Cooperation in Dynamic Games: An Experimental Investigation*

Job Market Paper

Emanuel I. Vespa

New York University

November 28, 2011

Abstract

Dynamic games describe intertemporal environments in which stage-game payoffs change depending on an underlying state variable. Folk-theorem arguments hold for dynamic games with infinitely repeated interactions, so that there are equilibria that support the first best. However, applications of dynamic games overwhelmingly focus on equilibria that are history-independent (Markov) and typically do not support cooperative outcomes. In this paper, I study experimentally whether cooperation arises in a stylized version of a canonical dynamic game, the dynamic commons problem. In addition, I uncover which strategies are used to support cooperative agreements, and inquire whether behavior can be captured with Markov strategies. Experimental findings show that most behavior is consistent with Markov strategies typically characterized in the literature. However, when the degree of opportunism is constrained, there is a significant presence of cooperation, at levels comparable to previous evidence from prisoners' dilemma experiments.

*I want to thank CESS at NYU for funding this project and Ryan Booth, Andrew Caplin, Ignacio Esponda, Alessandro Lizzeri, Ariel Rubinstein, David Pearce, Andy Schotter, Matan Tsur, Sevgi Yuksel, Alistair Wilson, CREED-CESS seminar participants for very useful comments. Special thanks to my adviser Guillaume Fréchette for very helpful guidance, support and patience along the development of this project. Address: 19 W.4th St. Off. 502, New York, NY 10012 , emanuel.vespa@nyu.edu.

1 Introduction

Dynamic games describe intertemporal environments in which stage-game payoffs change depending on an underlying state variable that evolves endogenously.[1] Folk-theorem arguments hold for dynamic games with infinitely repeated interactions, so that cooperative equilibria can be supported (Dutta 1995). However, applications of dynamic games overwhelmingly focus on equilibria that are history-independent and typically do not support cooperative outcomes. In this paper, I study experimentally whether cooperation arises in a stylized version of a canonical dynamic game, the dynamic commons problem. In addition, I uncover which strategies are used to support cooperative agreements, and inquire whether behavior can be captured with history-independent strategies.

A prominent question motivating this study relates to the connection between repeated and dynamic games. Repeated games can be thought of as dynamic games in which the state variable does not change. The repeated games literature has long recognized theoretically and experimentally that cooperation can be supported if behavior conditions on past play and punishes defection. Applications of dynamic games, though, center their attention on Markov perfect equilibria (MPE): subgame-perfect equilibria in which agents use strategies that only depend on the state and ignore past play (Markov strategies). Efficient outcomes typically cannot be supported with Markov strategies and resulting inefficiencies can be quite large.[2] The experimental design captures this tension between Markov behavior and the first best, as the efficient outcome will only be supported with history-dependent strategies. Experimental evidence identifying under which conditions cooperation is likely to occur can be useful to assess when the Markov restriction is appropriate.

To study cooperation in a dynamic environment, I implement a stylized dynamic commons problem in the laboratory. The model studies the behavior of agents that share the use of a productive asset. It has proved to be a useful tool to study economic

[1]Dynamic games have been widely studied by applied researchers in several fields. A few examples include: Industrial Organization (Ausubel and Deneckere 1989; Ericson and Pakes 1995), Public Finance (Battaglini and Coate 2007, 2008), Labor (Coles and Mortensen 2011), Political Economy (Coate and Morris 1999; Dixit et al. 2000; Acemoglu et al. 2010), Macroeconomics (Laibson 1997), Environmental Economics (Dutta and Radner 2006), Economic Growth (Bernheim and Ray 1987; Benhabib and Rustichini 1996) and Applied Theory (Rubinstein and Wolinsky 1990; Bergemann and Valimaki 2003; Hörner and Samuelson 2009)

[2]For example, the unique subgame perfect equilibrium using Markov strategies in the prisoners' dilemma prescribes permanent defection.

problems (i.e. renewable natural resources) and typically serves as the example for specialized textbooks to introduce dynamic games (see Fudenberg and Tirole 1991; Mailath and Samuelson 2006; Acemoglu 2011).[3] In each period of time, two agents simultaneously decide to extract either a high ($a_H = \frac{1}{2}$), medium (a_M) or low (a_L) percentage of a common access resource. The stock remaining after consumption reproduces at rate $r > 0$ and determines the stock available next period. In other words, the stock captures the 'state' of the game. In this environment, the current decision of one agent does not influence the current payoffs of the other.[4] However, the lower aggregate extraction is in the present, the higher the stock will be the next period, leading to higher future payoffs. Parameters are set so that efficient consumption involves always selecting a_L. Both players extracting at a low rate with no punishments upon deviation is not subgame-perfect, as there are high incentives to free-ride, consume at a higher current level, and foster the 'tragedy of the commons'. To assess the magnitude of inefficiencies, the literature has focused on comparisons to MPEs. One such MPE prescribes extraction at a_M for all levels of the stock. There is also an MPE in which both agents select a_H in the first period and the stock is immediately depleted. The tragedy of the commons can be avoided with history-dependent strategies that punish deviations from a_L, for instance, with permanent extraction at a_M. The purpose of the construction is to make it feasible to support cooperation with simple subgame-perfect punishments.

A prominent feature of the environment is the presence of opportunistic behavior at different levels. Higher extraction leads to higher inefficiencies. In common access resources, excessive free-riding can be controlled for with monitoring policies that deter high exploitation activities.[5] The experimental design captures the aim of

[3]The dynamic commons problem cast as a differential game was first presented by Lancaster (1973) to study investment decisions in an economy with two large agents: workers and capitalists. The discrete-time formulation is from Levhari and Mirman (1980), whose motivation is to better understand conflicts in fisheries, such as the then contemporary Cod Fish war between England and Iceland. Since then their model has also been used to study other exhaustible natural resources (see Clark (1990)) and other phenomena in economics. For example, Tornell and Velasco (1992) use it to study differences in capital flows between countries, Benhabib and Rustichini (1996) to model economic growth and Reinganum and Stokey (1985) to understand oligopoly behavior. In anthropology, the work of Diamond (2006) highlights that depletion of natural resources is the most important human cause for the collapse of past societies.

[4]Instant utility is given by the log of the extracted share, but if the stock is depleted agents receive no payoff. Future periods are discounted and each agent maximizes their own intertemporal payoffs.

[5]For example, establishing 'Regional Fishery Management Councils to steward fishery resources

such policies, in a very stylized way, with a treatment in which a_H is not allowed. Precluding the choice of a_H may help cooperation as there are less alternatives for opportunistic behavior. However, since a_H can be used as a harsh punishment to discourage defection, its presence can also foster cooperation. Additionally, the treatment with two available actions (a_L, a_M) will allow for a more direct comparison to previous experiments in the prisoners' dilemma.

The reproduction rate r influences how soon the benefits and costs of cooperation are experienced. As a benchmark, consider an agent who decides to cooperate in the prisoners' dilemma. The consequences of the other's actions are readily experienced in the present: benefits if the other reciprocates, costs if the other defects. In the dynamic commons game, behavior in the present affects agents in the future. When the reproduction rate is low, the consequences of excessive extraction are promptly borne in the future, while those effects are less pressing if the reproduction rate is high. This tension will be studied in the laboratory with treatments on the reproduction rate. Does a low reproduction rate help cooperation? Perhaps a delicate environment induces extreme over-exploitation (choices at a_H). Evidence from marine fisheries, for example, suggests that higher reproduction rates lead to higher exploitation (Hutchings 2000).

The setup with two players, at most three actions, and an integer-valued state space provides an environment that allows for computational challenges in the implementation of dynamic games to be overcome. Calculations required to make a choice in a dynamic game can be quite demanding, even with simple state and action spaces. The software interface is designed so that subjects at any point in the experiment can easily compute and compare future hypothetical paths before they make a decision. The infinite time horizon is implemented in the laboratory as an uncertain time horizon using a procedure developed by Cabral et al. (2011).[6] To help learning, subjects

through the preparation, monitoring, and revising of plans, which (A) enable stake holders to participate in the administration of fisheries, and (B) consider social and economics needs of states' is one of the purposes of the Magnuson–Stevens Fishery Conservation and Management Act, which governs marine fisheries management in the U.S. The Regional Fishery Management Councils establish catch monitoring programs, which are then implemented by the U.S. Coast Guard.

[6]See Section 3 for details. In the standard random termination procedure, first developed by Roth and Murnighan (1978), there is no discounting, but in each period the game is over with probability δ. With this protocol, even for high discount rates, the likelihood of having very short games is high. Since cooperation involves low payoffs in the first periods and high returns once the stock is large enough, it is important that subjects experience long games. The modified random termination method involves T periods played with certainty and discount δ, and random termination with

will play nine repetitions of the supergame, being randomly rematched with a new partner each time.

The recovered information allows to estimate the presence of a large set of strategies using a procedure first developed by Dal Bó and Fréchette (2011) and also used in Fudenberg et al. (2011). However, since many strategies can implement the first best, in principle none of them can be distinguished from each other if they succeed in supporting cooperation. I deal with this identification problem by introducing a one-period ahead strategy method that allows to recover some of the choices that would have been made had the history of play been different.

The experimental design involves four treatments, combining high and low reproduction rates and whether a_H is allowed or not. Using a between subjects design I conducted three sessions of each treatment for a combined total of 160 NYU undergraduates.

The main results show that most behavior in all treatments can be properly explained with Markov strategies typically characterized by the literature. However, when opportunistic behavior is constrained (a_H not allowed), I observe cooperation at levels comparable to previous experiments on the prisoners' dilemma. To provide an overview of cooperation patterns I first focus on choices at the aggregate level. The data shows clear differences across treatments in the frequency with which a_L is chosen (cooperation rate). Cooperation rates are relatively low when there are more alternatives for opportunistic behavior. In the second half of the experiment, the cooperation rate is below 8% for both levels of the reproduction rate when a_H is allowed. When the levels of opportunism are constrained, the cooperation rates are two or three times higher despite the fact that eliminating a_H substantially reduces the potential gains from cooperation (i.e. from 96% to 29% if the reproduction rate is low).[7] In particular, when the reproduction rate is low, the cooperation rate is at 25% and this is comparable to the cooperation level in prisoners' dilemma experiments where the potential gains from cooperation are similar.[8] In addition, evidence shows that cooperation may be higher when the reproduction rate is low. For instance,

probability δ and no discount starting at round $T + 1$.

[7]Gains from cooperation can be computed by contrasting the discounted payoffs of cooperation and permanent defection (see section 3 for details). The measure can be used to compare with other experimental studies of cooperation. For example, for gains of 92%, Dal Bó and Fréchette (2011) report cooperation rates of 76% for the infinitely repeated prisoners' dilemma.

[8]For gains from cooperation at 28%, Dal Bó and Fréchette (2011) report a cooperation rate of 25.6%.

when a_H is not allowed, the cooperation rate drops from 25% to slightly above 13% as the reproduction rate increases. This suggests that higher reproduction rates may lead to higher exploitation.

In all treatments the most popular strategy is Markovian, namely extraction at the a_M level, suggesting that the degree of the 'tragedy of the commons' is close to the standard theoretical prediction. However, across treatments there are differences in the magnitude of its popularity. When high extraction is allowed, there is no evidence of systematic presence of history dependent strategies. Permanent extraction at the a_M level explains approximately 70% of observed behavior. When a_H is not allowed and the reproduction rate is low, extraction at a_M drops to about 50% and the second most popular strategy (at approximately 30%) is grim-trigger with permanent reversion to a_M upon any defection. If the reproduction rate is high, the most popular history dependent strategy is tit-for-tat.[9] This suggests that in environments that are more fragile (lower reproduction rates) cooperative agreements are also more fragile (no return to cooperation upon any deviation).

Previous related experiments, which focus on a different set of questions, have both, static and dynamic sources of inefficiencies. The game under study in this paper focuses on a tension introduced by dynamic games. If each stage-game was considered on its own, there would be no disagreement between a planner and players acting in their own best interest: both would select the highest feasible extraction level. The trade-off at play results because decisions in the present lead to different states in the future. Herr et al. (1997) study a finite-horizon commons problem with a dynamic and a static externality: higher extraction in the present increases current extraction costs for all agents. They find high levels of myopic behavior that lead to higher inefficiencies compared to the case when only the static externality is present. Battaglini et al. (2011) are the firsts to examine the provision of a durable public good under legislative bargaining. In each period, committee members decide to either allocate money to private consumption or a public good that accumulates over time. They compare different committee decision rules in the laboratory and report behavior close to standard Markov predictions.[10] Saijo et al. (2009) implement in

[9]Upon a deviation, a subject following tit for tat punishes by choosing their partner's last period choice. This strategy is not subgame perfect.

[10]The sources of inefficiencies in Battaglini et al. (2011) can be static and dynamic. In each period of time one member of the committee is randomly selected to make a proposal on how to divide a fixed budget into private allocations for each member and contributions to a public good. The voting

the laboratory the global warming model of Dutta and Radner (2006), which shares some features with the dynamic commons game. Their subjects play one repetition of the supergame and the paper reports choices close to the first best even if high levels of opportunistic behavior are allowed. In the treatment that allows for a_H, I also find high cooperation rates in the first few repetitions, but it quickly vanishes with learning.[11]

Finally, this paper also contributes to the literature on experiments in common pool problems, which deals primarily with the static version. Ostrom et al. (1994) report evidence consistent with overextraction (see chapter 5), but an important finding in this literature is that cheap-talk communication allows to substantially reduce inefficiencies (see Ostrom and Walker 1991).[12] This paper explores an alternative channel to cooperation, namely repeated interactions. Indeed, the findings suggests that if excessive opportunistic behavior is not allowed, repeated interactions allow for higher levels of cooperation.

The theory underlying the experimental design is presented in detail in section 2. Section 3 presents the experimental design, while the data analysis is conducted in Section 4. Section 5 discusses the main findings and presents future avenues of research.

rule is a treatment variable and establishes the number of votes required for the proposal to pass. If the proposal does not pass, a status-quo is enacted. Consider first a stage-game in isolation. The choices of a planner and a committee member will differ, in a static sense, as the former maximizes aggregate welfare and the latter is concerned only with the qualified majority required to pass the proposal. With an infinite time horizon and a public good that accumulates in time, dynamic sources of inefficiencies also appear. It is encouraging that experimental results in this environment with several opportunities for free-riding also provide support for Markov behavior.

[11]Chermak and Krause (2002) and Fischer et al. (2004) implement common pool problems in which players are replaced in time to mimic intergenerational features. Both setups use a finite time horizon and report subjects making suboptimal choices. Other experiments have implemented dynamic decision problems in the laboratory (Ramsey growth model). These papers find support for the comparative statics, although subjects do not perfectly optimize (see Hey and Dardanoni 1988, Noussair and Matheny 2000 and Lei and Noussair 2002).

[12]A large list of experiments that study such environment is surveyed in Sturm and Weimann (2006) and Ostrom (2010).

2 Theory

This section briefly develops a stylized version of the dynamic commons problem as presented by Levhari and Mirman (1980).[13] I first discuss some aspects of the game using the original framework and then introduce some modifications.

In each period $t \in \{0, 1, ...\}$ two agents, $i \in \{1, 2\}$, decide simultaneously on an extraction share ($a_t^i \in A = [0, 1]$) from the current stock level of a resource, $s_t \in \mathbb{R}_+$. Consumption in each period (C_t^i) is thus given by: $C_t^i = a_t^i s_t$. If there is not enough stock to satisfy both agents' claims ($a_t^1 + a_t^2 > 1$), then the existing stock is split evenly and the game ends.[14] The stock starts at known $s_0 \in (0, \infty)$ and evolves according to:

$$s_{t+1} = (1 + r)\left[(1 - a_t^1 - a_t^2)s_t\right] \tag{1}$$

The term in brackets captures the period savings, which grow by a a given reproduction rate $r \in (0, \infty)$ to determine next period's stock level, so that $s_t \in S = \mathbb{R}_0^+$ for all t. Consumption in t is valued according to $U(C_t^i) = \ln(C_t^i)$. The agents discount the future at a fixed period rate given by $\delta \in (0, 1)$. Let $\mathbf{a_t} = (a_t^1, a_t^2) \in A^2$, then the history of the game at period t is given by: $h_t = (s_0, (\mathbf{a_0}, s_1), ..., (\mathbf{a_{t-1}}, s_t)) \in S \times (A^2 \times S)^t \equiv \mathcal{H}_t$. Let $\mathcal{H}_0 \equiv S$ and $\mathcal{H} = \cup_{t=0}^{\infty} \mathcal{H}_t$. A pure strategy for player i is a mapping: $\sigma^i : \mathcal{H} \rightarrow A$, associating an action with each history. Let Φ^i define the space of all such mappings for subject i. The payoff associated to a strategy profile $\sigma = (\sigma^1, \sigma^2)$ is given by:

$$U^i(\sigma) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \ln C_t^i(\sigma) \tag{2}$$

Where $C_t^i(\sigma)$ is the consumption level for subject i in period t implied by σ.

Definition. A strategy profile σ is a Nash equilibrium if $U^i(\sigma) \geq U^i(\tilde{\sigma}^i, \sigma^{-i})$ for all $\tilde{\sigma}^i \in \Phi^i$ and for all i. A strategy profile σ is a subgame-perfect equilibrium if for any history $h_t \in \mathcal{H}_t$ ending in state s, the continuation strategy $\sigma|_{h_t}$ is a Nash equilibrium of the continuation game, that starts in state s.

Markov Perfect Equilibria are of central interest in this game and defined next.

[13]For a detailed analysis of this game see Sundaram (1989), Dutta and Sundaram (1992) and Dutta and Sundaram (1993).

[14]This rule has no effect on the solution concepts used later.

Definition. A strategy profile σ is a Markov strategy if for any two histories $h_t, \tilde{h}_s$ terminating in the same state, $\sigma(h_t) = \sigma\left(\tilde{h}_s\right)$. The strategy profile σ is a Markov perfect equilibrium (MPE) if σ is a Markov strategy and a subgame-perfect equilibrium.

A Markov strategy, thus, requires the player to choose an action depending only on the current level of the state.[15] There are many MPEs in this game (Sundaram 1989), but as with most dynamic games attention is focused on symmetric equilibria ($\sigma = \sigma^i$ for $i = 1, 2$). Levhari and Mirman (1980) characterize a symmetric MPE using first backwards induction to solve the finite case and then showing that the construction is also an equilibrium in the game with an infinite horizon. Since then this MPE has received special attention in the literature and is characterized next.[16]

Proposition 1. *Levhari-Mirman MPE (LM-MPE): The strategy profile $\sigma(s) = a_M s$, with $a_M = \frac{1-\delta}{2-\delta}$ is a MPE.*

The Markov strategy in the LM-MPE is a linear function of the state with no intercept and a slope that depends on the discount rate δ. A higher discount rate implies a more valuable future and thus less extraction in the present. As $\delta \to 0$, $a_M \to \frac{1}{2}$, so that aggregate consumption depletes the stock in the first period. Notice, in particular, that a_M does not depend on the reproduction rate r. As can be verified in Appendix B this is a property that results from using a logarithmic instant utility function.

Both players announcing the maximum extraction rate for all s is also an MPE since there are no profitable deviations for any agent.[17] Equilibria in which there is positive consumption in all periods are usually referred to as 'interior'. Sundaram (1989) shows that any other 'interior' MPE of this game (symmetric or not) satisfies that the evolution of stock must be monotone. In particular, this implies that there are no MPEs in which the stock cycles.

[15]Notice that, as defined, a Markov strategy ignores the length of a history. In some cases authors add the word 'stationary' to identify cases in which a Markov strategy ignores length, and thus allow for Markov strategies to depend on history length and the state. Also, it is common to define the set of payoff-relevant states exogenously, as a central feature of the environment. However, the notion of a state is not required for the analysis of dynamic games. See Maskin and Tirole (2001).

[16]This optimization problem has been widely studied in the literature. For those interested, proofs for this and the next results are given in Appendix B, but for a more detailed analysis see Sundaram (1989).

[17]This equilibrium involves players having infinite negative utility in periods 1 and after.

Let aggregate welfare be defined by adding the intertemporal utility of both players: $U^1 + U^2$. Efficient extraction maximizes aggregate welfare and is characterized next.

Proposition 2. *Extraction is efficient if each player consumes a proportion* $a_L = \frac{1-\delta}{2}$ *of the stock in each period.*

Efficient extraction is also a linear function of the stock, but with a different slope. Since $\delta \in (0,1)$, the efficient extraction rate is lower than what is prescribed in the LM-MPE ($a_L < a_M$). To see that permanent extraction at a_L cannot be supported as a MPE assume that one agent chose to always extract at the efficient rate and consider the problem of the other agent, who knows this. Agent 1's best response to any proportion chosen by agent 2 is given by: $a^1 = (1-\delta)(1-a^2)$.[18] Consequently, if $a^2 = (1-\delta)/2$, the best response is $a^1 = (1-\delta^2)/2$, so that agent 1 would choose to extract at a much higher rate. A free-riding problem creates the difference between the efficient solution and the LM-MPE, leading to 'the tragedy of the commons'.

From now the analysis focuses in the case when the action space is given by $\tilde{A} = \{a_L, a_M, a_H\}$ with $a_H = \frac{1}{2}$. The results in propositions 1 and 2 hold in the constrained action space.

Proposition 3. *If the action space is* $\tilde{A} = \{a_L, a_M, a_H\}$*, then 1) LM-MPE: The strategy profile* $\sigma(s) = a_M s$ *is a MPE; and 2) Extraction is efficient if each player consumes a proportion* a_L *of the stock in each period.*

Since $a_H < \frac{1+\delta}{2}$, if one player is selecting a_H, the other always can deviate to a_L and avoid infinite negative payoffs associated with depletion. This means that there is no 'corner' MPE with depletion in the first period.

Several history-dependent strategies are successful in supporting the efficient extraction level as the outcome of a subgame-perfect equilibrium. These strategies have two phases. The 'normal' phase is common to all strategies and prescribes efficient extraction in period 1 and in any period for which the outcome in the previous period was efficient: (a_L, a_L).[19] The 'punishment' phase distinguishes between strategies and describes choices when one of the players deviated from efficient extraction in

[18] See Appendix B.

[19] The pair (a^i, a^j) indicate the extraction rates selected by player i and player j in a particular period of the game.

the previous period. The list below describes the punishment phases for a set of history-dependent strategies that are frequently found in the literature:

- Grim LM-MPE: If the outcome of the previous period is not (a_L, a_L), select a_M from that period onwards.
- $T-$Period Punishment (TPP): If the outcome of the previous period is not (a_L, a_L), select a_M. If (a_M, a_M) has been the choice in the previous T periods, select a_L in $T+1$.
- Win Stay-Lose Shift: If $a^i_{t-1} \neq a^j_{t-1}$, select a_M in t. Otherwise select a_L.
- Tit-for-Tat: If the outcome of the previous period is not (a_L, a_L), select the partner's previous period choice.

The following proposition summarizes under which conditions these strategies can sustain efficient extraction.

Proposition 4. *The efficient outcome is supported as the normal phase of a subgame perfect equilibrium if the punishment phase follows: 1) Grim LM-MPE and $\delta > 0.476$, or 2) TPP with $T = 1$ and $\delta > 0.765$, or 3) TPP with $T = 8$ and $\delta > 0.476$, or 4) Win Stay-Lose Shift and $\delta > 0.765$. Tit for Tat can support efficient extraction as a Nash equilibrium if $\delta > 0.558$.*

These results do not depend on the extraction rate r. In other words, for any positive r, the punishment phases described in Proposition 4 would support efficient extraction in the stylized game.

Evolution of the stock

Since consumption is higher in the LM-MPE than in any equilibrium supporting efficient extraction, savings are lower and the growth rate of the stock is also lower. Starting at s_0 given, the following equations characterize the evolution of the stock respectively under the LM-MPE and any equilibrium supporting efficient extraction (EE).

$$s_{t+1}^{LM-MPE} = \left((1+r)\frac{\delta}{2-\delta}\right) s_t^{LM-MPE}$$

$$s_{t+1}^{EE} = ((1+r)\delta)\, s_t^{EE}$$

		Extraction rate	
		$r^{Low} = 0.57$	$r^{High} = 0.92$
Actions	a_H allowed	3 sessions (42 subjects)	3 sessions (38 subjects)
	a_H not allowed	3 sessions (40 subjects)	3 sessions (40 subjects)

Table 1: Treatments

In both cases, for given r and δ, the stock either grows to infinity, tends to zero or remains unchanged.[20]

3 Experimental Design

The experiment consists of a 2×2 factorial design in which one of the treatment variables is the reproduction rate, with $r^{High} > r^{Low}$ and the other is the number of strategies in $\widetilde{A}$ (see Table 1). When a_H is not allowed, treatments will have an action space given by $\{a_L, a_M\}$. Omitting a_H does not affect the results in Proposition 3. The discount rate in the lab will be set at $\delta = 0.75$, and thus, the values for a_M and a_L are fixed at 0.2 and 0.125 respectively. In what follows in this section I first describe the experimental game, which involves a few required adjustments to the stylized game, followed by other implementation issues.

Experimental game: adjustments to the stylized game

The experimental game involves adjustments to the utility function and the state space. With respect to payoffs, the instant utility function will still be logarithmic as long as the stock is positive, but compensation will be zero otherwise. The state space will be simplified. Let $\underline{s}$ solve $\ln(a_L \underline{s}) = 0$, that is, $\underline{s} = 1/a_L = 8 > 0$. If consumption decisions at t are such that $(1 + r)\left[(1 - a_t^1 - a_t^2)s_t\right] < \underline{s}$, then $s_{t+1} = 0$. Further, the state space is limited to positive integers so that: $\tilde{S} = \{0, 8, 9, ...\}$.

These adjustments achieve several goals. The modification in the utility function avoids the infinite negative payoff that would be otherwise associated with depletion. The positive lower bound $\underline{s}$ guarantees that there will be no negative payoffs, but more importantly, allows for a_H to be used as a credible punishment to support cooperation. To see why, recall first that with a continuous action space there is

[20] A steady state would arise if the transition function (1) was strictly concave.

a 'corner' MPE in which both players select to consume all the stock in the first period. Any deviation would involve a lower payoff in the current period and would not change future payoffs. The lower bound introduces similar incentives without the need to set a_H to one. Consider a case in which after one player selects a_H the stock that would result next period is below $\underline{s}$ and thus leads to depletion. In that case it is a best response for the other player to also select a_H. When the stock is high enough so that one player selecting a_H does not deplete the resource in the next period, there are profitable deviations. In other words, the lower bound will allow for MPEs in which players use threshold strategies that for some subset of the state space prescribe the selection of a_H.

Simplifying the state space to integers substantially reduces the number of possible history paths and thus further simplifies computations for subjects. More importantly, the modification will significantly aid to identify whether people use Markov strategies. Although there will always be the possibility for the stock to grow indefinitely, most of the data will indeed be contained within 10 possible states. This means that subjects will repeatedly make choices over the same states, but coming from different histories. In practical terms, the adjustment of the state space means that the outcome of (1) is rounded to the closest integer.

$$s_{t+1} = round\left((1+r)\left[(1-a_t^1-a_t^2)s_t\right]\right) \tag{3}$$

With these adjustments, the results in Proposition 3 will hold for selected values of r and a_H. First, there is a lower bound for the choice of a_H. Without the adjustments, paths with high extraction lead to low stock levels rapidly, which make it likely for the instant utility function to return negative values. This is part of what makes high extraction unattractive. If the instant utility function is not allowed to take on negative values, all strategies prescribing high extraction become relatively more attractive. Thus, to guarantee that deviation to a_H is not profitable in the LM-MPE, a_H needs to be relatively high. In particular, given other selected parameters it must be that $a_H > 0.43$, but for values close to that boundary the benefit of not deviating to a_H is small. For these reasons I set $a_H = 0.5$. The goal of introducing a_H is to test if behavior is affected when subjects are given the possibility of high extraction, it is not computationally demanding to verify that high extraction is not profitable and high extraction can be used to support cooperation.

Second, the fact that depletion of the stock is feasible constrains the particular selection of r. As a reference, consider a case with action space in the unit interval, with r and δ such that if players extract according to the LM-MPE the stock tends to zero. In each period players extract a fixed a_M proportion of the stock and mathematically it is never depleted. At any point in time there is an infinite future, so that no player would want to deviate to full depletion no matter how small the actual stock is. If this was implemented in the lab, even if small numbers are allowed for, at some point the stock will be sufficiently close to zero and thus effectively depleted. This naturally raises a problem, because if subjects expect for the stock to be depleted it is as if the logic of a finitely repeated game is introduced, thus changing the incentives with respect to the original game. In practice, restoring the proper incentives translates into a constraint in the choice of r. In other words, r^{Low} needs to be set so that along the LM-MPE path the stock always moves down from s_0 towards $\underline{s}$, but consumption at a_M levels never leads to depletion. This can be done by selecting r^{Low} so that once the stock reaches $\underline{s}$ if players choose to consume at the a_M level, rounding in (3) takes the next period's capital level back to $\underline{s}$. Therefore, with such choice of r^{Low} , the incentives are restored. Along the LM-MPE path there is never depletion and always an infinite future ahead. Next, I first formally define threshold strategies and then summarize the theoretical predictions with respect to the experimental game.[21]

Definition. A strategy profile σ is a Markov threshold strategy if $\sigma(s) = \begin{cases} \underline{a} \in \tilde{A} & \text{if } s < \overline{s} \\ \overline{a} \in \tilde{A} & \text{otherwise} \end{cases}$ for $\overline{a} \neq \underline{a}$ and some $\bar{s} \in S$.

Summary of Predictions: *Let the action space given by $\tilde{A} = \{a_L, a_M, a_H\}$, the state space in $\tilde{S} = \{0, 8, 9, ...\}$, transition to other states defined by (3) and payoffs be logarithmic if $s > 0$ and zero otherwise. Parametrization given by $r^{Low} = 0.57$, $r^{High} = 0.92$, $\delta = 0.75$, $a_H = \frac{1}{2}$ and $s_0 = 12$. Then,*

1. *LM-MPE: The strategy profile $\sigma(s) = a_M s$ is a MPE.*

2. *Extraction is efficient if both players consume a fraction a_L of the stock in each period.*

3. *x-Threshold-MPE: The threshold strategy profile $\sigma(s) = a_H$ if $s \leq x$ and a_M otherwise is a MPE for any $x \in \{8, ..., \tilde{s}\}$ with $\tilde{s} = 15$ for r^{Low} and $\tilde{s} = 13$ for*

[21]The definition of a strategy σ should be adjusted correspondingly to the action and state spaces of the experimental game.

r^{High}.

4. *The efficient outcome cannot be supported with pure Markov strategies.*

5. *The efficient outcome is supported as the normal phase of a subgame perfect equilibrium if the punishment phase follows: 1) Grim ML-MPE, 2) TPP with* $T \geq 2$, *3) Grim* $x-$*Threshold-MPE.*

Appendix B provides support for these claims. The first, second and last claim relate to propositions 3 and 4. The third claim explicitly identifies the threshold strategy equilibria generated by the adjustments, and the last claim also acknowledges that punishments with reversion to such threshold equilibria can be used to support efficiency. This means that the presence of a_H allows for a larger set of (harsher) punishment strategies to support cooperation.

To see the reason behind the fourth claim, first let S^E be the set of 'efficient' states, reached if both players always select a_L. Efficient extraction using Markov strategies must prescribe the choice of a_L for all states in S^E. It is straightforward to find $s \notin S^E$ such that for any $a \in \tilde{A}$ there is a profitable deviation.

Implementation in the laboratory

Three implementation issues are discussed next, related to easing the computational burden, the implementation of an infinitely repeated game in the laboratory, and the identification of strategies from observed choices. The changes previously discussed significantly reduce the required computations to make decisions in the game, but even in a three-action space and integer-valued state space calculations can be quite demanding. To eliminate computational problems, the interface allows subjects to explore hypothetical futures before they make a decision.

Figure 1(a) presents the first screen shot that subjects see in one of the treatments with a_H.[22] The screen is divided in two parts, with general information to the left and the decision-making panel to the right. At the top on the right side of the screen subjects are reminded of the current stock. The stock is referred to as 'points'. The game is presented in a 3×3 matrix with the subject's actions in the rows and their partner's in the columns. Each of the alternatives subjects have are described as

[22] A similar screen shot is used to explain the interface, which was written using z-Tree (Fischbacher 2007). See the instructions in Appendix A for further details.

taking some percentage of the current stock and, as described, those percentages are kept constant throughout the experiment. Each alternative presents subjects with the corresponding level of points they would take. For each of the 9 possible outcomes the interface presents the payoff for subjects, the payoff to their partner and the stock of points that would result for the next round. Information related to the subject is presented in blue and they are reminded that their payoffs (the top left number in each cell) are computed using the formula in the top left corner of the matrix.

To aid subjects at any time, before they make a decision, they can click on any of the 9 possible outcomes of the current round and compute a hypothetical future.[23] When they are computing hypothetical futures the background of the screen turns light red to remind them that they have yet to make a decision in the current round. They can keep on making hypothetical choices for future rounds. Figure 1(b) presents an example of someone who is considering a hypothetical future starting in round 1, has made hypothetical decisions for rounds 2 through 5, and is currently considering how things would look in round 6. The history for the chosen hypothetical future is summarized for subjects in the general panel so that they can trace the path they are considering. Subjects can keep track of payoffs for each hypothetical round and, more importantly, the cumulative payoffs for that path. At any point they can click on the 'Back' button if they want to return to the screen where they make a decision or they want to start a new hypothetical future. Subjects are provided with pen and paper in case they want to write down payoff comparisons across hypothetical futures.

Roth and Murnighan (1978) were the first to induce an infinitely repeated game in the lab with random termination. Under risk neutrality an infinitely repeated game with discount δ is equivalent to a game in which after each round the game finishes with probability $(1-\delta)$ and continues for one more round with probability δ. Although, in principle, random termination would allow to implement the game in the lab, in practice it will hinder learning. Any strategy that implements efficient extraction will return low payoffs at the beginning, as the stock is building, and larger payoffs in the future. Figure 2 shows cumulative payoffs for the r^{Low} treatment. The solid line tracks payoffs along the LM-MPE and adds each round's payoff to the amount already accumulated from previous rounds. The dotted line reports a similar computation in case players were to extract at the efficient level. As can be seen upon inspection, the cumulative payoff of efficient extraction is higher than the payoffs from

[23] Periods are referred to as 'rounds' and each repetition of the supergame is called a 'match'.

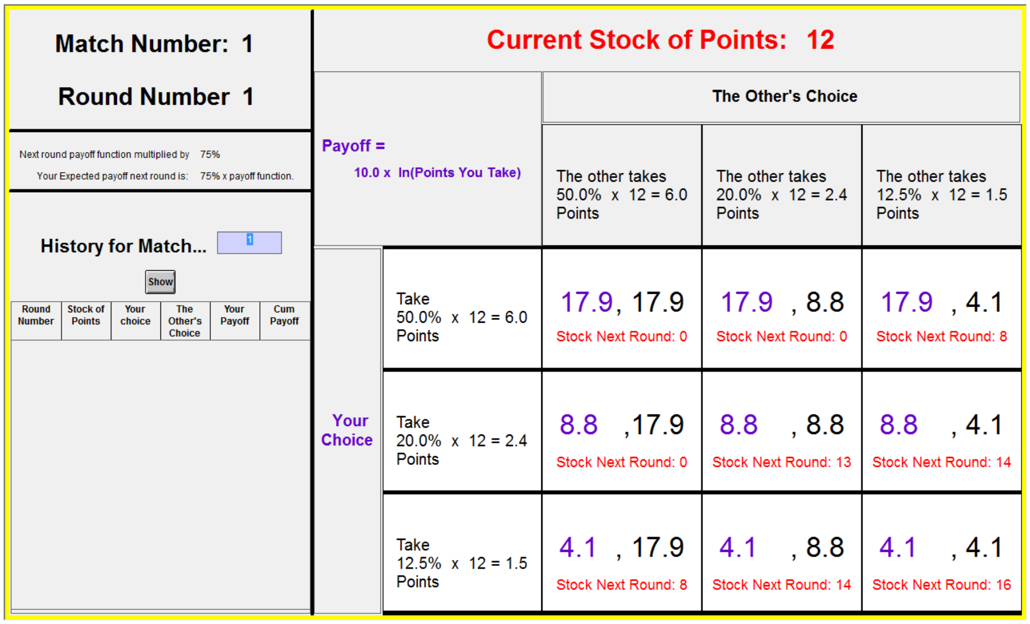

(a) Initial Screen Shot r^{High} treatment

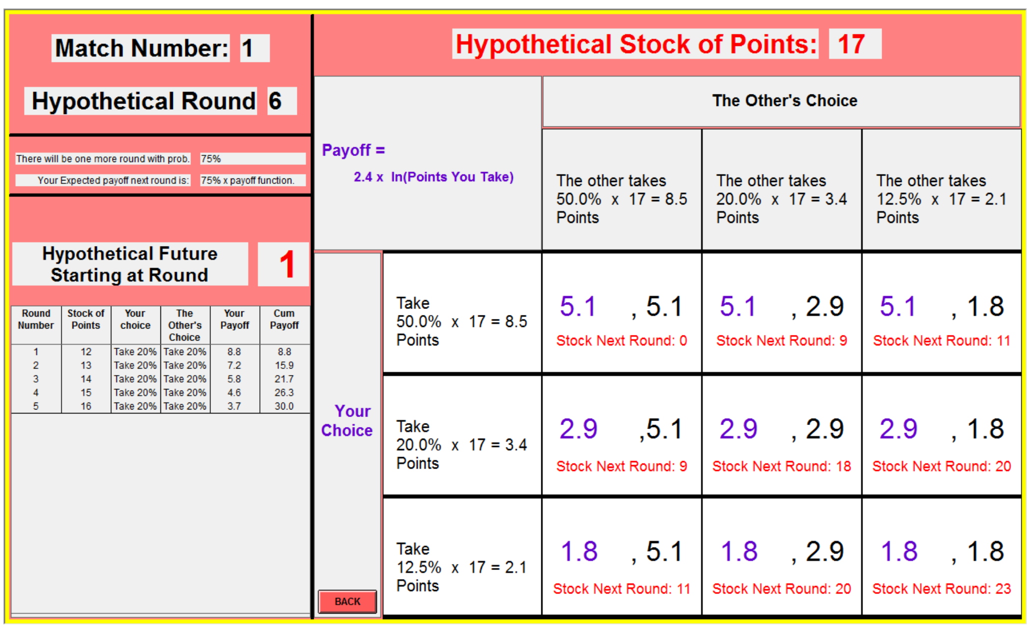

(b) Example of Hypothetical Futures

Figure 1: Interface Screen Shots

following the LM-MPE only if the game is played for more than 8 rounds. In other words, for subjects to experience repetitions in which efficient extraction pays off more than extraction at the MPE, repetitions should have at least 8 rounds. With random termination, even for high values for δ, subjects are likely not to experience repetitions that are long enough.

Instead, I implement the infinitely repeated game following Cabral et al. (2011). Subjects will play the first T rounds with certainty, discounting payoffs in each period by δ. Starting in round T there is no more payoff discounting in future rounds, but the game continues for one more round with probability δ. Under risk neutrality this implementation delivers payoffs equivalent to those of the infinitely repeated game. Moreover, Fréchette et al. (2011) compare across methods that implement infinitely repeated games in the laboratory and find that this is the method that produces more stable cooperation rates in a prisoners' dilemma. I set $T = 6$, so that on average a supergame lasts for 10 rounds.[24,25] The fact that each repetition is necessarily rather long constrains the total number of supergames that a session can have. Each session will consist of 9 repetitions of the supergame ('matches'), with subjects being randomly re-matched each time.

Finally, the information from observed choices may not be enough to identify between candidate strategies. Any strategy that successfully implements cooperation is indistinguishable from each other if there is no access to the choices that would have been made off the cooperating phase. Moreover, extraction at the LM-MPE level may be implemented with threats to a_H upon deviation. To gather extra information that would give access to some decisions off the observed path I implement a one-period

[24]With $\delta = 0.75$ the average extra expected rounds is 4.

[25]To explain random termination, subjects are told that at the end of each round after round 6 the computer will toss a 100 sided die and end the game if the number that comes up is higher than 75. For each round the random draw is displayed on the screen. Each treatment will have 3 sessions and the matlab random number generator was used to create three sequences of random numbers, one for each session. The random numbers are kept constant across treatments, so that for each sequence of random numbers there is a session of each treatment. The seed of the random number generator is set to the day and time of the first session in which a sequence was used. The instructions (see Appendix) provides an explanation with respect to the fact that random termination does not affect their expected payoffs. The middle box in the general information pane works as a reminder. Before random termination is implemented subjects are reminded that the payoff function will be multiplied by 75% next round (see Figure 1(a)). Starting at round 6 the interface reminds subjects that there will be one more round with probability 75% (see Figure 1(b)). In both cases, before and after random termination, subjects are reminded that their expected payoff for the next round are the same.

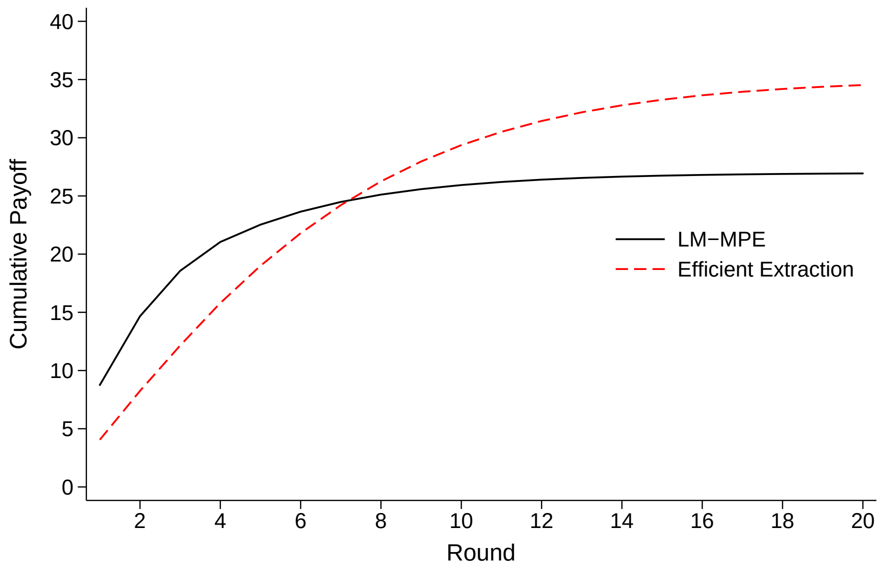

Figure 2: Cumulative payoffs in r^{Low} treatments

ahead strategy method. Sessions are divided in two parts, with part 1 consisting of the first 5 matches. The one-period ahead strategy method is introduced in the last 4 matches (part 2), once subjects have become familiarized with the interface. In part 2, subjects make a choice for round 1 and are subsequently asked to make a choice for round 2 for each possible round 1 decision of their partner . After they make their choices the software lets them know what their partner actually selected for round 1 and implements their corresponding decision in round 2. The procedure is repeated for all later rounds.

Summary

Table 2 summarizes the aspects of the experimental design discussed in this section. Three sessions for each treatment in Table 1 were run at CESS lab between the months of May and September 2011. Payments were set by adding payoffs from all 9 matches and then multiplying the total by \$0.065. In order to compare with experiments on the prisoners' dilemma, gains from cooperation can be measured with respect to the stage Nash equilibrium. In the treatments with a_H such gains are 96 and 230%

		Description
Random Termination		6 rounds with discounting $\delta = 0.75$. Round 7 onwards no discounting, match over with probability δ.
State space		$\tilde{S} = \{0, 8, 9, \ldots\}$
Parts	Part 1	5 matches
	Part 2	4 matches + one-period ahead strategy method
Extraction Predictions	LM-MPE	All treatments: Extract $a_M = \frac{1-\delta}{2-\delta} = 0.2$ of s_t
	Efficient extraction	All treatments: Extract $a_L = \frac{1-\delta}{2} = 0.125$ of s_t. Supported by strategies in Prop. 4.
Evolution of the stock	r^{Low} treatments	Extraction always at a_M, then $s \to \underline{s} = 8$. Extraction always at a_L then $s \to \infty$.
	r^{High} treatments	Extraction always at a_M, then $s \to \infty$. Extraction always at a_L then $s \to \infty$.

Table 2: Experimental Design

respectively for the r^{Low} and r^{High} treatments. When a_H is not a feasible choice, gains are 29 and 16% depending on whether the reproduction rate is low or high. In the treatments without a_H, permanent extraction at the stage Nash equilibrium is identical to the LM-MPE. It can be easily verified that percentage gains from cooperation computed with respect to the LM-MPE depend on r and δ. Maximizing those gains would involve a selection of 0.73 and 0.70 as discount rates for the r^{High} and r^{Low} treatments respectively. Thus, setting the discount rate at 0.75 comes very close to maximizing the percentage gains from cooperation with respect to the LM-MPE. There was no show-up fee, sessions lasted approximately 120 minutes, and the average subject collected $26. The average difference between the subject earning the maximum and the minimum in each session is 25.4%.

4 Empirical Analysis

4.1 Aggregate Data

Brief overview of general patterns

The evolution of the average stock gives a brief overview of the main aggregate patterns. The short and long-dashed lines in Figure 3 show, respectively, the predicted evolution of the stock according to the LM-MPE and efficient extraction. The solid lines display the average level of the stock for the first 10 rounds, averaged across

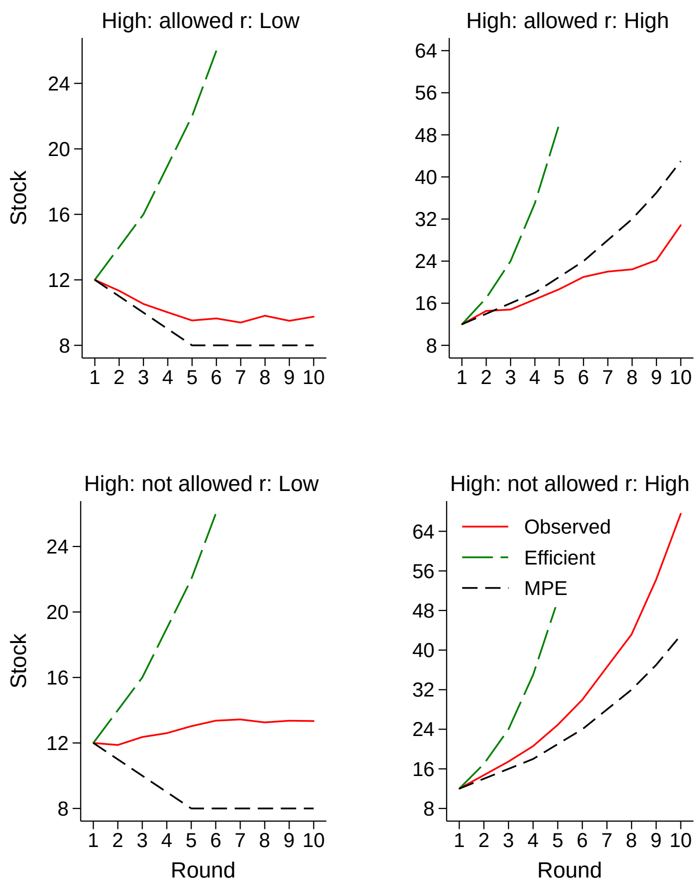

Figure 3: Evolution of the Stock (Last 4 Matches)

the last 4 matches.[26,27] First, consider the treatments with a_H, presented in the top row. For both levels of the reproduction rate, the trend observed in the data is quite close to the LM-MPE prediction. The median (not displayed in the Figure) perfectly coincides with the LM-MPE prediction for all rounds in the r^{Low} treatment. In the r^{High} treatment it also coincides until round 6 and is slightly lower for later rounds.

The evidence in Figure 3 points towards higher presence of efficient extraction

[26]Given random termination the number of observations over which the average is taken deceases with the number of rounds. The graphs display the first 10 rounds, for which there is a substantive amount of observations in all treatments.

[27]Recall that each repetition of the supergame will be referred to as a 'match' and that periods within each match are called 'rounds'.

when a_H is not available. For both treatments the observed average is above the LM-MPE prediction. Additionally, in treatments with a low reproduction rate, the observed stock average is closer to the efficient extraction prediction when compared to the high reproduction rate counterparts. These patterns suggest that efficient extraction rates are higher when a_H is not allowed and when the reproduction rate is low. I now turn to these questions.

Is cooperation higher when high extraction is not allowed?

In order to study the presence of cooperation I compute two measures: the cooperation rate and the cooperative outcome rate. The cooperation rate captures the proportion of times a_L is selected. Computations for all treatments are presented in columns 1 and 3 of Table 3. Some patterns are common across treatments.[28] First, cooperation rates are on average higher in the first round when compared against other rounds, but differences are not statistically significant. This holds for all treatments if the focus is on all matches (First v. Third column of Panel A) or on the last four matches (Panel B). Second, a_L is selected more often in the first part of the sessions than towards the end. This difference is statistically significant and holds for selections of a_L in the first round (First column: Panel A v. Panel B) or in all rounds (Third column: Panel A v. Panel B). Later in this section I will study the patterns within a match and within the session in more detail.

To answer the motivating question I focus on the comparison across treatments in the last four rounds (Panel B). Keep the reproduction rate constant and compare cooperation rates depending on whether a_H is allowed or not. Mean cooperation rates are statistically higher when a_H is not allowed, regardless if focus is on the first or on all rounds.

The cooperative outcome rate is defined as the proportion of times outcome (a_L, a_L) results and computations are presented in columns 2 and 4 of Table 3. Focusing in all rounds of the last four matches, using this measure also leads to a positive answer: higher levels of cooperative outcomes are observed when a_H is not allowed.

[28]Statistical statements are made with respect to Mann-Whitney tests at standard confidence levels.

		First Round		All Rounds	
		% a_L	% (a_L, a_L)	% a_L	% (a_L, a_L)
Panel A: All Matches					
a_H	r				
allowed	r^{Low}	12.70	1.59	11.93	4.32
	r^{High}	20.47	5.26	11.17	2.95
not allowed	r^{Low}	31.11	11.11	28.00	12.26
	r^{High}	24.44	5.56	19.42	5.99
Panel B: Last 4 Matches					
allowed	r^{Low}	8.74	1.19	7.90	1.88
	r^{High}	13.82	1.32	6.47	0.61
not allowed	r^{Low}	23.13	5.00	24.94	10.16
	r^{High}	18.75	1.25	13.70	4.05

Table 3: Cooperation rates across treatments (in %)

Is cooperation higher when the reproduction rate is low?

To evaluate the evidence, keep the allowed levels of opportunistic behavior constant and compare cooperation rates across levels of r. I focus on the last four matches (Panel B). When a_H is allowed, cooperation rates are not statistically different between treatments. When high extraction is not allowed, the differences are statistically significant for the first round and for all rounds. Similar findings hold for cooperative outcome rates.

Not allowing for a_H has a larger effect on cooperation rates than a lower reproduction rate. If a_H is not an option, the average cooperation rate triples when the reproduction rate is low and doubles when it is high (all rounds, last 4 matches).

A more detailed look into cooperation can be achieved by focusing on the behavior of those who select a_L when it has not been selected before in the match. A cooperative-intent case captures the event in which a_L is selected in round $t \geq 1$, if (a_L, a_L) was not the outcome of the previous round (or play is in round 1) and a_L is also selected in round $t + 1$. This event identifies situations in which a subject shows the intention to start a cooperative phase: there has not been a cooperative outcome immediately before, the subject chooses to cooperate in round t and keeps the choice for the next round, so that if their partner responds with a_L a cooperative phase can begin. Figure 4 shows the cooperative-intent patterns for all treatments, by adding all cooperative intent cases for each round.

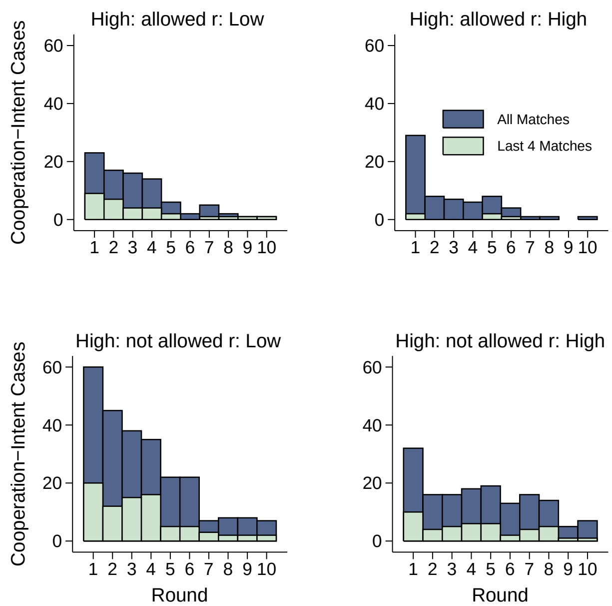

Figure 4: Distribution of Cooperation-Intent Cases within matches

Similar conclusions are reached when this information is used to answer the motivating question: more cooperative-intent cases are observed when the reproduction rate is low and there are fewer opportunities for opportunistic behavior. The figure also highlights the previously discussed common patterns across treatments, to which I turn next.

How does cooperation evolve within a session?

In all treatments, as subjects gain experience, they select a_L less often and a_M more often. In treatments with high extraction available, the selection of that choice is almost constant throughout the session. These findings can be verified upon inspection of Figure 5, which presents the evolution of the proportion of times each choice is selected for different treatments. The decrease in the selection of a_L (and corresponding increase in a_M) is starker in the first half of the sessions and stabilizes in the last matches.

An additional question is whether choices are affected by the introduction of the one-period ahead strategy method in match 6. Although Figure 5 does not show

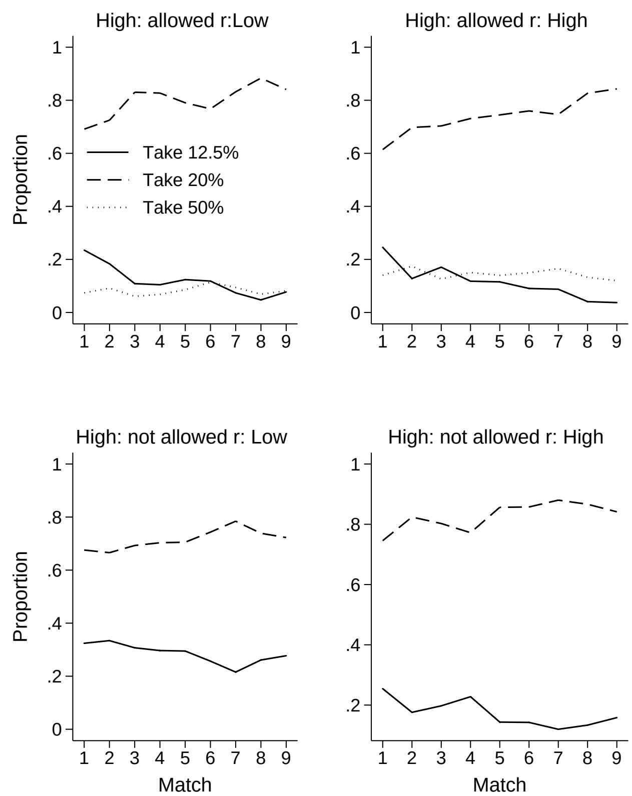

Figure 5: Evolution of choices within the session

clear evidence of any changes I look at this and the previous statements in more detail econometrically. I estimate a random effects probit model (order probit in the case when a_H is allowed) using the extraction choice at time t as the dependent variable. Dummy variables for different matches are included as part of the controls and the results are presented in Table 10 of Appendix C, where the estimation is discussed in more detail. For all treatments the coefficients do not show any significant difference between matches immediately before and after round 6. More importantly, the estimations show that there is no significant difference in choices after match 5, suggesting that most learning has already taken place. Moreover, results also confirm the patterns highlighted in Figure 5.

How does cooperation evolve within a match?

As a match evolves, subjects select a_L less frequently, in favor of higher extraction alternatives. The patterns for each treatment are presented in Figure 6, including all choices for the last four matches. When the reproduction rate is low, the selection of a_L increases between rounds 1 and 2. This pattern provides preliminary evidence for what Fudenberg et al. (2011) call 'exploitive' cooperative strategies. Subjects start by selecting a_M in round 1 and implement a history dependent cooperative strategy starting in round 2.

The selection of higher extraction rates in later rounds can result from strategic reasons (i.e. punishments upon deviation from cooperation), or may be induced by random termination. The econometric analysis in Table 10 includes dummies variables for different rounds. Indeed, in treatments when a_H is allowed there is evidence of a significant increase in higher extraction in round 6, perhaps with the intention of depleting the stock as soon as possible. Later in this section I will use an econometric procedure to uncover used strategies and I will specifically include arrangements that capture the behavior of subjects who are responding to the environment. Indeed, the analysis will show that there is a small proportion of subjects who respond to random termination. Before turning to individual behavior I look at choices that lead to depletion in more detail.

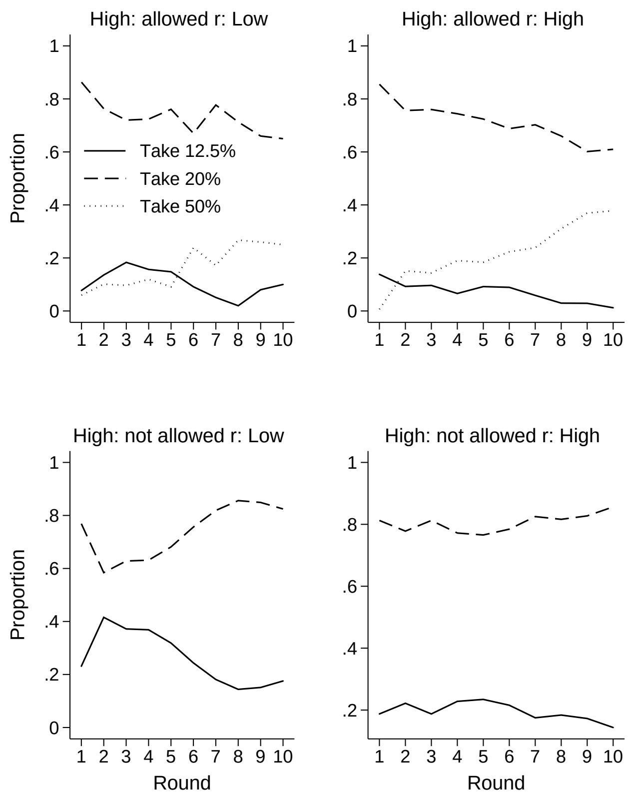

Figure 6: Distribution of Actions across Rounds (Last 4 Matches)

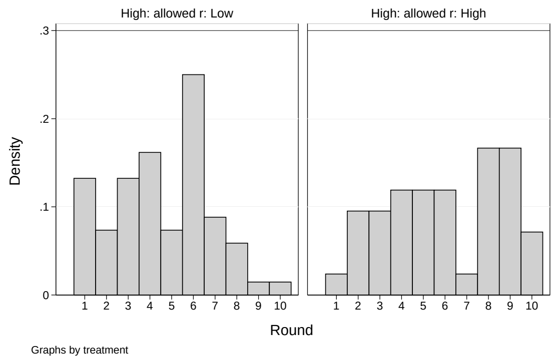

Figure 7: Depletion of the stock: Distribution of match length

Does a low reproduction rate trigger higher levels of depletion?

Only in the treatments with a_H can the stock be depleted.[29] With a total of 42 subjects and 9 matches per treatment there are 189 partnerships to study when the reproduction rate is low. Of those partnerships, 74% end up with a depleted stock before the match randomly terminates. When the stock is depleted, 92% of the times it is the result of a unilateral decision. Figure 7 tracks the distribution of match lengths when the stock is depleted.[30] In the r^{Low} treatment, the median length of a match is 5 rounds and although the mode is at 6 rounds, in most cases depletion takes place before random termination.[31] When the reproduction rate is high, 54.4% of all partnerships (171 in total) end with a depleted stock and in a third of those cases both subjects decided to select a_H in that round. As can be seen upon inspection of the graph to the right of Figure 7 depletion occurs mostly within rounds 2 and 9. The distribution almost remains unchanged towards the end of the session.

To better understand the differences between treatments Table 4 shows the results of a counter factual exercise. All extractions decisions are kept constant, but the reproduction rates between treatments are switched. In other words, I use extraction decisions as submitted by subjects, but compute the level of the stock using the reproduction rate of the other treatment. Observing depletion levels that do not

[29] When a_H is not available, extraction at a_M implies a lower value of the stock permanently at 8.

[30] Although matches last for as many as 17 rounds only the distribution for the first 10 rounds is displayed, as all matches that lasted for more than 10 rounds did not end up with the stock depleted.

[31] The distribution slightly moves to the left if the last 4 matches are considered. The median length of a match in part 2 is 4.5 rounds.

Treatment	Stock Depleted?		Counter factual in r: Stock Depleted?	
r	No	Yes	No	Yes
r^{Low}	49	140	162	27
	26.0%	74.0%	85.7%	14.3%
r^{High}	78	93	24	147
	45.6%	54.4%	14.0%	86.0%

Table 4: Depletion of the stock: Distribution of Partnerships across treatments

change in the r^{Low} treatment when the evolution of the stock is computed using r^{High} would suggest no difference in behavior across treatments. Results are reported in Table 4. The last two columns show the result of the counter factual exercise. If r^{High} was used to compute the evolution of the stock, 85.7% of partnerships in the r^{Low} treatment would not end with depletion. Moreover, if r^{Low} is used to compute growth in the r^{High} treatment, 86% of the partnerships would see their stocks depleted.

The findings suggest in favor of an important difference in behavior across treatments. Because the environment is more prone to deplete the stock such events are observed more frequently when the reproduction rate is low. However, subjects make significantly greater efforts to avoid depletion when the reproduction rate is low.

4.2 Identifying Strategies

For each match and each subject there is a specific history of decisions, totaling 1440 such histories in all sessions. This section looks for systematic patterns in those histories, as the ones suggested by well known strategies. On average, each history vector involves more than 10 choices, so that it is very likely to find histories that do not perfectly fit into a particular strategy. A taxonomy with perfectly fitted strategies is presented in Appendix C. In this section, I focus on estimating the presence of a large set of strategies allowing for errors and use the procedure in Dal Bó and Fréchette (2011). Consider a history vector for one subject in one match. The first step of the procedure consists of computing a vector of the choices that would be prescribed to that subject by each conceivable strategy. The econometric procedure acts as a signal detection method and estimates via maximum likelihood how close the actual choices are from the prescriptions of each strategy. The key estimates obtained are the proportion in which each strategy is observed in the population sample.

Which strategies are used?

Define c_{imr} as the choice of subject i in round r of match m, $c_{imr} \in \{a_L, a_M, a_H\}$. Consider a set of K strategies that specify what to do in round 1 and in later rounds for given levels of the stock and past history. Thus, for each history h, the decision prescribed by strategy k for subject i in round r of match m can be computed: $h_{imr}(h^k)$. A choice is a perfect fit for a history if $c_{imr} = h_{imr}(h^k)$ for all rounds of the history. The procedure allows for mistakes and models the probability that the choice corresponds to a strategy k as:

$$Pr(c_{imr} = h_{imr}(h^k)) = \frac{1}{1 + exp\left(\frac{-1}{\gamma}\right)} = \beta. \tag{4}$$

In (4) $\gamma > 0$ is a parameter to be estimated. As $\gamma \to 0$, then $Pr(c_{imr} = h_{imr}(h^k)) \to 1$ and the fit is perfect. Define y_{imr} as a dummy variable that takes value one if the subject's choice matches the decision prescribed by the strategy, $y_{imr} = 1\left\{c_{imr} = h_{imr}(h^k)\right\}$. If (4) specifies the probability that a choice in a specific round corresponds to strategy k, then the likelihood of observing strategy k for subject i is given by:

$$p_i\left(s^k\right) = \prod_m \prod_r \left(\frac{1}{1 + exp\left(\frac{-1}{\gamma}\right)}\right)^{y_{imr}} \left(\frac{1}{1 + exp\left(\frac{1}{\gamma}\right)}\right)^{1-y_{imr}} \tag{5}$$

Aggregating over subjects: $\sum_i \ln\left(\sum_k \phi_k p_i\left(s^k\right)\right)$, where ϕ_k represents the parameter of interest, the proportion of the data which is attributed to strategy s^k. I will compute the standard deviations for the estimates bootstrapping 1000 repetitions.[32] The procedure recovers an estimate for γ and the corresponding value of β can be calculated using (4). The estimate of β can be used to interpret how noisy the estimation is. For example, with only two actions a random draw would be consistent with $\beta = 0.5$.

The estimation results depend on the set of strategies that are considered. All the strategies that are used in the analysis are described in Table 5. There are four groups of strategies, presented in the following order: Markov, Cooperating, Mixed, and Round dependent. Markov strategies include those that always prescribe the same

[32] The procedure leaves unidentified the standard error for the K-th strategy.

choice and threshold strategies. Cooperating strategies involve history dependent strategies that aim at sustaining cooperation. Some strategies combine behavior that depends on the current level of the stock and past decisions. Such strategies will be referred to as 'Mixed'. Finally, 'Round' dependent strategies capture behavior that conditions on particular rounds within the supergame.

With the available information, distinguishing between strategies that often coincide is difficult. Compare, for example, a threshold strategy prescribing the choice of a_H if the state is at or below 12 and a_M otherwise, and always selecting a_H. Because the starting level for the stock is at 12, a subject following either of those strategies will not experience states above 12. Identification, however, is indeed possible for most strategies in Table 5, but strategies that are not very different from each other will show higher standard deviation estimates. For instance, 'Exploitive Grim LM-MPE' captures the notion that the subject starts by defecting from cooperation, but shifts to a cooperative phase in the next round. Clearly, this strategy has many observations in common with Grim LM-MPE and this will translate into imprecise estimates. In such cases, I will discuss the joint significance of those strategies.

Most strategies in Table 5 have been previously considered in the literature, except for those that combine state and past history. A strategy that results from inspection of the data is 'Tat for Tit and a_M'. This strategy compensates if the stock is below a certain threshold: higher extraction in the present if the partner over extracted in the previous period and vice versa. Once the stock is above the threshold, extraction is at a_M. In principle, this strategy can be estimated for different thresholds. To avoid identification problems, I will use a common threshold for all subjects. To determine the common threshold, I first compute the level of the threshold that is most successful in explaining choices for each history according to this strategy. I use as common thresholds the median values of the optimal thresholds (11 and 14 for r^{Low} and r^{High} respectively).

Table 6 reports the estimates for strategies with positive coefficients. The strategy 'Always a_M', prescribed by the LM-MPE, is in all cases the most popular one. The proportions that correspond to this strategy are higher when the reproduction rate is high. In the treatment with a_H and high reproduction rate there is evidence of Markov threshold behavior, at 13%, although the coefficient is not statistically significant.

With the exception of 'Always a_H', which shows a small and significant estimate, all other strategies in the treatments with a_H are not significant. In the r^{High} treat-

Strategy	Abbreviation	Description
Always Cooperate	Always a_L	Always play a_L
Levhari-Mirman MPE	Always a_M	Always play a_M
Permanent High Extraction	Always a_H	Always play a_H
Threshold (a_H, a_M)	Threshold $a_H \leq \tilde{s}$	a_H if $s \leq \tilde{s}$, a_M otherwise. For $\tilde{s} \in \{8, .., 16\}$
Threshold (a_M, a_H)	Threshold $a_M \leq \tilde{s}$	a_M if $s \leq \tilde{s}$, a_H otherwise. For $\tilde{s} \in \{8, .., 16\}$
Threshold (a_L, a_M)	Threshold $a_L \leq \tilde{s}$	a_L if $s \leq \tilde{s}$, a_M otherwise. For $\tilde{s} \in \{8, .., 16\}$
Grim LM-MPE	Grim LM-MPE	Play a_L until either deviates. Then always play a_M
Exploitive Grim LM-MPE	Exploitive Grim LM-MPE	Play a_M in round 1. Play Grim LM-MPE starting in round 2
Grim a_H	Grim a_H	Play a_L until either deviates. Then always play a_H
Exploitive Grim a_H	Exploitive Grim a_H	Play a_M in round 1. Play Grim a_H starting in round 2
Grim Threshold (a_H, a_M)	Grim Threshold $a_H \leq \tilde{s}$	Play a_L until either deviates. Then play Threshold (a_H,a_M)
Exploitive Grim Threshold (a_H, a_M)	Expl Grim Threshold $a_H \leq \tilde{s}$	Play a_M in round 1. Play Grim Threshold (a_H,a_M)
Tit for Tat	Tit for Tat	Play a_L in round 1. From round 2 what partner's previous choice
Exploitive Tit for Tat	Exploitive Tit for Tat	Play a_M in round 1. Then play Tit for Tat
Lenient Grim 2	Grim2	Play a_L until 2 subsequent rounds occur in which either player chose differently. Then play a_M forever
Lenient Grim 3	Grim2	Play a_L until 3 subsequent rounds occur in which either player chose differently. Then play a_M forever
Win Stay-Lose Shift	WSLS	Play a_L if both choice a_L last round. Otherwise select a_M
T2PP	T2PP	Play a_L until either deviates, then play a_M until the outcome (a_M, a_M) repeats twice and return to a_L
T3PP	T2PP	Play a_L until either deviates, then play a_M until the outcome (a_M, a_M) repeats three times and return to a_L
Exploitive T2PP	Exploitive T2PP	Play a_M in round 1. Then play T2PP
Tit for 2 Tats	TF2T	Play a_L unless partner deviated in last 2 rounds. In that case select partner's previous choice
Grim (a_H,a_M)	Grim a_H,a_M	Play a_M until either deviates. Then always play a_H
Tat for Tit and a_M	Tat for Tit and a_M	a_M in round 1. If $s \leq \tilde{s}$: 1) partner chose a_L last, choose a_M 2) partner did not choose a_L, play a_M. Otherwise a_M
Tat for Tit and a_M 2	Tat for Tit and a_M 2	a_M in round 1. If $s \leq \tilde{s}$: 1) partner chose a_L last, choose a_H 2) partner did not choose a_L, play a_L. Otherwise a_M
Stock adjustment	Stock adjustment	a_M in round 1. If stock decreased, play a_L. Otherwise a_M
Stock adjustment 2	Stock adjustment 2	a_M in round 1. If stock decreased, play a_M. Otherwise a_L
False cooperator	FC	Plays a_L in round 1, then a_M forever
a_M and a_H if Round $\geq x$	a_M and a_H if R $\geq x$	a_M if Round $\leq x$. Otherwise play a_H. For $x \in \{5, ...10\}$
a_L and a_M if Round $\geq x$	a_L and a_M if R $\geq x$	a_L if Round $\leq x$. Otherwise play a_M. For $x \in \{5, ...10\}$

Table 5: Description of strategies considered. First group: Markov, Second: Cooperating, Third: Mixed, Fourth: Round dependent

ment, estimates for history dependent strategies are not only not significant, but are also quite small. Adding cooperating strategies accounts for slightly more than 8%. Contrarily, there are higher estimates for cooperation strategies in the r^{Low} treatment (although still not significant), with exploitive tit for tat accounting for around 11%. Finally, approximately 7% of strategies can be explained by subjects who turn towards a_H starting in round 6.

In the treatments without a_H, as expected, there is presence of strategies that sustain cooperation. Most notably, Grim LM-MPE is significant at the 1% level and accounts for approximately 16% in the r^{Low} treatment. If both, Grim LM-MPE and the exploitive version are taken together, they are jointly significant and explain approximately 30%. When attention centers on the r^{High} treatment, tit-for-tat turns out as the most popular cooperative strategy at almost 10%. Jointly significant when combined with the 'exploitive' version, they account for approximately 20%.

Do used strategies change during the session?

According to Figure 5 there is evidence of higher cooperation rates in the first few matches. To study this in more detail, Table 7 provides the estimation for strategies using information from matches 1-2. Identification is now more difficult as one-period ahead information is not available. I, thus, reduce the set of estimated strategies and combine the original and 'exploitive' version of several history-dependent strategies into one. Prior to the estimation, I determine for each observed history whether the original or the 'exploitive' version better captures their behavior. I use this information to construct a unique prescribed choice for each history-dependent strategy.[33]

When attention centers in the first matches, Grim LM-MPE is statistically significant in all treatments. In some cases (i.e. a_H not allowed, r^{Low}) it is even the most popular strategy. Notice that, although not statistically significant, tit-for-tat accounts for approximately 15% of strategies in both treatments when a_H is allowed. In other words, there is evidence of history-dependent strategies present at significant levels even if a_H is allowed. The strategy 'Always a_M' receives an important share in all treatments, but is in all cases below the values reported in Table 6. To allow for a comparison with later matches, Table 7 also presents the estimation results from using matches 6-9, but excluding information from the one-period ahead strategy

[33]Table 12 in Appendix C shows that the conclusions in Table 6 are not changed with this modification to the definition of history-dependent strategies.

Strategies	a_H allowed		a_H not allowed	
	r^{Low}	r^{High}	r^{Low}	r^{High}
Always a_H	0.021***	0.000		
	(0.007)	(0.000		
Always a_M	0.647***	0.730***	0.554***	0.722***
	(0.132)	(0.108)	(0.112)	(0.070)
Always a_L	0.026	0.000	0.010	0.025
	(0.051)	(0.011)	(0.020)	(0.027)
Threshold $a_H < 10$	0.021	0.131		
	(0.030)	(0.095)		
Grim LM-MPE	0.022	0.000	0.163***	0.002
	(0.026)	(0.014)	(0.080)	(0.000)
Exploitive Grim LM-MPE	0.000	0.000	0.125	0.048
	(0.000)	(0.000)	(0.100)	(0.040)
Grim a_H	0.043	0.050		
	(0.033)	(0.043)		
Exploitive Grim a_H	0.000	0.000		
	(0.019)	(0.007)		
Grim (a_M,a_H)	0.034	0.031		
	(0.037)	(0.031)		
Tit for Tat	0.000	0.000	0.051	0.099*
	(0.000)	(0.004)	(0.054)	(0.060)
Exploitive Tit for Tat	0.114	0.000	0.034	0.096
	(0.073)	(0.001)	(0.035)	(0.066)
Tat for Tit and a_M	0.000	0.000	0.060	0.008
	(0.000)	(0.000)		
a_M and a_H if $R \geq 6$	0.065	0.056		
γ	0.692	0.940	0.758	0.534
β	0.809	0.743	0.789	0.867

Bootstrapped standard errors in parentheses. Data from Matches 6-9 including unimplemented choices from one-step ahead strategy method

Table 6: Estimation of Used strategies by treatment

method in the estimation.[34] As the session evolves, the presence of history-dependent strategies that support cooperation vanishes in the treatments with a_H and, although still significant, it drops when a_H is not allowed.

Can behavior be explained only using Markovian strategies?

One way to answer this question is by constraining the estimation to include only linear Markov strategies and tracking the value of β. If the values of β that result are close to a random draw, then behavior is better explained by also including other strategies. The results are presented in Table 8. In treatments without a_H, a random draw is at 0.5 and estimates of β are above 0.75. In cases with high extraction allowed, random draws are at 1/3 and estimates for β are also relatively high, suggesting that linear Markovian strategies are a good predictor of behavior.

Estimations in Appendix C suggest that the introduction of the one-period ahead strategy method did not alter selected strategies. The analysis compares the strategy estimation output using only information from matches 4 and 5 to the results when only information coming from matches 6 and 7 is used. Qualitative conclusions are similar.

5 Conclusion

This paper implements the dynamic commons problem in the laboratory with the aim of understanding how the presence of an endogenously evolving state variable may affect the incentives to cooperate. Findings show that cooperation is higher when there are limits to opportunistic behavior and when the reproduction rate of the stock is low. If high levels of opportunistic behavior are feasible, cooperation almost disappears, but the most frequent extraction levels are not the highest. Most behavior can be rationalized with linear Markov strategies, as suggested by prominent MPE characterized in the literature. When cooperation arises, different arrangements are used to sustain it, depending on the reproduction rate. When the reproduction

[34]In the comparison between the results in Table 7 and Table 6 it stands out that when only actual choices are used there is almost no evidence of tit-for-tat in treatments without a_H. Using data only from the observed path leads to significant estimates only for 'Always at a_M' or Grim LM-MPE (once at a_L, then at a_M). The knowledge that comes from off path decisions, however, shows that some subjects would be willing to select a_L if their partner selected a_L. This leads to some mass of histories to be better explained by tit-for-tat than by either Grim LM-MPE or 'Always a_M'.

Strategies	a_H allowed				a_H not allowed			
	r^{Low}		r^{High}		r^{Low}		r^{High}	
	Matches 1-2	Matches 6-9	Matches 1-2	Matches 6-9	Matches 1-2	Matches 6-9	Matches 1-2	Matches 6-9
Always a_H	0.000	0.027*	0.000	0.000				
	(0.002)	(0.013)	(0.007)	(0.000)				
Always a_M	0.558***	0.621**	0.265*	0.598***	0.191*	0.396***	0.612***	0.723***
	(0.127)	(0.138)	(0.154)	(0.087)	(0.114)	(0.100)	(0.101)	(0.109)
Always a_L	0.038	0.045	0.000	0.000	0.035	0.000	0.075	0.025
	(0.043)	(0.047)	(0.022)	(0.000)	(0.048)	(0.003)	(0.055)	(0.029)
Threshold $a_H < 10$	0.000	0.000	0.210***	0.252*				
	(0.000)	(0.000)	(0.023)	(0.138)				
Grim LM-MPE	0.228**	0.023	0.399***	0.063	0.772***	0.419***	0.311***	0.210***
	(0.120)	(0.036)	(0.175)	(0.083)	(0.125)	(0.101)	(0.105)	(0.077)
Grim a_H	0.000	0.023	0.057	0.017				
	(0.004)	(0.019)	(0.048)	(0.000)				
Grim (a_M,a_H)	0.011	0.000	0.000	0.000				
	(0.019)	(0.000)	(0.016)	(0.028)				
Tit for Tat	0.149	0.198	0.071	0.000	0.000	0.109	0.000	0.041
	(0.126)	(0.113)	(0.072)	(0.020)	(0.036)	(0.074)	(0.100)	(0.054)
Tat for Tit and a_M	0.000	0.000	0.000	0.000	0.000	0.073	0.000	0.000
	(0.020)	(0.000)	(0.002)	(0.000)				
a_M and a_H if $R \geq 6$	0.013	0.060	0.000	0.068				
γ	0.647	0.526	0.909	0.788	0.798	0.675	0.567	0.434
β	0.824	0.870	0.750	0.781	0.777	0.814	0.820	0.909

Bootstrapped standard errors in parentheses. Data includes only actual choices for matches 6-9.

Table 7: Estimation of Strategies Used in Matches 1-2 and 6-9

Strategies	a_H allowed		a_H not allowed	
	r^{Low}	r^{High}	r^{Low}	r^{High}
Always a_H	0.018***	0.0000		
	(0.006)	(0.009)		
Always a_M	0.910***	0.974***	0.974***	0.900***
	(0.052)	(0.036)	(0.030)	(0.047)
Always a_L	0.070	0.025	0.025	0.099
γ	0.806	1.044	0.906	0.619
β	0.775	0.723	0.750	0.834

Bootstrapped standard errors in parentheses. Data from Matches 6-9 including unimplemented choices from one-step ahead strategy method

Table 8: Estimation of Used strategies by treatment

rate is low, agreements involve large punishment periods, while subjects are promptly ready to forgive their defecting partners if the reproduction rate is high.

The findings in this paper raise new questions. In the developed parametrization, the MPE typically characterized in the literature involves the same choice for all states. A clear benefit of this construction is that it allows for simple cooperative equilibria (i.e. cognitively as demanding as in the prisoners' dilemma). However, the parametrization is a limited test of Markov behavior. Although there is some evidence of subjects using threshold strategies, a more convincing test would involve MPEs that require agents to change choices depending on the state. The environment presented in this paper can be readily modified for such a test. In fact, preliminary evidence from treatments with that property suggest that the choices of a considerable number of subjects is consistent with Markov behavior.

Attention to Markov strategies is often justified as a means for agents to deal with complex strategy spaces. In fact, one frequent argument in favor of Markov strategies holds that since the state determines the stage game agents face, it summarizes all the information from past decisions that is relevant to determine payoffs at time t and the details of how that particular state was reached can be ignored. The fact that there is almost no cooperation when agents have three choices, but positive levels when two choices are available is not inconsistent with subjects dealing with a more complicated action space by ignoring history. This environment can also be modified to test how subjects deal with increased complexity.

Appendix A: Instructions

This appendix contains the instructions for the r^{High} treatment when a_H is allowed.

Welcome

You are about to participate in a session on decision-making, and you will be paid for your participation with cash vouchers, privately at the end of the session. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. Please turn off pagers, mp3 players and cellular phones now. Please close any program you may have open on the computer. The entire session will take place through computer terminals, and all interaction between you will take place through the computers. Please do not talk or in any way try to communicate with other participants during the session. We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the session and will be shown how to use the computers. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear. This experiment has two parts; we will start with part 1. Once this part is over, instructions for part 2 will be given to you. Your decisions in this part have no influence on part 2.

General Instructions: Part 1

1. In this experiment you will be asked to make decisions in several rounds. You will be randomly paired with another person for a sequence of rounds. Each sequence of rounds is referred to as a match. Once a match ends, you will be randomly paired with another person for a new match.

2. At the beginning of a match you and the person you are matched with start with a stock of 12 points. Your choices are:

   - Take 50% of the points (50% x 12 = 6.0 points)
   - Take 20% of the points (20% x 12 = 2.4 points)
   - Take 12.5% of the points (12.5% x 12 = 1.5 points)

   Your payoff for the round is 10 times the natural logarithm of the points you decide to take:

   - If you take 50% of the points, your payoff would be: 10 x ln(6.0) = 17.9

- If you take 20% of the points, your payoff would be: 10 x ln(2.4) = 8.8
- If you take 12.5% of the points , your payoff would be: 10 x ln(1.5) = 4.1

The software will make this calculations for you and will always round up the points you take and your payoff to one decimal point. The choices and the payoffs in Round 1 are as follows:

Your Choice	the other's choice		
	Take 50% of the points	Take 20% of the points	Take 12.5% of the points
Take 50% of the points	17.9,17.9	17.9,8.8	17.9,4.1
Take 20% of the points	8.8,17.9	8.8,8.8	8.8,4.1
Take 12.5% of the points	4.1,17.9	4.1,8.8	4.1,4.1

As you can see, this shows the payoff associated with each choice. The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are matched with. That is, if: You select "Take 50%" and:

- The other selects "Take 50%", you each make 17.9.
- The other selects "Take 20%", you make 17.9, the other 8.8.
- The other selects "Take 12.5%", you make 17.9, the other 4.1.

You select "Take 20%" and:

- The other selects "Take 50%", you make 8.8, the other 17.9.
- The other selects "Take 20%", you each make 8.8.
- The other selects "Take 12.5%", you make 8.8, the other 4.1.

You select "Take 12.5%" and:

- The other selects "Take 50%", you make 4.1, the other 17.9.
- The other selects "Take 20%", you make 4.1, the other 8.8.
- The other selects "Take 12.5%", you each make 4.1.

3. The stock of points in the Next Round depends on the choices you and the other make This Round. The remaining points after you and the other made your choices is:

   Remaining Points = (1-% you take - % the other takes) ×Current Stock of Points

   The Stock of Points next round increases the Remaining Points by 92%:

   $$\text{Stock of Points Next Round}=1.92\times\text{Remaining Points}$$

   The software will make the computation for you and always round the points to the closest whole number. If the stock of points next round is lower than 8 points the software considers that the stock will be depleted and will display 'Stock Next Round: 0'. The Stock of Points in Round 1 is 12. Depending on the choices you and the other make in Round 1 the table below shows the resulting Stock in Round 2.

Your Choice	the other's choice		
	Take 50% of the points	Take 20% of the points	Take 12.5% of the points
Take 50% of the points	Stock Next Round:0	Stock Next Round:0	Stock Next Round:8
Take 20% of the points	Stock Next Round:0	Stock Next Round:13	Stock Next Round:14
Take 12.5% of the points	Stock Next Round:8	Stock Next Round:14	Stock Next Round:16

   That is, if: You select "Take 50%" and:

   - The other selects "Take 50%", the Stock Next Round is 0 Points.
   - The other selects "Take 20%", the Stock Next Round is 0 Points.
   - The other selects "Take 12.5%", the Stock Next Round is 0 Points.

   You select "Take 20%" and:

   - The other selects "Take 50%", the Stock Next Round is 0 Points.
   - The other selects "Take 20%", the Stock Next Round is 11 Points.
   - The other selects "Take 12.5%", the Stock Next Round is 13 Points.

You select "Take 12.5%" and:

- The other selects "Take 50%", the Stock Next Round is 0 Points.
- The other selects "Take 20%", the Stock Next Round is 13 Points.
- The other selects "Take 12.5%", the Stock Next Round is 14 Points.

4. On your computer screens the information with respect to payoffs and the stock next Round will be displayed in a single table, as you can see on the screenshot next page and in the projected slide as well. On the left of your screens you will have general information: the Match number, the Round number and history on your past decisions.[35]

   - On the right side you have the information on payoffs and future stock of points.
   - At the beginning of each Match (Round 1) the Stock of Points is 12. You can see the Current Stock of Points at the top of the screen in Red.
   - Depending on the decisions for this round, each cell presents in Red the Stock of points Next Round.
   - In Blue you can see information related to your payoffs. At the top left corner of the matrix you are reminded that your payoff is equal to 10 times the natural logarithm of the points you decide to take. For each cell, the computation is presented to you in blue.
   - When you click on top of one of your three choices a submit button will appear. You can change your decision by clicking on any other choice, but once you click on the submit button your decision will be final.
   - Once you and the person you are paired with have made your choices, those choices will be highlighted and your payoff for the round will appear.

5. Once Round 1 is over play moves to Round 2. In Round 2, you and the other will face exactly the same choices (Take 50%, 20%, or 12.5% of the points). However, the current stock of points from which you and the other will be taking points will be the one that resulted from the choices in the previous

[35]Subjects see a projected slide of Figure 1(a), that is also printed in the next page of the instructions, but omitted on this appendix.

round. Your payoff in round 2 will be: 75% of 10 times the natural logarithm of the points you decide to take, that is 7.5 times the natural logarithm of the points you take.

6. This process repeats itself. Each time you move to a new Round, the current stock of points will be the one that resulted from your choices in the previous round and the payoff function is multiplied by 75%. The evolution of the payoff function is specified in the following table:

Round	Payoff	
2	75% × 10 × ln(Points you take) =	7.5 × ln(Points you take)
3	75% × 7.5 × ln(Points you take) =	5.6 × ln(Points you take)
4	75% × 5.6 × ln(Points you take) =	4.2 × ln(Points you take)
5	75% × 4.2 × ln(Points you take) =	3.2 × ln(Points you take)
6	75% × 3.2 × ln(Points you take) =	2.4 × ln(Points you take)

The interface will clearly present how your payoffs are computed in each round. The corresponding formula will be displayed at the top left corner of the matrix.

7. After Round 6 , there is a 75% probability that the match will continue for at least another round. At the end of each round after Round 6 the interface will simulate a 100-sided die and will end the match if a number higher than 75 comes up. If a number equal or lower than 75 comes up there will be another round. So, if you are in round 6, the probability there will be a seventh round is 75% and if you are in round 18, the probability there will be another round is also 75%.

8. If play moves to Round 7 you will make your choices just as you did before and the stock of points will be the one that resulted from choices in Round 6. Your payoffs in Round 7 will be: 2.4 × ln(Points you take). That is, from Round 7 onwards your payoffs are no longer multiplied by 75% each time a new round starts. From Round 7 onwards your payoffs will always be computed as 2.4 × ln(Points you take). Once you and the other have made your choices for Round 7 with 75% probability there will be a new Round and the process repeats itself.

9. In the first 6 rounds, payoffs shrink by 75% every round, while in rounds 7 and above, payoffs don't shrink, but the match continues with probability 75%.

However, the computation of your expected average payoff is the same in all rounds of the match. In rounds 2 through 6 the average payoff is 75% × payoff function in the previous round. In rounds 7 and above the expected payoff before you know whether round 7 will take place is 75% × payoff function in the previous round (in case Round 7 takes place) + 25% × 0 (in case it doesn't). This means that the average payoff in rounds 7 and above is 75%× expected payoff function in the previous round. In other words, for all rounds of the match before and after Round 7 your expected payoff is 75% × payoff function in the previous round.

10. The match also ends if the Stock of Points Next Round is lower than 8 points. If the stock of points Next round is lower than 8 points the software considers that the stock will be depleted and will display 'Stock Next Round: 0'.

11. Between matches, you will be told the history of your decisions and the decisions of the person you were matched with in the last match. The box on the left of your screens will also display the history of your decisions for the current Match. If you wish to see your choices for a previous Match, simply type the Match number in the box after "History for Match...".

12. Part 1 will finish once 5 Matches have been played. At the end of the experiment we will add your payoff from every round and multiply it by $0.065 to determine your payoff for Part 1. There is no show-up fee for this experiment.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. The first 6 rounds will be played with certainty, but your payoffs in each new round are 75% of what they were in the previous round. After round 6, payoffs do not shrink but there is a 75% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match.

- Practice: As you can see your payoff may differ substantially depending on the decisions you and the person you were matched with make. In order to help you see a few Rounds ahead at any time when you make a decision the software allows you to compute hypothetical futures. If you RIGHT click with your mouse on top of one of the nine possible outcomes you can see what happens next round if that outcome is selected in the current round. Again you can repeat this procedure further to build hypothetical futures and compute cumulative payoffs for different outcomes. Your screens will change when you are computing hypothetical futures: the background is light Red to remind you that you are computing a hypothetical future and not yet made a decision in the current round. Whenever you want to go back to make a decision or start a new hypothetical future, click on the BACK button. Before we start with the experiment you will have 3 minutes to experiment with hypothetical futures. This will not count for payoffs but is just meant to familiarize you with the interface.

General Instructions: Part 2

1. The basic structure of Part 2 is similar to Part 1: you will play 4 matches as you did in Part 1 and you will be randomly matched with a new partner in each match.

2. In Round 1 of each match you will have to make a choice just as you did in Part 1.

3. In Part 2 you will make your choices for Round 2 before you learn what your partner selected in Round 1. That is, you will have to make THREE choices for Round 2: Make a choice if your partner selects "Take 12.5%" in Round 1. Make a choice if your partner selects "Take 20%" in Round 1. Make a choice if your partner selects "Take 50%" in Round 1.

4. After you submit your choices for Round 2, you will learn what your partner selected for Round 1 and the payoffs for Round 1.

5. In Round 2 the interface will implement what you previously specified. In other words, the software looks at both: what your partner chose in Round 1 and what you chose for Round 2 for that particular Round 1 decision of your partner. Then the software implements your decision for Round 2.

6. In all subsequent Rounds you will be asked to make THREE choices for the Next Round before you learn what your partner selected in that Round.

- Are there any questions?

Before we start, let me remind you that:

- The length of a match is randomly determined. The first 6 rounds will be played with certainty, but your payoffs in each new round are 75% of what they were in the previous round. After round 6, payoffs do not shrink but there is a 75% probability that the match will continue for at least another round. You will play with the same person for the entire match.

- After a match is finished, you will be randomly paired with another person for a new match.

Appendix B: Proofs

The optimization problem in this appendix has been widely studied in the literature. For a more detailed analysis see Sundaram (1989)

Proof of Proposition 1

The restriction to Markov strategies allows to introduce a value function $V(s)$ identifying the equilibrium value in any continuation game induced by a history ending in state s. Let $\sigma_t(s_0)$ identify the equilibrium strategy for each period t as a function of the initial stock. Then, the function $V(s_0)$ identifies equilibrium utilities and must satisfy:

$$V(s_0) = (1-\delta)\sum_{t=0}^{\infty} \delta^t \ln(\sigma_t(s_0)).$$

According to the Markov restriction, current actions depend only on the current state and consequently each player's strategy can be expressed by a function $\sigma(s)$. Suppose that the equilibrium strategy is linear: $\sigma(s) = a_M s$. Since the linear strategy is used by both players (1) implies: $s_t = \left[(1+r)(1-2a_M)\right]^t s_0$. The value function can be expressed as:

$$V(s_0) = \ln(a) + \frac{\delta}{(1-\delta)} \ln\left((1+r)(1-2a_M)\right) + \ln(s_0) \tag{6}$$

Using the one-shot deviation principle, $\sigma(s)$ solves the Bellman equation,

$$\sigma(s) \in \arg\max_{\hat{a}\in A}(1-\delta)\ln(\hat{a}s) + \delta V((1+r)(1-\hat{a}-a)s),$$

where $\hat{a}s$ represents i's consumption, and the other player adheres to the equilibrium strategy $\sigma(s) = as$. Assume that the value function is differentiable (it is straightforward from (6) that it will indeed be the case in equilibrium), the first-order condition is:

$$\frac{(1-\delta)}{\sigma(s)} = \delta(1+r)V'\left((1+r)(s-2\sigma(s))\right). \tag{7}$$

Use (6) to get $V'(s) = 1/s$, and after some algebra $a_M = \frac{1-\delta}{2-\delta}$. Notice that in (7) an increase in r has two effects on the right-hand side, which represents the marginal cost of a current increase in extraction. On the one hand an increase in r directly raises the marginal cost, leading to lower extraction in the present. On the other hand (via V') it reduces the marginal cost. When the utility is $\ln(C)$ both effects cancel each other and r does not affect the extraction rate. Finally, the conditions of Theorem 4.5 in Stokey et al. (1989) are satisfied, guaranteeing that working with the Bellman equation effectively leads to a solution of the recursive problem.

Proof of Proposition 2

To determine efficient extraction consider the maximization of aggregate utility by a benevolent planner. By concavity it is clear that a planner would allocate consumption equally across all individuals. The value function is given by:

$$V^P(s_0) = (1-\delta)\sum_{t=0}^{\infty}\delta^t 2\ln(P_t(s_0)),$$

where $P_t(s_0)$ identifies the equilibrium choice of a planner for a given initial stock of the resource. Representing the dynamic programming problem recursively, the first-order condition of the corresponding Bellman equation implies:

$$\frac{2(1-\delta)}{P(s)} = 2\delta(1+r)V^{P'}\left((1+r)(s-2P(s))\right)$$

Guessing that the solution to the problem is linear ($P(s) = a_L s$), algebra shows that

$a_L = \frac{1-\delta}{2}$. It is readily verifiable that the conditions of Theorem 4.5 in Stokey et al. (1989) are satisfied. Moreover, uniqueness results from the problem being strictly concave.

Proof of Proposition 3

With $\tilde{A} = \{a_L, a_M, a_H\}$ (6) still holds. Assume that player j is always consuming at a_M and consider a one-shot deviation from $a_M = \frac{1-\delta}{2-\delta}$ to $\tilde{a} \in \{a_L, a_H\}$ by i. Next period's stock would be given by: $s^D = (1+r)(1 - a_M - \tilde{a})s$. Subject i would not want to deviate if $V(s) \geq (1-\delta)\ln(\tilde{a}s) + \delta V(s^D)$, or using (6) and simplifying:

$$(1-\delta)\ln\left(\frac{a_M}{\tilde{a}}\right) \geq \delta \ln\left(\frac{1 - a_M - \tilde{a}}{1 - 2a_M}\right) \tag{8}$$

It is readily verifiable that (8) holds for any $\tilde{a} \in \{a_L, a_H\}$, with $a_L = \frac{1-\delta}{2}$ and $\frac{1+\delta}{2} > a_H > a_M$.

Similarly, a_L was the prescription of a social planner when a_H and a_M were among the alternatives, it will also be the choice when only those other alternatives are available.

Proof of Proposition 4

At any period $\tilde{t}$ (6) can be used to compute the value the normal phase, $V^{N=}V\left(s_{\tilde{t}}; a_L\right)$. Under defection from any considered alternative the agent will enjoy instant consumption at the optimal defection level (a^D) and from next period onwards the corresponding punishment phase begins, with the stock starting at $s^D = (1+r)\left(1 - a_L - a^D\right) s_{\tilde{t}}$. This allows to compute the corresponding value of defection, V^D. The considered strategy is successful in supporting efficient extraction given δ if for all $s_{\tilde{t}}$, $V^C \geq V^D$.

Consider a Grim LM-MPE strategy:

$$a_t^G = \begin{cases} a_L & \text{if } t = 0 \text{ or if } t > 0 \text{and } a_{\tilde{t}}^i = a_L \text{ for } i = \{1, 2\} \text{ and all } \tilde{t} < t \\ a_M & \text{otherwise} \end{cases}$$

The continuation value of defection is given by:

$$V^G(s) = (1-\delta)\ln(a^D) + \delta\ln((1+r)(1-a_L-a^D)) + \delta\ln(1-\delta) - \frac{\delta}{(1-\delta)}\ln(2-\delta) + \frac{\delta^2}{(1-\delta)}\ln((1+r)\delta) + \ln(s)$$

If any deviation in the unit interval was available, the optimal defecting extraction action would be given by $a^D = \frac{1-\delta^2}{2}$. When the action space is $\tilde{A}$, it is straightforward to show that if $\delta > 0.293$, then $a^D = a_M$ and $a^D = a_H$ otherwise. The inequality $V^N(s; a_L) \geq V^G(s)$ depends only the discount rate. Computations show that cooperation is sustainable if $\delta > 47.6\%$.

A T period punishment (TPP) strategy is defined by:

$$a_t^{TPP} = \begin{cases} a_L & \text{if } t = 0 \\ a_L & \text{if } a_{t-1}^i = a_L \text{for } i = \{1,2\} \text{or } a_{\tilde{t}}^1 = a_{\tilde{t}}^2 = a_M \text{for } \tilde{t} = t - T > 0 \\ a_M & \text{otherwise} \end{cases}$$

The corresponding continuation value of defection is given by:

$$V^{TPP}(s) = (1-\delta)\ln\left(a^D\right) + \delta\ln\left((1+r)(1-a_L-a^D)\right) + \delta\ln(1-\delta) - \frac{\delta(1-\delta^T)}{(1-\delta)}\ln(2-\delta) - \delta^{T+1}\ln(2) + \frac{\delta^2}{(1-\delta)}\ln((1+r)\delta) + \ln(s)$$

The optimal defecting action is identical to the one characterized for the Grim LM-MPE case. Notice that as $T \to \infty$, $V^{TPP}(s) \to V^G(s)$. Naturally, the discount rate that makes efficient extraction sustainable depends on T. Computations show that for $T \geq 8$ efficient extraction is sustainable if $\delta > 47.6\%$, just as in the $T = \infty$ case. For $T = 1$, the requirement is $\delta > 76.5\%$. If $T = 2$, $\delta > 55.8\%$.

The Win-Stay Lose-Shift strategy is defined as:

$$a_t^{WSLS} = \begin{cases} a_L & \text{if } t = 0 \text{ or } t > 0 \text{and } a_{t-1}^i = a_{t-1}^j \text{for } i,j = \{1,2\} \\ a_M & \text{otherwise} \end{cases}$$

Notice that after one period of punishment this strategy goes back to cooperation.

Therefore the continuation value $V^{WSLS}(s)$ is equal to $V^{TPP}(s)$ with $T = 1$.

Tit for Tat (TfT) is defined by:

$$a_t^{TfT} = \begin{cases} a_L & \text{if } t = 0 \\ a_{t-1}^{-i} & \text{otherwise, } -i \text{ identifies the other player} \end{cases}$$

The continuation value of defection is given by:

$$V^{TfT}(s) = \frac{\left(\ln(a^D) + \delta \ln(a_L)\right)}{(1+\delta)} + \frac{\delta}{(1-\delta)} \ln\left((1+r)(1-a_L-a^D)\right) + \ln(s)$$

Notice that the optimal deviation would be different than previous cases if actions could be selected in the unit interval. With the action space in $\tilde{A}$, a_M is the optimal deviation. If $\delta > 55.8\%$ it can be verified that $V^{TfT}(s) > V^N(s)$. Unlike the previous strategies considered, there are profitable deviations from Tit for Tat's punishment phase. A one-shot deviation in the first punishment period, leads to permanent consumption at the efficient level, which is more rewarding for all δ. Therefore, Tit for Tat is not subgame-perfect.

Experimental Game

The evolution of the stock is computed as:

$$\hat{s}_{t+1} = \text{round}\left((1+r)\left[(1-a_t^1-a_t^2)s_t\right]\right) = \left((1+r)\left[(1-a_t^1-a_t^2)s_t\right]\right)\varepsilon_{t+1} \qquad (9)$$

where $\varepsilon_{t+1} \in \mathbb{R}_+$ and $\varepsilon_0 = 1$. If $\hat{s}_{t+1} \geq \underline{s} = 8$, then $s_{t+1} = \hat{s}_{t+1}$, otherwise $s_{t+1} = 0$.

Claim 1: Extracting a_M constitutes a symmetric MPE of this modified game provided r is selected so that if extraction is always at a_M the stock is never depleted. Indeed that is the case for r^{Low} and r^{High}. For any starting s, the value of always extracting at a_1 is given by:

$$V_{EG}(s) = \ln(a_M) + \ln(s) + \frac{\delta}{(1-\delta)} \ln\left((1+r)(1-2a_M)\right) + \left[(1-\delta)\sum_{t=0}^{\infty} \delta^t \ln\left(\prod_{k=0}^{t} \varepsilon_k(s,r)\right)\right] \qquad (10)$$

Notice that the only difference with respect to the value function in (6) is the addition of the term in brackets, which captures the adjustments due to rounding. The adjustment values depend on s and r. If there was no rounding all, ε_k values would equal one, but given the state space, $\varepsilon_j \in [0.95, 1.07]$ for all j. This means that for large values of t the argument of the logarithm in the term in brackets is close to one, but there may be non-zero terms in the summation for low t.

Assume that subject j is always extracting at a_M and consider a one-shot deviation from a_M to $\tilde{a} \in \{a_L, a_H\}$ by subject i. If $s^D = \text{round}\left((1+r)(1-a_1-\tilde{a})s\right) \geq \underline{s}$, subject i would not want to deviate if $V_{EG}(s) \geq (1-\delta)\ln(\tilde{a}s) + \delta V_{EG}(s^D)$, or using (10) and simplifying:

$$\ln\left(\frac{a_M}{\tilde{a}}\right) - \frac{\delta}{1-\delta}\ln\left(\frac{1-a_M-\tilde{a}}{1-2a_M}\right) \geq \sum_{t=0}^{\infty}\delta^t \ln\left(\prod_{k=0}^{t}\frac{\varepsilon_k(s^D, r)}{\varepsilon_k(s, r)}\right) \tag{11}$$

Notice that the left-hand side of (11) is constant, while the right-hand side depends on the starting value of s. If, instead $s^D = 0$, then a deviation would not be profitable if $V_{EG}(s) \geq (1-\delta)\ln(\tilde{a}s)$, or equivalently:

$$\begin{aligned} \ln(a_M) + \frac{\delta}{(1-\delta)}\ln\left((1+r)(1-2a_M)\right) - (1-\delta)\ln(\tilde{a}) \geq \\ -\ln(s) - \left[(1-\delta)\sum_{t=0}^{\infty}\delta^t \ln\left(\prod_{k=0}^{t}\varepsilon_k(s, r)\right)\right] \end{aligned} \tag{12}$$

For treatments with $r^{Low} = 0.57$ and $r^{High} = 0.92$ it is straightforward to verify numerically that subject i will not want to deviate from choosing a_M for any s in $S^{Lab} = \{\underline{s}, ..., \bar{s}\}$, where $\bar{s} = 10,000$. *Claim 2:* In a similar fashion it can be verified that efficiency involves extracting always at $a = a_L$. *Claim 6:*Adding the same modifications and using the same procedure it can be verified that efficient extraction can be supported as a subgame-perfect equilibrium if the punishment phase is that of Grim LM-MPE and TPP with $T \geq 2$.

Claim 3: Consider a threshold strategy that prescribes a_H if $s \leq \bar{s}$ and take any $\tilde{s} \leq \bar{s}$. The payoff from not deviating are: $(1-\delta)(\ln(1/2\,\tilde{s}))$ and the stock is depleted. If the stock is depleted for any alternative action, there is clearly no profitable deviation. So, consider the case in which a deviation to $\tilde{a}$ is such that $\tilde{s}' = \text{round}\left((1+r)(1-\tilde{a}-1/2)\tilde{s}\right) > 8$. For the levels of r considered in the parametrization whenever one agent selects a_H the stock next period is lower: $\tilde{s}' < \tilde{s}$.

Treatment	Matches	Actions			Outcomes					
		a_H	a_M	a_L	a_H, a_H	a_M, a_M	a_H, a_L	a_M, a_M	a_M, a_L	a_L, a_L
a_H allowed r^{Low}	All	8.23	79.84	11.93	1.05	13.28	1.16	70.92	9.27	4.32
	Last 4	8.98	83.12	7.90	1.41	14.31	1.41	74.35	6.59	1.88
a_H allowed r^{High}	All	14.48	74.63	11.17	3.05	21.14	1.43	61.81	9.62	2.95
	Last 4	14.30	79.23	6.47	2.45	21.60	0.61	68.30	6.95	0.61
a_H not allowed r^{Low}	All	-	72.00	28.00	-	-	-	63.95	23.79	12.26
	Last 4	-	75.06	24.94	-	-	-	67.56	22.27	10.16
a_H not allowed r^{High}	All	-	83.35	19.42	-	-	-	79.22	14.79	5.99
	Last 4	-	86.30	13.70	-	-	-	83.24	12.70	4.05

Table 9: Distribution of Actions and Outcomes across Treatments (in %)

If there is no $\tilde{a} \in \{a_L, a_M\}$ such that $(\ln(\tilde{a}\tilde{s}) + \delta \ln(1/2\tilde{s}')) > (\ln(1/2\tilde{s}))$, then there is no profitable deviation from a_H at $\tilde{s}$. For r^{Low} such a profitable deviation exists if $\bar{s} > 15$. Naturally, there are no profitable deviations when the strategy prescribes the choice of a_M, thus there are x-threshold MPE for $\bar{s} \in \{8, 9, ...15\}$ if $r = 0.57$. Likewise if $r = 0.92$, then there are x-threshold MPE for $\bar{s} \in \{8, 9, ...13\}$. In particular, notice that since $s_0 = 12$, there are equilibria such that there is depletion in the first period. Any of these threshold MPE support cooperation as a subgame-perfect equilibrium for $\delta = 0.75$.

Claim 4: Let S^E be the set of 'efficient' states, reached if both players always select a_L. Efficient extraction using Markov strategies must prescribe the choice of a_L for all states in S^E. For r^{High}, consider $s = 27$ and for r^{Low}, $s = 13$. In either case the selected state does not belong to S^E. If the Markov strategy prescribes the choice of: 1)a_H, then any deviation is profitable, 2) a_M, then a deviation to a_L is profitable, and 3) a_L, then a deviation to a_M is profitable.

Appendix C: Additional Empirical Analysis

Evolution of choices within the session

Table 9 presents in detail the proportion of times each choice and outcome was selected for all matches and the last 4.

Table 10 provides the estimates of random effects probit models in which the dependent variable is the specific choice of the subject. The set of controls includes

	Random Effects Ordered Probit Estimations a_H allowed				Random Effects Probit Estimations a_H not allowed			
	r^{Low}		r^{High}		r^{Low}		r^{High}	
Variable	Coeff	Std. Err.	Coeff	Std. Err.	Coeff	Std. Err.	Coeff	Std. Err.
Stock of Points	-.020	.016	.002	.004	-.027***	.009	-.001***	.000
$Own_{t-1} = a_L$	-2.062***	.401	-.268*	.142				
$Own_{t-1} = a_M$	-1.192***	.384	.200*	.110	.085	.056	.439***	.074
$Other_{t-1} = a_L$	-1.414***	.382	-.501***	.141				
$Other_{t-1} = a_M$	-.923***	.360	-.290***	.110	.208***	.057	.509***	.077
$Own_1 = a_M$	.087	.128	.263***	.088	.357***	.068	.636***	.076
$Other_1 = a_M$	-.017	.128	.090	.077	-.058	.064	.312***	.078
Match 2	.239*	.146	.328***	.127	-.063	.106	.350***	.129
Match 3	.243	.154	.055	.128	.099	.112	.205	.132
Match 4	.312**	.157	.282**	.135	.155	.112	.073	.132
Match 5	.390**	.154	.139	.125	-.008	.106	.275***	.130
Match 6	.452***	.155	.267***	.129	.099	.106	.186	.126
Match 7	.489***	.151	.197	.125	.106	.103	.363***	.123
Match 8	.373**	.158	.190	.128	.097	.110	.256*	.133
Match 9	.410***	.157	.236*	.135	-.016	.110	.247*	.130
Round 3	.075	.109	-.036	.101	-.157	.089	-.184	.125
Round 4	.101	.115	.018	.104	-.016	.090	-.350***	.122
Round 5	.209*	.122	-.105	.109	.141	.094	-.249**	.122
Round 6	.585***	.125	.127	.112	.325***	.099	-.253**	.124
Round 7	.907***	.152	.461***	.124	.444***	.107	-.197	.129
Round 8	.673***	.181	.605***	.139	.680***	.123	-.252	.140
Round 9	.687***	.254	.695***	.172	.641***	.139	.048	.169
Round 10 and above	1.122***	.349	.771***	.173	.701***	.103	.182	.157
Constant					.375*	.197	-.094	.164
Constant 1	-3.533***	.547	-.985***	.208				
Constant 2	-.377	.544	1.542***	.211				
Number of Subjects	42		38		40		40	

*** significant at 1%, ** at 5%, * at 10%.

Table 10: Probit Estimations: Evolution of Choices within the session and rounds

round and match dummies (when it is feasible to identify them), the past own and partner's decisions, the current stock and the choices for round 1. The inclusion of round 1 decisions control for dynamic unobserved effects.

Perfectly fit strategies

Table 11 presents, for each treatments, the proportion of histories that are perfectly explained by some strategy. If some history can be explained with two alternative strategies it will be counted twice. In practice there are very few such cases.

Unified history dependent strategies

History dependent strategies and their exploitive versions can be re-interpreted as a single strategy. For example, for each individual I estimate whether Grim or Exploitive Grim LM-MPE better explain their choices and redefine Grim LM-MPE accordingly. Table 12 presents the estimates following this adjustment. The conclusions

Strategies	with a_H		without a_H		Total
	r_L	r_H	r_L	r_H	
Always a_H	18	6			24
	(4.76)	(1.75)			(1.67)
Always a_M	125	49	57	141	372
	(33.07)	(14.33)	(15.83)	(39.17)	(25.83)
Always a_L	6	2	1	10	19
	(1.59)	(0.58)	(0.28)	(2.78)	(1.32)
Grim LM-MPE	15	0	13	17	45
	(3.97)	(0.00)	(3.61)	(4.72)	(3.13)
Grim a_H	4	1			5
	(1.06)	(0.29)			(0.35)
Grim (a_M, a_H)	7	10			17
	(1.85)	(2.92)			(1.18)
Tit for Tat	13	12	13	17	55
	(3.44)	(3.51)	(3.61)	(4.72)	(3.82)
Tat for Tit and a_M	29	32	13	31	105
	(7.67)	(9.36)	(3.61)	(8.61)	(7.29)
a_M and a_H if $R \geq 6$	44	33			77
	(11.64)	(9.65)			(5.35)
Tat for Tit and a_H	18	10			28
	(4.76)	(2.92)			(1.94)
a_M and a_H if s High	18	10			28
	(4.76)	(2.92)			(1.94)
Total Explained Histories	297	165	97	216	775
	(78.57)	(48.25)	(26.94)	(60.00)	(53.82)
Total Histories in Treatment	378	342	360	360	1440

As a percentage of Total Histories in Treatment in parentheses.

Table 11: Taxonomy of observed strategies by treatment

Strategies	a_H allowed		a_H not allowed	
	r^{Low}	r^{High}	r^{Low}	r^{High}
Always a_H	0.022***	0.000		
	(0.009)	(0.003)		
Always a_M	0.661***	0.710**	0.482***	0.657***
	(0.126)	(0.093)	(0.106)	(0.098)
Always a_L	0.009	0.000	0.000	0.025
	(0.049)	(0.003)	(0.000)	(0.029)
Threshold $a_H < 10$	0.000	0.173*		
	(0.000)	(0.155)		
Grim LM-MPE	0.139*	0.015	0.381***	0.117*
	(0.084)	(0.069)	(0.079)	(0.071)
Grim a_H	0.048	0.000		
	(0.039)	(0.009)		
Grim (a_M,a_H)	0.000	0.026		
	(0.001)	(0.031)		
Tit for Tat	0.123	0.025	0.072	0.188***
	(0.069)	(0.028)	(0.052)	(0.075)
Tat for Tit and a_M	0.000	0.000	0.063	0.011
	(0.000)	(0.006)		
a_M and a_H if $R \geq 6$	0.074	0.049		
γ	0.695	0.983	0.708	0.520
β	0.808	0.744	0.804	0.872

Bootstrapped standard errors in parentheses. Data from Matches 6-9 including unimplemented choices from one-step ahead strategy method

Table 12: Estimation of Used strategies by treatment

presented in the text remain unchanged, but this change allows for higher precision in the estimations.

Does the strategy method affect used strategies?

I now present several robustness exercises. One concern is that part 2 of the experiment with the one-period ahead strategy method changes the strategies selected by subjects. Table 13 presents the estimations for all treatments constraining to use either matches 4 and 5 (before part 2 starts) or matches 6 and 7 (the first 2 matches of part 2). For matches 4 and 5 only actual choices are available. In the provided estimations I have restricted analysis only to actual choices in matches 6 and 7 for the comparison to be valid. There are small quantitative changes, but significant strategies remain the same and the ordering and relative importance of non-significant strategies does not change.

Strategies	a_H allowed				a_H not allowed			
	r^{Low}		r^{High}		r^{Low}		r^{High}	
	Matches 4-5	Matches 6-7	Matches 4-5	Matches 6-7	Matches 4-5	Matches 6-7	Matches 4-5	Matches 6-7
Always a_H	0.000	0.020*	0.000	0.000				
	(0.000)	(0.011)	(0.000)	(0.000)				
Always a_M	0.693**	0.604***	0.701**	0.831***	0.298***	0.414***	0.623***	0.725***
	(0.148)	(0.184)	(0.150)	(0.143)	(0.117)	(0.112)	(0.101)	(0.127)
Always a_L	0.048	0.042	0.000	0.000	0.019	0.000	0.050	0.025
	(0.045)	(0.056)	(0.004)	(0.001)	(0.044)	(0.000)	(0.043)	(0.029)
Threshold $a_H < 10$	0.000	0.000	0.057	0.016				
	(0.000)	(0.000)	(0.057)	(0.066)				
Grim LM-MPE	0.000	0.000	0.139	0.067	0.484***	0.469***	0.285***	0.199*
	(0.014)	(0.035)	(0.136)	(0.143)	(0.131)	(0.116)	(0.098)	(0.107)
Grim a_H	0.000	0.016	0.028	0.015				
	(0.000)	(0.024)	(0.032)	(0.017)				
Grim (a_M,a_H)	0.071	0.029	0.057	0.016				
	(0.064)	(0.053)	(0.057)	(0.066)				
Tit for Tat	0.085	0.181	0.000	0.000	0.142	0.073	0.040	0.049
	(0.089)	(0.136)	(0.052)	(0.045)	(0.124)	(0.075)	(0.062)	(0.071)
Tat for Tit and a_M	0.024	0.000	0.000	0.000	0.055	0.042	0.000	0.000
	(0.032)	(0.000)	(0.000)	(0.000)				
a_M and a_H if $R \geq 6$	0.077	0.104	0.072	0.068				
γ	0.513	0.600	0.771	0.829	0.801	0.664	0.463	0.449
β	0.875	0.841	0.785	0.769	0.777	0.818	0.898	0.902

Bootstrapped standard errors in parentheses. Data includes only actual choices for matches 6-7.

Table 13: Estimation of Strategies Used before and after the start of part 2

References

Acemoglu, D., "Political Economy Lecture Notes," *mimeo*, 2011.

_ **, G. Egorov, and K. Sonin**, "Political Selection and Persistence of Bad Governments," *The Quarterly Journal of Economics*, 2010, *125* (4), 1511–1575.

Ausubel, L.M. and R.J. Deneckere, "Reputation in bargaining and durable goods monopoly," *Econometrica*, 1989, pp. 511–531.

Battaglini, M. and S. Coate, "Inefficiency in Legislative Policymaking: A Dynamic Analysis," *The American Economic Review*, 2007, pp. 118–149.

_ **and** _ , "A Dynamic Theory of Public Spending, Taxation, and Debt," *The American Economic Review*, 2008, *98* (1), 201–236.

_ **, S. Nunnari, and T. Palfrey**, "Legislative Bargaining and the Dynamics of Public Investment," *mimeo*, 2011.

Benhabib, J. and A. Rustichini, "Social conflict and growth," *Journal of Economic Growth*, 1996, *1* (1), 125–142.

Bergemann, D. and J. Valimaki, "Dynamic common agency," *Journal of Economic Theory*, 2003, *111* (1), 23–48.

Bernheim, B.D. and D. Ray, "Economic Growth with Intergenerational Altruism," *The Review of Economic Studies*, 1987, pp. 227–241.

Cabral, L., E.Y. Ozbay, and A. Schotter, "Intrinsic and instrumental reciprocity: An experimental study," *mimeo*, 2011.

Chermak, J.M. and K. Krause, "Individual response, information, and intergenerational common pool problems," *Journal of Environmental Economics and Management*, 2002, *43* (1), 47–70.

Clark, C., *Mathematical bioeconomics*, New York, NY (USA); John Wiley and Sons Inc., 1990.

Coate, S. and S. Morris, "Policy persistence," *The American Economic Review*, 1999, *89* (5), 1327–1336.

Coles, M.G. and D.T. Mortensen, "Dynamic Monopsonistic Competition and Labor Market Equilibrium," *mimeo*, 2011.

Dal Bó, P. and G.R. Fréchette, "The evolution of cooperation in infinitely repeated games: Experimental evidence," *The American Economic Review*, 2011, *101* (1), 411–429.

Diamond, J.M., *Collapse: How societies choose to fail or succeed*, Penguin Group USA, 2006.

Dixit, A., G.M. Grossman, and F. Gul, "The dynamics of political compromise," *Journal of Political Economy*, 2000, pp. 531–568.

Dutta, P.K., "A folk theorem for stochastic games," *Journal of Economic Theory*, 1995, *66* (1), 1–32.

_ **and R. Radner**, "Population growth and technological change in a global warming model," *Economic Theory*, 2006, *29* (2), 251–270.

_ **and R.K. Sundaram**, "Markovian equilibrium in a class of stochastic games: existence theorems for discounted and undiscounted models," *Economic Theory*, 1992, *2* (2), 197–214.

_ **and** _ , "The tragedy of the commons?," *Economic Theory*, 1993, *3* (3), 413–426.

Ericson, R. and A. Pakes, "Markov-perfect industry dynamics: A framework for empirical work," *The Review of Economic Studies*, 1995, *62* (1), 53.

Fischbacher, U., "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Fischer, M.E., B. Irlenbusch, and A. Sadrieh, "An intergenerational common pool resource experiment," *Journal of Environmental Economics and Management*, 2004, *48* (2), 811–836.

Fréchette, G., A. Schotter, and S. Yuksel, "Implementing Infinitely Repeated Games In the Laboratory," *mimeo*, 2011.

Fudenberg, D. and J. Tirole, "Game Theory, 1991," 1991.

_ , **D.G. Rand, and A. Dreber**, "Slow to anger and fast to forgive: Cooperation in an uncertain world," *The American Economic Review*, 2011, *forthcoming.*

Herr, A., R. Gardner, and J.M. Walker, "An Experimental Study of Time-Independent and Time-Dependent Externalities in the Commons* 1,* 2," *Games and Economic Behavior*, 1997, *19* (1), 77–96.

Hey, J.D. and V. Dardanoni, "Optimal consumption under uncertainty: An experimental investigation," *The Economic Journal*, 1988, *98* (390), 105–116.

Hörner, J. and L. Samuelson, "Incentives for Experimenting Agents," *mimeo*, 2009.

Hutchings, JA, "Collapse and recovery of marine fishes," *Nature*, 2000, *406*, 882–885.

Laibson, D., "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Lancaster, K., "The dynamic inefficiency of capitalism," *The Journal of Political Economy*, 1973, pp. 1092–1109.

Lei, V. and C.N. Noussair, "An experimental test of an optimal growth model," *American Economic Review*, 2002, pp. 549–570.

Levhari, D. and L.J. Mirman, "The great fish war: an example using a dynamic Cournot-Nash solution," *The Bell Journal of Economics*, 1980, pp. 322–334.

Mailath, G.J. and L. Samuelson, *Repeated games and reputations: long-run relationships*, Oxford University Press, USA, 2006.

Maskin, E. and J. Tirole, "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 2001, *100* (2), 191–219.

Noussair, C. and K. Matheny, "An experimental study of decisions in dynamic optimization problems," *Economic Theory*, 2000, *15* (2), 389–419.

Ostrom, E., "Revising theory in light of experimental findings," *Journal of Economic Behavior & Organization*, 2010, *73* (1), 68–72.

— **and J. Walker**, *Communication in a commons: Cooperation without external enforcement*, In: Palfrey, T.R. (Ed.), Laboratory research in political economy, Ann Arbor: University of Michigan, 1991.

— **, R. Gardner, and J. Walker**, *Rules, games, and common-pool resources*, Univ of Michigan Pr, 1994.

Reinganum, J.F. and N.L. Stokey, "Oligopoly extraction of a common property natural resource: The importance of the period of commitment in dynamic games," *International Economic Review*, 1985, *26* (1), 161–173.

Roth, A.E. and J.K. Murnighan, "Equilibrium behavior and repeated play of the prisoner's dilemma," *Journal of Mathematical Psychology*, 1978, *17* (2), 189–198.

Rubinstein, A. and A. Wolinsky, "Decentralized trading, strategic behaviour and the Walrasian outcome," *The Review of Economic Studies*, 1990, *57* (1), 63.

Saijo, T., K. Sherstyuk, N. Tarui, and M. Ravago, "Games with Dynamic Externalities and Climate Change Experiments," *mimeo*, 2009.

Stokey, N.L., R.E. Lucas, and E.C. Prescott, *Recursive methods in economic dynamics*, Harvard Univ Pr, 1989.

Sturm, B. and J. Weimann, "Experiments in environmental economics and some close relatives," *Journal of Economic Surveys*, 2006, *20* (3), 419–457.

Sundaram, R.K., "Perfect equilibrium in non-randomized strategies in a class of symmetric dynamic games," *Journal of Economic Theory*, 1989, *47* (1), 153–177.

Tornell, A. and A. Velasco, "The tragedy of the commons and economic growth: why does capital flow from poor to rich countries?," *Journal of Political Economy*, 1992, pp. 1208–1231.