view of agents' symmetry, we focus on the unique symmetric solution out of the set of all Pareto-optimal outcomes. In other words, we consider as objective the sum of the two agents' utilities, as each is weighted by 1/2. This is clearly equivalent to the single-agent or monopoly problem:

$$Max \sum_{t=0}^{\infty} \delta^t \log c_t^i \; subject\ to\ x_{t+1} = (x_t - c_t)^\alpha , \quad t = 0, 1 \ldots \tag{28}$$

The solution may be obtained from the previous section by setting $\lambda \equiv 0$. The resulting optimality, envelope and Euler relations are respectively

$$V(x) = \max_{0 \le c \le x} \left\{ \log c + \delta V [(x - c)^\alpha] \right\}, \tag{29}$$

$$V'(x) = \frac{1}{g(x)}, \tag{30}$$

and

$$\frac{1}{g(x)} = \frac{\delta\alpha (x - g(x))^{\alpha - 1}}{g[(x - g(x))^\alpha]}. \tag{31}$$

Following the same method as before, the optimal consumption policy is then

$$g_{po}(x) = (1 - \delta\alpha) x, \tag{32}$$

and the corresponding total utility per agent is

$$V_{po}(x) = \frac{\log x}{1 - \delta\alpha} + \frac{\log(1 - \delta\alpha) + \frac{\delta\alpha}{1 - \delta\alpha} \log \delta\alpha 2}{1 - \delta}. \tag{33}$$

The steady-state equilibrium level of resource stock is easily seen to be

$$\overline{x}_{po} = (\delta\alpha)^{\frac{\alpha}{1-\alpha}}. \tag{34}$$

6. The trigger-strategy equilibria

Cave (1987) extends the analysis of Levhari & Mirman (1980) by incorporating history-dependent behavior in the specific form of threat or trigger strategies. We here review the main results of Cave's analysis. Unlike closed-loop (no memory) strategies, agents now have access to the entire past history of play and condition their actions on past observations at each time period. A trigger-strategy equilibrium is characterized by two phases. The first is a cooperative phase where players specify a mutually beneficial extraction path (e.g. a Pareto-optimal path) for the entire length of the game. The second is a punishment phase where a player would 'pull the trigger' and revert to the punishment part of his strategy as soon as he detects a deviation by the rival from the cooperative path. Following Cave, we consider the equilibria that result from the threat of reversion to the unique feedback equilibrium of Levhari-Mirman as calculated in Section 4 of the present paper, so that the resulting equilibria are clearly overall subgame-perfect.[18]

Cave exploits the special structure of this model by specifying consumption by agents in fractions h^i of the current resource stock. The agents select a pair $h = \left(h^1, h^2\right)$ of extraction rates in each period from the set of feasible extraction rate vectors given by $Y = \left\{h \in R_+^2 : \sum h^i \leq 1\right\}$, so that at date t, $c_t^i = h_t^i x_t$. By solving recursively for the resource stock at time t as a function of a given initial stock, x_0, we find $x_t\,(h, x_0) = (1 - H)^{\lambda(t)}\, x_0^{\alpha^t}$, where $H = h^1 + h^2$ and $\lambda\,(t) = \alpha\left(1 - \alpha^t\right) / (1 - \alpha)$. The present value to agent 1 for a given extraction rate vector $h = \left(h^1, h^2\right)$ is then

$$V\,(h, x_0) = \sum_{t=0}^{\infty} \delta \log\left[h^1 x_t\,(h, x_0)\right] = \frac{\Psi^1\,(h) + (1-\delta)\log\,(x_0)}{(1-\delta)\,(1-\alpha\delta)}. \tag{35}$$

where $\Psi^1\,(h)\,(1-\alpha\delta)\log\left(h^1\right) + \alpha\delta\log\,(1-H)$. A strategy for an agent is an infinite sequence of extraction rates $(f_0, f_1, \ldots)$, such that each f_t is a map from the history of states and actions up to time t to the set of feasible consumption levels.

Specifically, let $b = (b_1, b_2, \ldots)$ be a sequence of extraction pairs representing cooperative behavior and let z_t be the history of play, that is, the pairs of extraction rates during the first t periods. A closed-loop supported equilibrium strategy supporting cooperative behavior for player 1 is then

$$f_t^1\left(z_{t-1}\right) = \begin{cases} b_t^1, & if\ and\ only\ if\ \ z_{t-1} = \left(b_1, b_2, \ldots, b_{t-1}\right), \\ \frac{g_{cl}(x_t)}{x_t}, & otherwise, \end{cases} \tag{36}$$

where g_{cl} is agent 1′s equilibrium closed loop (no memory) strategy as given by (24). Given this, cooperative behavior is closed-loop supported for player 1 if and only if

$$V\,(h, x) \geq \max_c \left\{\log c + \delta V_{cl}\left[x - c - b^2 x\right]\right\}, \tag{37}$$

where the right-hand side of (37) stands for the maximal gain an agent would obtain by deviating from the cooperative path for one period, given that he would be punished from the next period on. If a pair of cooperative extraction rates for agent's 1 and 2, $d = \left(b^1, b^2\right)$, is closed-loop supported for both players, d belongs to the set of closed-loop supported extractions, CL. From (26) and (35) we have that a pair of extraction rates $h \in CL$ if and only if $(i, j = 1, 2;\ j \neq i)$

$$F^i\,(h) = \Psi^i\,(h) - (1-\delta)\log\left(1 - h^j\right) \geq K = F^i\left(\frac{g_{cl}\,(x)}{x}, \frac{g_{cl}\,(x)}{x}\right), \tag{38}$$

where $K = (1-\alpha\delta)\log\,(1-\alpha\delta) + \alpha\delta\log\alpha\delta - \delta\log\,(2-\alpha\delta)$.

The set CL of extractions can be determined by varying H in (38). This leads to a non-empty, convex, compact, and symmetric set. It follows from the very definition of CL, i.e. (38), that the closed loop equilibrium extractions $\left(\frac{g_{cl}(x)}{x_t}, \frac{g_{cl}(x)}{x_t}\right)$ always constitute the largest total extraction rate in CL. (See Cave (1987) for details.)

By symmetry, it is now clear from (32) and (38) that there exists a Pareto optimal solution in the set of closed loop supported extraction rates if and only if

$$F^1\left(\frac{1-\alpha\delta}{2},\frac{1-\alpha\delta}{2}\right)=F^2\left(\frac{1-\alpha\delta}{2},\frac{1-\alpha\delta}{2}\right)\geq K, \tag{39}$$

or

$$\delta\log(2-\alpha\delta)-\delta(1-\alpha)\log 2-(1-\delta)\log(1-\alpha\delta)\geq 0. \tag{40}$$

As trigger-strategy equilibria form an open set, there are also a priori uncountably many steady-state equilibra of the resource stock, $\overline{x}_{ts}$, all of which are clearly greater than $\overline{x}_{ol}$.

7. Comparative results

We state the main result of this paper in the form of:

Proposition 1. (i) The open-loop equilibrium coincides with the symmetric Pareto-optimal solution in that both lead to the same consumption path and utility levels.

(ii) The closed-loop (no memory) equilibrium leads to overconsumption (at every possible stock level) and to a lower total discounted utility level for each agent relative to the symmetric Pareto-optimal solution.

(iii) The trigger-strategy equilibrium set is such that any of its selections leads to a consumption level and a total discounted utility level that are respectively lower and higher than the corresponding levels under a closed-loop equilibrium. Furthermore, this equilibrium set includes a Pareto-optimal point if and only if (40) holds.

Proof. We provide only an outline of proof, as the actual analysis was conducted in the previous section. (i) follows from comparing (11) and (32). Total consumption is clearly the same in the OL equilibrium and the Pareto optimal solution. Hence, consumption per agent is also the same in the two solutions. Hence, the OL equilibrium and the Pareto optimal solution have the same total discounted utility per player. (ii) follows by comparing (24) with (32) and (26) with (33), respectively. Clearly, total consumption in the CL equilibrium is higher than the Pareto optimal level while total discounted utility is highest in the Pareto optimal solution. (iii) follows from the very definition of the trigger strategies used here.

Some important consequences of this Proposition are given next. The first compares the steady-state levels of resource under the different regimes (the proof is omitted, as it amounts to a simple comparison of the three quantities).

Corollary 2. The steady-state equilibrium levels of the resource stock satisfy:

$$\overline{x}_{cl}\leq\overline{x}_{ts}\leq\overline{x}_{po}=\overline{x}_{ol}. \tag{41}$$

Observe that Proposition 1 (ii) does not imply that actual consumption levels are higher at every time period in the CL than in the OL solution, the reason being that although consumption levels form a higher proportion of the current stock in the CL case, the stock gets consumed faster in the CL case, at least initially. In other words, under the CL regime, the agents consume a larger proportion of a resource stock that is shrinking faster or growing slower than in the OL regime. In fact, it is worthwhile to note the obvious fact that consumption levels are always unambiguously comparable at steady-state.

Corollary 3. Steady-state consumption levels are higher in the open-loop (or in the Pareto-optimal) solution than in the closed-loop solution.

Proof. The result follows from a simple computation, the steady-state consumption levels for the OL and the CL solutions being respectively

$$\frac{1-\delta\alpha}{2-\delta\alpha}\left(\frac{\delta\alpha}{2-\delta\alpha}\right)^{\frac{\alpha}{1-\alpha}} \quad \textit{and} \quad \frac{1-\delta\alpha}{2}(\delta\alpha)^{\frac{\alpha}{1-\alpha}}. \tag{42}$$

While OL equilibria typically fail the desirable property of subgame-perfection, the present case forms an exception. Indeed, due to the coincidence of the open-loop and the Pareto-optimal solutions, and the subgame-perfection of the latter, we have:

Corollary 4. The open-loop equilibrium of our dynamic game is subgame-perfect.

We now provide an extensive discussion of our conclusions. The main result sheds light on the important issue of how the extent of the tragedy of the commons depends on the information structure of the agents in common-property resource extraction. For the well-known model of Levhari & Mirman (1980), noncooperative open-loop behavior is outcome-equivalent to perfect cooperation, and thus leads to a first-best or Pareto-optimal outcome. Feedback or Markovian behavior leads to the most pronounced version of the tragedy of the commons, of all cases considered. Finally, history-dependent behavior in the form of trigger-strategies leads as usual (due to Folk-Theorem-type arguments) to a myriad of solutions, all reflecting an intermediate outcome, ranging from the feedback solution all the way to possibly the first-best situation (depending on the parameters of the problem).

We now attempt to provide an intuitive understanding of these results. In the open-loop case, agents are deprived of the possibility of observing the current stock and reacting in a dynamic sense to the catches of the rival. This lack of awareness of the period-by-period effects of the rival's catches on the resource stock leads the agents to avoid the over-consumption that is otherwise inherent to noncooperative resource extraction. That this leads all the way to Pareto optimality is rather surprising, and is due to the special structure of the model (the sole presence of the dynamic externality). Under feedback behavior, the agents are fully aware of the consequences of the rival's catch on the resource level, and react accordingly, causing the standard tragedy of the commons to be fulfilled. Feedback information, on the other hand, is not sufficient to allow agents to make threat-secured agreements, as these require history-dependent information for the agents to be able to detect deviations from agreed-upon extraction paths. Increasing the agents' information levels as Cave did allows them to make such cooperative agreements and reduce the extent of the tragedy of the commons, possibly all the way to first-best.

Our main result may be succinctly phrased as follows: The efficiency of equilibria in common-property resource extraction is not monotonic in the level of information on which the players' decisions are based. Rather it is U-shaped, with a global maximum at the lowest level of information (open-loop strategies) and possibly another at the maximum level of information (i.e. history-dependent strategies).

While the coincidence of the open-loop and the Pareto-optimal solutions is easily shown to extend to general utility and growth functions and asymmetric agents, the trigger-strategy equilibrium is much harder to characterize at such a level of generality, as suggested by Cave's specific analysis on which we relied here. Aside from the latter difficulties, the main conclusions of the present analysis are robust provided the only externality at work is the dynamic externality.

The subgame-perfect property of the open-loop equilibrium turns out to be of possible practical relevance to a possible regulatory application of our analysis.

8. A policy implication

A natural question at this point is whether our conclusions offer any policy-relevant insights. Consider a government regulator with a mandate to prevent or limit over-extraction of a natural resource (such as fish or game) over a specified time horizon. One possible way to proceed[19] would be to require the consuming agents to (independently and simultaneously) submit individual consumption plans over the entire time horizon with the understanding that they will remain committed to these announcements, and penalized for any future deviations. According to our analysis, if the dynamic externality is the only one present and if uncertainty plays no role in the context at hand,[20] this noncooperative mechanism would lead to a Pareto-optimal solution over the specified time horizon. This offers an interesting possible alternative to the usual centrally-negotiated cooperative solution consisting of allocating quotas to the agents, which is worth exploring.

If the specified time horizon is long enough, and in view of the coincidence of the open-loop and the Pareto-optimal solutions, we know from Easley & Spulber (1981) that this rolling-horizon solution should remain fairly close to the Pareto-optimal solution of the infinite-horizon problem. More broadly, while we agree with the now well-known assessment that open-loop behavior in a long-term game is generally inappropriate as an approximation of real-life volontary behavior in most economic settings, we argue here that open-loop behavior may quite naturally arise from Pareto-improving regulation in some specific settings. Indeed, forcing a commitment from all agents and policing their compliance with the announcements may be welfare-improving for all, in some situations characterized in particular by the independence of the one-period rewards on the rival's actions.

A key property of such a mechanism is that it would essentially be cheat-proof, at least in a unilateral sense, since if all the other agents employ open-loop strategies, a given agent cannot improve his payoff by following some more complex strategy, so he might as well do the same (this fact is an important property of open-loop equilibrium and was discussed in Section 3.1). Thus the monitoring regulator need only worry about deviations from the announced paths that involve more than one agent. This potentially makes the monitoring problem less demanding than in the case of regulation via centrally-allocated extraction quotas. Another desirable feature of this mechanism is that it is likely to be less demanding to administer, as it allows policy-makers to avoid the typically lengthy and strenuous negotiations that are often associated with quota negotiations.[21] In the language of industrial regulation (see e.g. Perry 1984), our proposal can be regarded more as an instance of structural policy than of behavioral policy. In other words, the aim of the regulator is to influence the strategic environment of the regulated agent by constraining his allowable strategy set, and not to interfere directly with the agents' postulated behavior of individual payoff maximization (taking the rival's action as given). In contrast, central quota assignment is an instance of the latter mode of regulation.

The notion that commitment is important in many classes of noncooperative games has been around for quite some time, and this is only a new instance (at least to us) of this general fact.[22] Yet, as simple as this idea is in the context of noncooperative dynamic resource exploitation, it does not seem to have been proposed before.

9. Information structure and resource depletion

In this section, we elaborate on the observation that the differences in noncooperative resource extraction between open-loop and Markovian behavior can in some cases be so pronounced as to lead to the long-run resource sustainability in the former case and depletion in the latter case. To this end, it is necessary to postulate another form of natural growth function, as (1) always leads to a strictly positive resource stock at steady-state, regardless of the information structure.

Consider the log utility function (2) and a linear growth for the resource evolution

$$x_{t+1} = ax_t, \text{ with } a > 1. \tag{43}$$

Under the OL and the CL information structures, the same steps as before lead to the respective unique linear equilibrium (per-agent consumption) strategies[23]

$$g_{ol}(x) = \frac{1-\delta}{2}x, \forall x \geq 0 \textit{ and } g_{cl}(x) = \frac{1-\delta}{2-\delta}x, \quad \forall x \geq 0. \tag{44}$$

The corresponding resource dynamics under the OL and CL equilibria is given respectively by the following two dynamical systems

$$x_{t+1} = \frac{a\delta}{2-\delta}x_t \textit{ and } x_{t+1} = a\delta x_t. \tag{45}$$

It is easy to verify, using the above equations, that under a moderate natural growth rate of the resource, i.e. $\frac{2-\delta}{\delta} < a < \frac{1}{\delta}$, we have $\lim_{t\to\infty} x_t = \infty$ under the OL solution, and $\lim_{t\to\infty} x_t = 0$ under the CL solution. In other words, the resource grows without bound under noncooperative open-loop extraction and is depleted over time under feedback extraction, a dramatic difference in many respects.

On the other hand, in case of a high growth rate $\left(a > \frac{1}{\delta}\right)$, $\lim_{t\to\infty} x_t = \infty$ in both cases, while in case of a low growth rate $\left(a < \left(\frac{2-\delta}{\delta}\right)\right)$, $\lim_{t\to\infty} x_t = 0$ in both cases. The underlying computational details are simple and left to the reader.

10. Conclusion

Adopting the simple framework proposed by Levhari & Mirman (1980), this paper has provided a comparative analysis of the effects of the information structure on the equilibrium extraction path and overall utility levels for a common-property natural resource. To allow a more intuitive understanding of the relationship between information and extraction paths, we derive all solutions using a dynamic programming approach. The main result is that equilibrium efficiency is U-shaped in the level of information available to the agents, with the open-loop equilibrium being Pareto-optimal. We argue that the main result extends to more general functional forms, provided the special structure of the Levhari-Mirman model is preserved, with the absence of a market-type externality. This result is of broader interest for the application of dynamic games in economics, where similar information effects on the behavior and welfare of agents are often investigated.

The results suggest a mechanism design approach to the problem of regulation of common-property resource extraction: If agents could be induced to behave according to open-loop strategies, the equilibrium would be self-enforcing, and thus require relatively minor monitoring. This proposal would also save on the usual tedious bargaining and dispense with the reliance on truthful information revelation, two of the features that always accompany the process of quota-setting.

It must of course be clearly stressed that this provocative proposition is valid only in rather special environments as reflected in the special structure of the model used here. On the other hand, there may exist other more complex environments where, for reasons similar to the ones driving our results here, the open-loop solution welfare-dominates the closed-loop solution without satisfying Pareto-optimality. Then, our proposed regulatory scheme would alleviate the tragedy of the commons without fully resolving it as in the present analysis. These issues would be important natural extensions for future research.

Acknowledgements

The authors are grateful to Charles Figuieres, Grzegorz Halaj, Jean-Francois Mertens, Philippe Michel and Leonard Mirman for helpful comments on the topic of this paper. We also thank an anonymous referee for helpful comments. Rabah Amir also thanks the Center for Industrial Economics at the University of Copenhagen for their stimulating hospitality as well as financial support while the research reported here was carried out.

Notes

1. In particular, the pioneering work of Gordon (1954) is one study that immediately comes to mind. See Clark (1999) for further references.
2. Henceforth we refer to this type of equilibrium using any of the three terminologies.
3. It is worthwhile here to fully specify what overconsumption means here: The resulting consumption functions are higher at every stock level for the Markovian equilibrium than for the first-best solution. Since this implies that the resource stock declines faster under the former regime than under the latter, actual consumption levels will eventually be lower under the Markovian equilibrium.
4. While the idea that strategic interaction over time with recall of past play can induce more cooperative behavior on the part of players was well-established in the theory of repeated games, the economic applications of the theory of dynamic games (with a state variable) have generally restricted attention to history-independent behavior. In this sense, Cave (1987) pioneered this type of analysis in economic applications of dynamic games. The complexities involved in such a task explain the specific nature of the model that Cave adopted, as will be seen below.
5. In certain continuous time formulations of resource extraction models, the Pareto optimality of open-loop equilibria has been known for some time, see Chiarella et al. (1984) and Dockner & Kaitala (1989). On the other hand, to the best or our knowledge, there is no discrete-time counterpart in the literature. Furthermore, it is well-known that fundamental differences often exist between the continuous-time and corresponding discrete-time formulations, including in particular with regard to the issue of existence of Markovian equilibrium.
6. See Reinganum & Stokey (1985) for an analysis of the role of the period of commitment to extraction paths in a dynamic game.
7. In other words, an open-loop Nash equilibrium (i.e. one that is derived by allowing only open-loop deviations for the players) remains an equilibrium if the players are allowed to use any more general classes of strategies, such as Markovian or history-dependent strategies.
8. It is worthwhile to point out here that, in view of the absence of a supranational governing body to ensure direct compliance with assigned quotas, OPEC has to rely on threat-secured compliance instead. Indeed, sustained deviations in the form of over-production in the past have sometimes triggered reactions (e.g. by Saudi Arabia) that resulted in a glut in the oil market. Thus, OPEC would fit quite nicely a Cave-style analysis. On the other hand, a model of OPEC would generally not

satisfy the result that the open-loop equilibrium is Pareto-optimal, since OPEC reflects a market externality and no dynamic externality (as oil wells are the private property of the states involved), in complete contrast to the Levhari-Mirman world.

9. Throughout, we denote agents by superscripts and calendar time by subscripts.
10. These are defined by the facts that the one-period reward and the next state are time-invariant and depend only on the current state and actions.
11. The model of Levhari & Mirman (1980) has been investigated in a general framework with a stochastic growth law for the resource in Amir (1996).
12. No general results are known about the comparison of equilibrium payoffs under open-loop vs. Markovian (or closed-loop no-memory) behavior. Examples can be given for either type of equilibrium to be better than the other type, in terms of resulting equilibrium payoffs.
13. In the economics of natural resource exploitation and sustainability, studies that rely on the open-loop information structure tend to be older. They include, among many others, Salant (1976), Lewis & Schmalensee (1980), and Dasgupta & Heal (1979).
14. Stating these strategies in this form is very convenient for the purposes of this paper, as this allows for direct comparisons with the other cases. The alternative, i.e. giving an infinite vector of consumption levels depending only on x_0, would make this comparison less transparent. Thus we deviate from standard practice in formulating an open-loop solution in this manner.
15. With the above terminology originating in the control theory literature, the standard game-theoretic way of referring to such strategies is as Markov-stationary, see e.g. Amir (1996). The latter terminology is more consistent with the way we defined the general class of games at hand, but is less prevalent in the resource economics literature.
16. It is of interest to observe that these important justifying arguments, as well as the so-called one-shot deviation principle, follow directly from the general properties of dynamic programming.
17. In other words, we are really considering a generalized game, in Debreu's (1952) terminology.
18. One may appropriately interpret this set-up as a self-enforcing cooperative agreement in a noncooperative framework. By its very construction, such a trigger-strategy equilibrium has the property that no agent wishes to defect from either the agreement or the threats associated with out-of-equilibrium behavior, hence Cave's 'Cold Fish War' reference.
19. There is an extensive literature on the regulation of natural resources via taxes or quotas: see e.g. Dasgupta & Heal (1979) and Bergstrom (1982).
20. Recall that, as can be intuitively expected, the open-loop information structure loses many of its equilibrium properties in the presence of uncertainty in the state transition law (see Section 3).
21. For resource conservation, the main approach taken by regulation authorities is to stipulate a total allowable catch for the whole fishery based on biological studies on the actual stock population. Individual quotas are subsequently awarded, usually, to agents with a history of participation in the fishery (and through negotiations at the international level.). To implement a cost-effective fishery individually transferable quotas are increasingly being used. Under this system total quotas being issued are still determined by the regulation authorities, while trade in quota units is allowed so that they go to agents who value them the most. Such management systems of centralized control are used for most regulated fish populations. Clark (2000) provides a broad discussion of traditional and alternative fishery management techniques.
22. For instance, within the theory of multiperiod mechanism design, it proves essential for the principal to be able to commit to a long term contract at the beginning of a relationship. Otherwise, serious ratchet effects may arise leading to substantial difficulties for contract theory, see Laffont & Tirole (1993 [chapter 9]).
23. Since $\log ax_t = \log a + \log x_t$, the equilibrium policies are independent of a. With $a = 1$, one can obtain these policies from the corresponding solutions in previous sections by setting $\alpha = 1$.

References cited

Amir, Rabah. 1996. Continuous stochastic games of capital accumulation with convex transitions. Games and Economic Behavior 16: 111–131.

Amir, Rabah. 2003. Stochastic games in economics II: an overview. Pp. 455–470 in A. Neyman & S. Sorin (ed.) Stochastic Games and Applications. Kluwer, Dordrecht.

Basar, Tamer & Geert Olsder. 1999. Dynamic noncooperative game theory. SIAM Classics (revised edition).

Bergstrom, Theodore C. 1982. On capturing oil rents with a national excise tax. American Economic Review 72: 194–201.

Cave, Jonathan. 1987. Long-term competition in a dynamic game: The cold fish war. Rand Journal of Economics 18: 596–610.

Chiarrella, Carl, Murray Kemp, NgoVan Long & Koji Okuguchi. 1984. On the economics of international fisheries. International Economic Review 25: 85–92.

Clark, Colin. 1999. Renewable resources: fisheries. Pp. 109–121 in J.C.J.M. den Bergh (ed.) Handbook of Environmental and Resource Economics. Edward Elgar Publishing Limited, Cheltenham.

Datta, Manjira & Leonard Mirman. 1999. Externalities, market power, and resource extraction. Journal of Environmental Economics and Management 37: 233–55.

Dasgupta, Partha & Geoffrey Heal. 1979. Economic theory and exhaustible resources. Cambridge University Press, Cambridge.

Debreu, Gerard. 1952. A social equilibrium existence result. Proceedings of the National Academy of Sciences of the USA.

Dockner, Engelbert & Vaijo Kaitala. 1989. On efficient equilibrium solutions in dynamic games of resource management. Resources and Energy 11: 23–34.

Easley, David & Daniel Spulber. 1981. Stochastic equilibrium and optimality with rolling plans. International Economic Review 22: 79–103.

Fisher, Ronald & Leonard Mirman. 1996. The complete fish wars: biological and dynamic interactions. Journal of Environmental Economics and Management 30: 34–42.

Fudenberg, Drew & Jean Tirole. 1986. Dynamic models of oligopoly. Harwood Academic Publishers, Chur, Switzerland.

Gordon, H. Scott. 1954. The economic theory of a common-property resource: the fishery. Journal of Political Economy 62: 124–42.

Laffont, Jean-Jacques & Jean Tirole. 1993. A theory of incentives in procurement and regulation. MIT Press, Cambridge, MA.

Levhari, David & Leonard Mirman. 1980. The great fish war: an example using a dynamic Cournot-Nash solution. Bell Journal of Economics 11: 322–344.

Lewis, Tracy & Richard Schmalensee. 1980. On oligopolistic markets for nonrenewable natural resources. Quarterly Journal of Economics 95: 475–491.

Perry, Martin. 1984. Scale economies, imperfect competition and public policy. Journal of Industrial Economics 32: 313–333.

Reinganum, Jennifer & Nancy Stokey. 1985. Oligopolistic extraction of a non-renewable common-property resource: the importance of commitment in dynamic games. International Economic Review 26: 161–173.

Salant, Steven. 1976. Exhaustible resources and industrial structure: a Cournot-Nash approach to the world oil market. Journal of Political Economy 84: 1079–93.