



Subject-specific performance information can worsen the tragedy of the commons: Experimental evidence

Mauricio G. Villena ^{a,*}, Franco Zecchetto ^b

^a School of Business, Adolfo Ibáñez University, Chile

^b School of Business and School of Engineering, Adolfo Ibáñez University, Chile

ARTICLE INFO

Article history:

Received 16 January 2007

Received in revised form 22 December 2010

Accepted 23 December 2010

Available online 31 December 2010

JEL classification:

C72

C91

D83

Q2

PsycINFO classification:

2260

2343

3020

Keywords:

Common property resources

Information

Learning and imitation

Experimental economics

ABSTRACT

The main aim of this article is to investigate the behavioral consequences of the provision of subject-specific information in the group effort levels chosen by players in an experimental CPR game. We examine two basic treatments, one with incomplete information and the other with complete information. In the former, subjects are informed only about their own individual payoffs and the aggregate extraction effort level of the group, and in the latter they are also informed about the individual effort levels and payoffs of each subject. Given this setting, the basic question we attempt to answer is: Will the provision of subject-specific performance information (i.e. individual's effort levels and payoffs) improve or worsen the tragedy of the commons (i.e. an exploitation effort level greater than the socially optimum level)? In order to motivate our hypotheses and explain our experimental results at the individual level, we make use of the theory of learning in games, which goes beyond standard non-cooperative game theory, allowing us to explore the three basic benchmarks in the commons context: Nash equilibrium, Pareto efficient, and open access outcomes. We use several learning and imitation theoretical models that are based on contrasting assumptions about the level of rationality and the information available to subjects, namely: best response, imitate the average, mix of best response and imitate the average, imitate the best and follow the exemplary learning rules. Finally, in order to econometrically test the hypotheses formulated from the theoretical predictions we use a random-effects model to assess the explanatory power of the different selected behavioral learning and imitation rules.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

The standard economic theory on common property resource (henceforth CPR) exploitation predicts a non-cooperative result in the commons implying the economic overexploitation of the resource. In a non-cooperative game theoretic context, this implies that the Nash equilibrium (henceforth NE) of the CPR game is inefficient, entailing an exploitation effort level greater than the Pareto efficient (henceforth PE), i.e. where marginal revenue equals marginal cost.¹ This economic prediction about CPR use has been empirically tested in practice with field studies and experimental works recreating the commons setting.²

* Corresponding author. Address: School of Business, Adolfo Ibáñez University, Chile. Av. Diagonal las Torres 2640, Peñalolén. Santiago, Chile. Tel.: +56 2 3311492; fax: 56 2 2784413.

E-mail address: mauricio.villena@uai.cl (M.G. Villena).

¹ For details about the economic theory on CPR exploitation see, inter alia, Dasgupta and Heal (1979), Gordon (1954), Seabright (1993), and Stevenson (1991).

² For reviews about CPR field work see, inter alia, Bromley (1992) and Ostrom (1990). For a review about CPR experimental work, see Ostrom (2006).

In particular, the experimental literature on the commons has received increasing attention from social scientists in the last two decades.³ Some laboratory experiments have been mainly designed to investigate whether experimental results conform to the economic predictions on CPR exploitation. These studies have concluded that at the aggregate level behavior does approximate the NE level, being therefore consistent with the basic non-cooperative game theoretic prediction (Ostrom, Gardner, & Walker, 1994; Walker, Gardner, & Ostrom, 1990: 209). Some studies have also examined what happens with this basic result when *face-to-face communication* is allowed, by having, at some point, an open group discussion before each round of the experiment. These works have shown that repeated face-to-face communication can be extremely successful in increasing group returns as well as subjects' overall performance (Cardenas, 2000, 2003; Ostrom & Walker, 1991; Ostrom et al., 1994; Sally, 1995). In addition to allowing cheap talk, some works have also investigated the *impact of sanctioning* in CPR experiments by giving subjects the opportunity to pay a fee in order to impose a fine on the payoffs received by another player (Casari & Plott, 2003; Ostrom, Walker, & Gardner, 1992, 1994). These studies have found that in contrast to what economic theory predicts (no costly sanctioning in a finitely repeated game), players do sanction, reacting both to the cost of sanctioning and to the fee-to-fine relationships, and that in general a sanctioning system greatly improves the efficiency of resource use. Finally, some papers have also assessed the willingness and ability of subjects to look for and adopt *different institutional arrangements* designed to improve social outcomes in a CPR setting (Cardenas, Stranlund, & Willis, 2000; Rodriguez-Sickert, Guzmán, & Cárdenas, 2008; Walker, Gardner, Herr, & Ostrom, 2000). These papers point the importance of collective choice, showing that giving participants the opportunity of proposing allocation rules can significantly increase efficiency and that the imposition of external regulatory rules, including external enforcement, can imply a poorer performance in terms of group returns, as it can crowd out group regarding behavior in favor of greater self interest.

Specifically, this paper relates to a recent experimental work by Apestequia (2006), who analyses the impact of information on players' decisions regarding the exploitation of a commonly shared resource, asking the question: *Does information matter in the commons?* In that article the author specially investigates the effects of the provision of information about the *payoff structure* in the context of a CPR game. He runs two treatments, one with complete information on the payoff structure and the other with none, the main result of the paper being that aggregate behavior is not significantly different between the two treatments. In both cases the aggregate effort levels converge to the NE.

In this work we want to add to the experimental literature on the commons by building upon and complementing the work of Apestequia. Particularly, we examine the behavioral consequences of the provision of *subject-specific information* on the aggregate effort levels chosen by players in an experimental CPR game. As did Apestequia (2006), we consider two basic scenarios, one with incomplete information, henceforth treatment 'incomplete', and another with complete information, henceforth treatment 'complete'. In the former, subjects are informed only about their own individual payoffs and the aggregate extraction effort level of the group, while in the latter they are also informed about the individual effort levels and payoffs of each subject. Given this setting, the basic question we attempt to answer is: *Will the provision of subject-specific performance information (i.e. individuals' effort levels and payoffs) improve or worsen the tragedy of the commons (i.e. an exploitation effort level greater than the socially optimum level)?*

Our paper adds to and complements Apestequia's and previous experimental works on the commons in several respects. First, although the provision of information regarding the payoff structure is quite relevant in the commons, the information about subject-specific performance, such as other individuals' effort levels and payoffs, is also crucial in the decision-making process frequently faced in the commons, adding a different perspective to the problem. Indeed, additional information about other players' actions and profits adds a *potential strategic and psychological component* to the analysis, which may well change Apestequia's basic result in terms of whether or not information matters in the commons. In fact, while the results of Apestequia's work show that aggregate behavior is not significantly different between the two treatments, converging in both cases to the NE effort level, we postulate as a main hypothesis of our paper that in treatment 'incomplete' the average aggregate use level will match the NE outcome, but in treatment 'complete' the average group effort level will be significantly higher than the NE. In other words, we expect to prove that more information in the commons can make a difference in terms of aggregate results.

Second, in order to motivate our hypotheses and explain our experimental results at the individual level, we make use of the theory of learning in games (see, e.g., Fudenberg & Levine, 1999), which goes beyond standard game theory, allowing us to explore three basic benchmarks in the commons context: NE, PE and Open Access (henceforth OA) outcomes. The rationality behind using learning and imitation models to motivate these benchmarks is that decision costs may imply that although agents have the required information to, for instance, derive best response strategies, they prefer not doing so because it is too costly. In fact, it can be assumed that agents try to find strategies, such as imitation rules, that save on cognitive effort and decision costs (Offerman, Potters, & Sonnemans, 2002). In particular, unlike Apestequia (2006), who considers mainly *best response dynamics* to explain individual behavior in both treatments, we use several learning and imitation theoretical models, which are based on contrasting assumptions about the level of rationality and the information available to subjects, namely: 'best response', 'imitate the average', mix of 'best response' and 'imitate the average', 'imitate the best' and 'follow the exemplary' learning rules.⁴

³ It can be noted that some important avenues for research on CPR experiments were pioneered by Elinor Ostrom, who was awarded the 2009 Nobel Prize in Economics 'for her analysis of economic governance, especially the commons'.

⁴ For details about: (i) the 'best response' rule see Huck, Normann, and Oechssler (1999), Huck, Normann, and Oechssler (2002), Offerman et al. (2002); (ii) the 'imitate the average' rule see Schlag (1999), Alós-Ferrer and Schlag (2007); (iii) the mixing of 'best response' and 'imitate the average' rules see Huck et al. (2002); (iv) the 'imitate the best' rule see Vega-Redondo (1997), Alós-Ferrer and Ania (2005); (v) the 'follow the exemplary' rule see Offerman et al. (2002).

Given the informational requirements of these rules, in treatment ‘incomplete’ agents could not apply ‘imitate the best’ or ‘follow the exemplary’, thus we would expect that in this treatment players use ‘best response’ or ‘imitate the average’ rules, or a mix of both. As we will show, the mix of ‘best response’ and ‘imitate the average’ leads towards an aggregate effort level converging globally to the NE of the constituent CPR game. Furthermore, since the cognitive efforts to apply ‘best response’ are relatively higher than those needed to apply ‘imitate the average’, we would also expect that the latter is employed more frequently than the former on average.

In treatment ‘complete’, players can also apply ‘best response’ and ‘imitate the average’ rules, or the mix of both, and besides, since subject-specific performance information from all subjects is now provided, players can also use ‘imitate the best’ and ‘follow the exemplary’. Hence, for this treatment we could expect that players can be influenced to some extent by the experience of others when deciding their effort levels. As shown by Vega-Redondo (1997), Alós-Ferrer and Ania (2005), if all subjects follow ‘imitate the best’, the unique stochastically stable state of the process is the OA outcome of the constituent CPR game, and, as proved by Offerman et al. (2002) if all subjects use the ‘follow the exemplary’ rule, the unique stochastically stable state of the process is the PE outcome of the constituent CPR game. Moreover, the ‘imitate the best’ rule can be used at minimal decision costs in comparison to ‘follow the exemplary’, since subjects receive the information processed in such a way that the comparison of payoffs across agents and the identification of the associated use level are straightforward. Hence, we would expect that the former is employed more frequently than the latter on average.

Finally, in order to econometrically test the hypotheses formulated from the theoretical predictions we use a random-effects model (see, e.g., Skrondal & Rabe-Hesketh, 2004), to assess the explanatory power of the different selected behavioral learning and imitation rules. One of the main reasons to use this technique is that each subject repeatedly interacts with the same group of players throughout the experiment, and therefore, in principle, we need to control within-subject dependence in effort decisions that may be introduced by unobserved between-subject heterogeneity. A random-effects model for analyzing subjects’ effort adjustment decisions is further justified given that the target of inference is the *population* of players and the interest is in analyzing potential subject *heterogeneity* in the employment of the learning and imitation rules.

The paper is structured as follows: Section 2 describes the experimental environment and the theoretical predictions are introduced in Section 3. In Section 4, we set forward the hypotheses of the work, and Section 5 presents the experimental results. Finally, a summary of conclusions is put forward.

2. The experimental environment

2.1. The constituent game⁵

Consider a community of $n \geq 2$ appropriators with shared access to a resource that is valuable but costly to extract, and let $I = \{1, \dots, n\}$ denote the set that indexes appropriators. Each appropriator $i \in I$ has an endowment y that can be invested either in the CPR or in an outside market. The marginal payoff obtained from the outside market is constant, exogenously given, and equal to w , while the payoff from investing in the CPR depends on the aggregate group effort and the share of individual effort on aggregate group effort. The appropriator i 's effort level in the CPR is denoted by x_i , where $0 \leq x_i \leq y$, and let X be the aggregate effort, defined as the sum of individual efforts, i.e. $X = \sum_{i=1}^n x_i$. Following the standard formulation of the dynamic of renewable resources, in which the group harvest is modeled as a parabola (see, e.g., Gordon, 1954), the aggregate group harvest is given by a concave function H which depends only on the aggregate group effort in the CPR,

$$H(X) := (v - kX)X, \quad (1)$$

with $v > 0$ and $k > 0$. The resource market is assumed competitive and the selling price is normalized to unity. Initially, an extra unit invested in the CPR pays more than investing in the outside market, and it is counterproductive for all individuals to invest all their endowment in the CPR. These assumptions are reflected by the conditions $H'(0) = v > w$, and $H'(ny) = v - 2nyk < 0$, respectively. As usually assumed, the part of the group harvest, obtained by each appropriator, is directly proportional to her share of effort in the group effort, i.e. $(x_i/X)H(X)$. Since the average harvest $H(X)/X$ is decreasing in X , it is clear that the individual harvest of any appropriator not only depends upon her effort, but also upon the effort introduced by the rest of the community.

Let $x = (x_1, \dots, x_n)$ be a vector profile of individual appropriators’ effort levels in the CPR. The payoff of an appropriator, $u_i(x)$, is given by

$$u_i(x) := w(y - x_i) + (x_i/X) H(X) = w(y - x_i) + x_i(v - kX), \quad i \in I. \quad (2)$$

Note from (2) that if appropriators invest all their endowment in the outside market by setting $x_i = 0$, they get a payoff of wy , whereas if they invest some of their endowment in the CPR, they get a payoff $w(y - x_i)$ from the outside market plus a payoff from the CPR that results from the sale of the amount of harvest obtained by the appropriator. From (2) it is also clear that an increase in the effort level by one appropriator decreases the share of harvest of *all* appropriators, thus imposing a negative externality on other individual payoffs.

⁵ Here we present an operationalized version of the classical static model of the commons that can be found in, inter alia: Dasgupta and Heal (1979), and Stevenson (1991).

Let us determine the conditions that characterize the group effort that maximizes the well-being of all the appropriators. Define $U(X)$ as the group payoff function, which is obtained summing individual payoffs $u_i(x)$ across all appropriators,

$$U(X) := w(ny - X) + H(X), \quad (3)$$

which is maximized subject to the constraints $0 \leq X \leq ny$. Given the stated assumptions on H , and denoting X^e as the PE group effort level, the unique solution of this problem is

$$-w + H'(X^e) = -w + (v - 2kX^e) = 0. \quad (4)$$

According to (4), at the PE group effort level the marginal return from the CPR should equal the foregone marginal payoffs from the outside market. Note that the solution does not depend on the endowment parameter y as long as it is sufficiently large. The PE aggregate group effort would be optimal from the point of view of the community if each individual were to invest an equal share of this effort. Nevertheless, self-motivated individuals would prefer to invest more if all appropriators choose the PE effort level, since this would increase their own individual payoffs.

In this context, the constituent CPR game can be formally described by $G = \{x_1, \dots, x_n; u_1, \dots, u_n\}$. It can be noted that G is a symmetric, non-cooperative static game with payoffs given by (2). Since this game is symmetric, there exists a symmetric NE at which each extractor invests x_i^* , $i \in I$ in the CPR, where

$$-w + H'(nx_i^*)/n + H(nx_i^*)(n-1)/n^2x_i^* = -w + (v - 2knx_i^*)/n + (v - knx_i^*)(n-1)/n = 0, \quad \forall i \in I. \quad (5)$$

As before, the endowment parameter does not affect the effort incentives if it is sufficiently large, and in the experiments it can be viewed as a fixed bonus to each subject for participation. Comparing the solutions of (5), $x_i^* = (v - w)/[(n + 1)k]$, and (4), $x_i^e = (v - w)/(2kn)$, $\forall i \in I$, we find that since each appropriator neglects the negative externality of her decisions on the payoffs of others, the outcome results in suboptimal extraction levels from the group perspective, i.e. $x^* > x^e$. This is the prediction of the classical model of the commons: *the level of extraction in a self-motivated, decentralized choice is inefficient since each appropriator could get higher payoffs if all were to limit their extraction, but no individual has the unilateral incentives to do so.*

Nonetheless, things are much worse in the case where the resource is characterized by *open access*, in which the number of appropriators increases into infinity and so individuals neglect the strategic interactions among them. However, if the number of appropriators is finite but they are poorly informed about and/or neglect the impact of their effort decisions on the average harvest (more precisely, if agents assume that $\partial[H(X)/X]/\partial x_i = 0$) then the model implies exactly the same behavior at the aggregate level. The latter assumption applied to the individual first-order condition of (2) characterizes the group OA outcome X^o ,

$$w = v - kX^o. \quad (6)$$

Eq. (6) yields the individual open access effort $x_i^o = (v - w)/(nk)$, $\forall i \in I$. This effort level is clearly higher than the individual NE outcome. Expression (6) states that at the OA outcome the average foregone payoffs from the outside market equal the average revenues from the CPR, thus dissipating all the rents from the CPR (for further details, see, e.g., Gordon, 1954; Stevenson, 1991).

2.2. Design and procedure⁶

The experiments conducted in this research follow the classical baseline design proposed in Walker et al. (1990), in whose study the problem of CPR exploitation is presented as an abstract investment decision. Eight identical subjects play the constituent CPR game repeatedly for 20 consecutive rounds. At the beginning of each round, subjects must solve the problem of distributing a set quantity of 30 tokens between two investment alternatives: a private market and a group market. The returns for each invested unit in the group market diminish as the total quantity invested by the group increases. This game characteristic seeks to represent the main problem of CPR exploitation by incorporating a negative externality into the group market. At the same time, the return per unit of investment in the private market introduces an opportunity cost for investment in the group market, which would represent the costs of exploiting the CPR. The specific parameters of the constituent game G used throughout our experiments correspond to those proposed in the experimental work of Casari and Plott (2003): consequently, we set $n = 8$, $y = 30$, $w = 5/2$, $v = 23/2$, and $k = 1/16$.

Two treatments, which differed in the information provided to subjects, were performed (Table 1 summarizes the experimental design). In treatment 'incomplete' each participant simultaneously makes a decision without knowledge about the individual decisions of the other subjects. After each round, participants were only informed about the group effort in the last round, and reminded about their own individual net payoffs and effort levels (this individualized information was not known by others, hence the name 'incomplete'). By contrast, treatment 'complete' was modified in relation to treatment 'incomplete' to enable agents to decide their effort levels based also on the experience of the population. Specifically, in addition to the information provided in treatment 'incomplete', at the beginning of each round participants were provided with subject-specific performance information in the earlier rounds, i.e. the complete profile of individual effort levels and associated payoffs.

⁶ For further details and an alternative analysis see: Salazar and Villena (2005).

Table 1
Summary of experimental design.

Treatment	Baseline information	Additional information
Incomplete	H, w_i, u_i, X	–
Complete	H, w_i, u_i, X	x_{-i}, u_{-i}

Note: information concerns the last round. Subscript $-i$ denotes all other subjects.

The experiments were performed in July 2004 with individuals recruited from a population of undergraduate students majoring in business and economics at the University of Concepción, Chile. Both treatments were conducted using the partner protocol, where each individual is confronted with the same group of players in each experiment. For each treatment we performed three sessions with a group of eight individuals in each, making a total of 48 subjects (six groups of eight individuals each) participating in the experiments. Recruited subjects were allocated randomly to each group. Communication among individuals was not allowed. Subjects did not know the exact number of rounds. However, it was common knowledge that experiments would end between rounds 15 and 25 and that the final round would be announced when two rounds remained.

Experiments followed the standard procedure of previous CPR experiments found in the economic literature (see, e.g., Plott & Smith, 2004). First, instructions were read aloud and explained with examples. Second, students were asked to respond to a set of exercises in order to verify their understanding of the game. Third, two test runs were performed to develop familiarity with the procedures and rules of the game. These test runs did not affect final gains. Forth, agents in both treatments were informed about the functional forms of the constituent CPR game G , and were told that per-round profits would be determined by average gains times the effort level by the agent, minus extraction costs and plus the fixed bonus. Finally, players knew how many competitors were in the game, although anonymity was guaranteed.

Participants were paid according to the payoffs obtained through the 20 rounds. The amount of money received by each participant expressed in Chilean pesos was six times the individual total gain achieved (where the ‘individual total gain achieved’ corresponds to the sum across rounds of individual payoffs given by (2) under the specific parameters of the experiment) and was paid privately in cash at the end of the experiment. Participants knew how total gains were converted to Chilean pesos before playing the experiment. Expressed in US dollars from July 2004,⁷ in treatment ‘incomplete’ the individual total gain varied from US\$ 15.89 to US\$ 20.92, with an average gain of US\$ 17.68, while in treatment ‘complete’ varied from US\$ 12.89 to US\$ 17.80, with an average gain of US\$ 15.54. Since the OA outcome or even higher investment levels were possible to emerge, the inclusion of the fixed bonus in each period (derived from the private market) served to ensure that participants would not be frustrated by low or negative payoffs. For treatments ‘incomplete’ and ‘complete’, experiments lasted on average 2 h and 2 h and 30 min, respectively, including the preliminary runs and the awarding of prizes.

3. Theoretical predictions

In this section we discuss simple theoretical learning and imitation models that are based on contrasting assumptions about the level of rationality and the information available to subjects. Moreover, these learning and imitation rules yield three very distinct outcomes: the PE, NE, and OA profiles. Before the discussion of the learning rules, we summarize the solution of the game under the proposed parameters.

Consider eight subjects that repeatedly play the static CPR game G . First, from (4) the individual PE outcome is obtained, which is denoted as x_i^e ,

$$x_i^e = (v - w)/(2nk) = 9, \quad \forall i \in I, \quad (7)$$

with an aggregate effort level $X^e = 72$ and individual payoffs $u_i^e = 115.5, \forall i \in I$. Second, from (5), the unique NE of the stage game, denoted as x_i^s , is obtained

$$x_i^s = (v - w)/[(n + 1)k] = 16, \quad \forall i \in I, \quad (8)$$

which yields an aggregate effort level $X^s = 128$ and individual payoffs $u_i^s = 91, \forall i \in I$. Finally, from (6) the individual OA outcome is obtained, which is denoted as x_i^o ,

$$x_i^o = (v - w)/(nk) = 18, \quad \forall i \in I, \quad (9)$$

with an aggregate effort level $X^o = 144$ and individual payoffs $u_i^o = 75, \forall i \in I$.

We compare the economic performance of the group effort outcomes observed in the experiments through an efficiency index ε , defined as group payoffs minus the group endowment money, normalized using the maximum possible group payoffs (Casari & Plott, 2003). Formally, we have:

$$\varepsilon(x) := [u(x) - nwy]/[u(x^e) - nwy], \quad (10)$$

⁷ The monthly average exchange rate at July 2004 was 632.28 Chilean pesos per US dollar.

Table 2
Benchmarks for the constituent CPR game.

Benchmarks	x_i	X	u_i	u	ε
Pareto efficient	9	72	115.5	924	100
Nash equilibrium	16	128	91	728	39.5
Open access	18	144	75	600	0

Note: The efficiency index is expressed in percents.

with $nwy = 600$. By definition, the efficiency index at the PE level is 100%, at the NE is 39.5%, and at the OA outcome is 0% since appropriators exploit the CPR up to a point at which the average foregone payoff equals the average revenue from the CPR. Table 2 and Fig. 1 summarize the theoretical benchmarks of the constituent game G.

Consider the standard prediction for the repeated CPR game obtained by backward induction: at each round subjects play the unique symmetric NE of the constituent game $x_i = 16 \forall i \in I$, which constitutes the unique *subgame-perfect Nash equilibrium*. In both treatments subjects had all the information needed to compute the NE. However, since the information processing can be too difficult or too costly for subjects in order to behave according to what prescribes the perfect rationality paradigm, we consider heuristics that subjects may adopt when confronted with informational and time constraints in an environment in which rationality is possibly bounded.

3.1. Best response learning

Consistently with our experimental design, let $\Gamma = \{0, 1, \dots, 30\}$ denote the choice set of effort levels. Formally, the 'best response' rule states that in round t subject i chooses her effort x_i^t as a myopic best reply to the group effort chosen by the other subjects in the previous round, denoted as X_{-i}^{t-1} . Thus, the 'best response' rule for subject $i \in I$ maps X_{-i}^{t-1} to the set

$$BR_i^t := \left\{ z \in \Gamma : u_i(z, X_{-i}^{t-1}) \geq u_i(z', X_{-i}^{t-1}), \quad \forall z' \in \Gamma \right\}. \quad (11)$$

Given the functional forms specified for the experiments and the discreteness and finiteness of the choice set Γ , we have

$$BR_i^t = \begin{cases} \{0\}, & \text{if } X_{-i}^{t-1} \geq 144 \\ \{30\}, & \text{if } X_{-i}^{t-1} < 84 \\ \left\{ 72 - \frac{1}{2} X_{-i}^{t-1} \right\}, & \text{if } 84 \leq X_{-i}^{t-1} < 144 \text{ and } X_{-i}^{t-1} \text{ is even} \\ \left\{ 72 - \frac{1}{2} X_{-i}^{t-1} - \frac{1}{2}, 72 - \frac{1}{2} X_{-i}^{t-1} + \frac{1}{2} \right\}, & \text{otherwise,} \end{cases} \quad (12)$$

where in the last case the effort level is chosen from BR_i^t according to some probability distribution with full support.

The 'best response' process defined in (12) yields a Markov chain over the state space Γ^8 , whose convergence to a stable equilibrium (the NE profile) cannot be assured globally.⁸ It is well known that the introduction of inertia⁹ stabilizes the best response dynamics, as shown by Huck et al. (1999), whose result can be readily applied to our linear CPR setting. Note that in our experimental design we did not introduce any inertia. Nevertheless, although the 'best response' process can be theoretically unstable, in some previous experimental work there is no empirical evidence of such instability (see, e.g., Huck et al., 2002).

To behave as a 'best response' learner, each subject needs to know the revenue and opportunity cost functions of the CPR market, the aggregate effort level of the remaining subjects from the last round, and of course how to compute a best response function. Hence, 'best response' is regarded as one of the most sophisticated forms of learning.

3.2. Imitate the average

The 'imitate the average' rule simply states that subjects imitate the average effort level of the other subjects observed in the last round. As indicated by Bosch-Domènech and Vriend (2003), there are two important reasons to consider this rule. First, it constitutes a reciprocating strategy in the sense of 'giving as good as you get' (equivalent to a *tit-for-tat* strategy). Second, as discussed by Huck et al. (1999), the rule seems reasonable because subjects may assume that while individual choices may be wrong, the 'average choice' cannot be 'too wrong'. Further reasons for applying this rule can be a preference for cautious behavior and a taste of conformity.

Neglecting the discreteness of the choice set Γ , the process is characterized by the set of difference equations $x_i^t = \frac{1}{7} X_{-i}^{t-1}$, $\forall i \in I$. Solving this system recursively we have

⁸ In order to see this, note that if the system under experimentation reaches one state of the absorbing set $a = \{(0, \dots, 0), (30, \dots, 30)\}$, then it will oscillate between both states and will never escape from a .

⁹ Inertia is introduced by assuming that in round t with independent probability θ subject i will stick to her previous effort x_i^{t-1} instead of following the best response. Huck et al. (1999) demonstrated that the resulting Markov process converges globally in finite time to the NE for any $\theta \in (0, 1)$.




Fig. 1. Basic benchmarks of the CPR game.

$$x_i^t = \frac{1}{8}(x_i^1 + X_{-i}^1) + \frac{1}{8}(7x_i^1 - X_{-i}^1)\left(-\frac{1}{7}\right)^{t-1}, \quad \forall i \in I. \tag{13}$$

It should be noted that the process in (13) converges to the average of all initial efforts, and that if all subjects follow this rule, the process is bounded above and below by the highest and lowest initial effort levels.

Even though subjects in treatment ‘incomplete’ were not provided with information about individual effort levels and payoffs of the other players, they could still apply some form of imitation (for instance, ‘imitate the average’) since they could observe aggregate, and therefore average, effort levels of the other subjects.

3.3. Mixed learning

The ‘best response’ and ‘imitate the average’ rules are based on opposite assumptions regarding the level of rationality and cognitive efforts needed to apply them. Hence, following Huck et al. (2002), it may be interesting to study some properties at the aggregate level of a learning framework in which subjects mix the ‘best response’ and ‘imitate the average’ learning rules.¹⁰

Let α be the weight given to the ‘best response’ rule and $1 - \alpha$ the weight given to the ‘imitate the average’ rule. Neglecting the discreteness of the choice set Γ , the difference equation that characterizes the mixed learning rule is

$$x_i^t = \alpha\left(72 - \frac{1}{2}X_{-i}^{t-1}\right) + (1 - \alpha)\frac{1}{7}X_{-i}^{t-1}, \quad \forall i \in I, \tag{14}$$

summing up across subjects, we obtain the process for aggregate investments

$$X^t = 576\alpha + \left(1 - \frac{9}{2}\alpha\right)X^{t-1}, \tag{15}$$

and solving this first-order difference equation yields the following proposition.

Proposition 1. *The equilibrium point of the dynamical process that characterizes the aggregate effort level under the mixed learning rule corresponds to the aggregate $NEX^* = 128$ if $\alpha \neq 0$. This equilibrium point is globally stable in the range $0 < \alpha < \frac{4}{9}$.*

Hence, at the aggregate level the mixing of the ‘best response’ with ‘imitate the average’ stabilizes the pure ‘best response’ process if the weight given to the ‘best response’ rule is positive, but not ‘too high’. Note that the mixed-learning process can be applied by subjects in both treatments.

3.4. Imitate the best

In ‘imitate the best’ learning, which is postulated by Vega-Redondo (1997), learning is driven by a mimicking dynamic: each subject identifies the highest payoff across all subjects in the last round and then selects her effort level from the set of effort levels that yielded the highest payoff. Formally, at every round t each subject i chooses her effort level x_i^t from the following set:

¹⁰ As stated by Huck et al. (2002), the mixed learning model can result more plausibly because some subjects are ‘best response’ learners and others are imitators. Empirically, this question is addressed formally in the analysis of experimental results at the individual level.

$$IB^t := \left\{ z \in \Gamma : \exists j \in I \text{ s.t. } x_j^{t-1} = z \text{ and } u_j^{t-1} \geq u_i^{t-1}, \forall i \in I \right\}, \quad (16)$$

according to some probability distribution with full support.

Note that any symmetric state $s = (z, \dots, z)$ is stable under the dynamics in (16). Hence, Vega-Redondo introduced some perturbation $\varepsilon > 0$ to the process, whose interpretation is that subjects either experiment with new strategies or make mistakes¹¹ from time to time. The following proposition follows directly from the theorem stated in Vega-Redondo (1997) in the context of a Cournot oligopoly and more generally from proposition 4 in Alós-Ferrer and Ania (2005), which we apply to our CPR setting:

Proposition 2. *If all subjects imitate the most successful one, the unique stochastically stable state of the process is the OA profile, namely $x^\varepsilon = (18, \dots, 18)$ as $\varepsilon \rightarrow 0$.*

Hence, if all players mimic the most successful subject, the system converges to the OA outcome of the constituent CPR game. The intuition is that if average returns of extraction are higher than foregone unit payoffs, the subject with the highest effort level will have the highest payoff and her decision will thus be imitated by other subjects up to the point where average returns equal foregone unit payoffs. From this point, continuing to increase efforts does not increase payoffs since average returns will be lower than foregone unit payoffs.

To ‘imitate the best’, subjects only need to know individualized information about efforts and payoffs from the last round. They are not required to know the payoff structure of the game, nor to have high rationality or calculation capabilities; they only need to be able to compare payoffs and mimic the most successful effort levels, making ‘imitate the best’ one of the simplest learning rules.

3.5. Follow the exemplary

The ‘follow the exemplary’ rule from Offerman et al. (2002) is based on imitation of the subject that sets the good example from the perspective of group payoffs. The exemplary subject (or one of the subjects) is defined as the one whose effort level in the last round would give the highest sum of payoffs if all subjects had followed it. Consider a learning rule in which the non-exemplary subjects imitate the effort level chosen by the exemplary subject, while the latter imitates the choice of the exemplary other subject, i.e. the subject that is exemplary if her own-exemplary effort level is disregarded.¹² Formally, at every round t each subject $i \in I$ chooses her use level x_i^t from the set

$$FE_i^t := \left\{ z \in \Gamma : \exists j \in I \text{ s.t. } x_j^{t-1} = z \text{ and } U(x_j^{t-1}, \dots, x_j^{t-1}) \geq U(x_k^{t-1}, \dots, x_k^{t-1}), \forall k \in I, j, k \neq i \right\}, \quad (17)$$

according to some probability distribution with full support.

Consider the introduction of some perturbation $\varepsilon > 0$ to the process, whose purpose, as in ‘imitate the best’, is to introduce some dynamics to the system. Then the following proposition follows from a result stated in Offerman et al. (2002) as applied to our CPR setting:

Proposition 3. *If every subject uses the ‘follow the exemplary’ rule, the unique stochastically stable state of the process is the PE profile, namely $x^\varepsilon = (9, \dots, 9)$ as $\varepsilon \rightarrow 0$.*

The intuition is the following: first, the non-exemplary subjects follow the exemplary effort level, while the exemplary subject will move towards a higher effort compared to the previous level. Then, the exemplary subject of the previous round will choose her original effort level, since now she follows. At this state, all subjects are exemplary, hence the resulting symmetric profile will be maintained until a subject makes a mistake or experiments a new effort level in the direction of the PE outcome.

In order to ‘follow the exemplary’, subjects need to know subject-specific effort levels, but not the payoffs. Subjects are not required to know the payoff structure of the game, although they must be able to determine the effort level from the last round that moves towards the PE outcome.

To sum up, Table 3 presents a summary of the learning and imitation rules, the informational requirements and the associated theoretical predictions.

4. Hypotheses

In treatment ‘incomplete’ subjects had the information available to apply a learning rule that yields outcomes towards the NE of the constituent CPR game G , while treatment ‘complete’ may lead towards the OA outcome through imitation of the most successful subjects. Let us elaborate on these arguments.

¹¹ States that are in the support of the limit invariant distribution of the process for $\varepsilon \rightarrow 0$ are known as *stochastically stable states*. For a textbook treatment, see, e.g., Fudenberg and Levine (1999).

¹² There are variations to this rule that differ in the way the exemplary subject conditions his choice based on those of the non-exemplary ones, which were used in the literature of experimental economics in a Cournot oligopoly setting (see, e.g., Offerman et al., 2002).

Table 3
Summary of learning and imitation rules.

Learning and imitation rules	Required information	Prediction (x_i, X)
Best response and imitate the average	X_{-i}, BR_i	(16.128)
Imitate the best	x_{-i}, u_{-i}	(18.144)
Follow the exemplary	x_{-i}	(9.72)

Note: Information concerns the last round. Subscript $-i$ denotes all other subjects.

First, recall that in treatment ‘incomplete’ each subject is informed about her individual effort level and payoffs, and of the group effort level in the last round. Hence, in principle subjects may behave according to what prescribes the ‘best response’ learning rule since they have the information to do so. Moreover, it is possible in treatment ‘incomplete’ that to some extent subjects imitate through ‘imitate the average’. Since subjects cannot ‘imitate the best’ or ‘follow the exemplary’ because they lack the subject-specific information needed to do so, we can expect that treatment ‘incomplete’ will involve some mix of ‘best response’ and ‘imitate the average’. As Proposition 1 states, the mix of ‘best response’ and ‘imitate the average’ leads towards an aggregate effort level that converges globally to the NE of the constituent CPR game. Thus, we can expect that treatment ‘incomplete’ yield outcomes close to the NE through some mix of ‘best response’ and ‘imitate the average’. Furthermore, since the cognitive efforts to apply the ‘best response’ rule are relatively higher than those needed to apply ‘imitate the average’, we can expect that on average the latter is employed more frequently than the former.¹³

Second, treatment ‘complete’ can be conducive to the use of the same rules discussed above, plus ‘imitate the best’, and ‘follow the exemplary’ rules, since subject-specific performance information from all subjects is provided. In this context, it is natural to postulate that subjects are influenced to some extent by the experience of others when deciding their effort levels. Although both imitation rules can be applied in treatment ‘complete’, ‘imitate the best’ can be used at minimal decisions costs in comparison to ‘follow the exemplary’, since subjects receive the information processed in such a way that the comparison of payoffs across subjects and the identification of the associated effort levels are straightforward. Further arguments pointing to the use of ‘imitate the best’, instead of ‘follow the exemplary’, are the potential preferences of players to ‘beat the opponent’ or to ‘behave spitefully’ (see, e.g., Bosch-Domènech & Vriend, 2003; Hamilton, 1970; Levine, 1998) and that agents are boundedly rational (see, e.g., Vega-Redondo, 1997). As Propositions 2 and 3 showed, if all subjects follow ‘imitate the best’, the unique stochastically stable state of the process is the OA outcome, and if all subjects use the ‘follow the exemplary’ rule, the unique stochastically stable state is the PE outcome. Consequently, we can expect outcomes ranging from NE to OA in treatment ‘complete’.

Hence, given these conjectures, we formulate the following hypotheses:

Hypothesis 1. In treatment ‘complete’ the average aggregate effort level is systematically different than the corresponding level in treatment ‘incomplete’.

Hypothesis 2. In treatment ‘incomplete’ the average aggregate effort level systematically approaches the NE outcome of the CPR game, while in treatment ‘complete’ the average aggregate effort approaches some level between the NE and the OA outcomes, such level being systematically different than these two outcomes.

It should be noted that Hypothesis 1 suggests that there is a systematic difference between the experimental treatments at the average group effort level that cannot be predicted by standard non-cooperative game theory arguments. Indeed, since the experimental game is a finitely repeated one with complete and imperfect information, the unique subgame-perfect Nash equilibrium in both treatments is the NE of the constituent CPR game played in each round. Thus, we postulate that, unlike Apesteguia (2006), aggregate behavior is significantly different between both informational treatments, whenever subject-specific performance information is at stake instead of payoff structure information. Moreover, Hypothesis 2 is a refinement of Hypothesis 1, since it posits that the average group effort level in treatment ‘incomplete’ is consistent with the NE, while the corresponding effort level for treatment ‘complete’, although higher than the NE, is systematically lower than the OA outcome. This follows since we expect that, in spite of the fact that ‘imitate the best’ can be an important driver of behavior, subjects can also be influenced by ‘best response’ and ‘imitate the average’, and so a certain balance between the main states of attraction, namely NE and OA outcomes, can be expected on average.

Hypothesis 1 and 2 were constructed from Propositions 1–3, which are in turn based on theoretical models at the individual level. Hence, we now formulate a hypothesis at the individual level, consistent with Propositions 1–3, regarding the explanatory power of the learning and imitation rules considered in this work:

¹³ Moreover, note from Proposition 1 that this is a necessary condition for convergence to the NE.

Hypothesis 3. In treatment ‘incomplete’ subjects on average mix ‘best response’ with ‘imitate the average’, while in treatment ‘complete’ subjects on average mix these two rules with ‘imitate the best’.

5. Experimental results

We now turn to the experimental analysis of the postulated hypotheses. This section is organized as follows: First, according to the qualitative hypotheses, we analyze the data at the aggregate level, in order to assess how the different informational structures affect aggregate use behavior of the resource. This will provide indirect evidence for the assessment of imitation in addition to individual learning when information about others’ experience is available. Nevertheless, a more detailed study of the prediction performance of the different learning theories must be based on analysis of the individual data. Henceforth, we formally study individual data by estimating a random-effects model to take into account potential heterogeneity in the prediction power of learning rules.

5.1. Aggregate behavior analysis

The average group effort levels for the 20 rounds in our experiments are presented in Fig. 2. Here it can clearly be seen that the average group effort levels in treatment ‘complete’ are systematically higher than those in treatment ‘incomplete’ in almost every round. Table 4 reports summary statistics for aggregate behavior in both treatments, namely average aggregate effort, average efficiency index, and average Gini coefficient over the entire experiment (20 rounds) and the last 10 rounds, and across all three groups (sessions) in each treatment. The corresponding standard deviations across rounds are also reported. First, note that the average group effort in treatment ‘incomplete’ (Mean20) is below the NE outcome, while the corresponding effort considering the last 10 rounds (Mean10) is interestingly close. Second, note that in treatment ‘complete’ (Mean20 and Mean10) the average aggregate effort is at a point somewhere between the NE and OA outcomes.

Moreover, Table 4 suggests that the average efficiency in treatment ‘complete’ (Effy20 and Effy10) is lower than in treatment ‘incomplete’, which is remarkably close to the NE efficiency when considering the last 10 rounds (Effy10). This result follows directly from the seemingly higher average group effort in treatment ‘complete’. The latter suggests that when information about individual efforts and payoffs is available to subjects, the resource is exploited at an aggregate level higher than in treatment ‘incomplete’, thus yielding an even lower efficiency index. Interestingly, the average Gini index seems to be higher in treatment ‘incomplete’ (Gini20 and Gini10) than in treatment ‘complete’, suggesting that in terms of efforts and payoffs, behavior is more symmetric when subject-specific information is provided. These results suggest that imitation of the best could significantly explain subjects’ behavior in treatment ‘complete’. We formalize our findings below.

Regarding the aggregate analysis, it is customary in the experimental literature (see, e.g., Apesteguia, 2006; Huck et al., 1999) to consider one observation per group (the average) in order to obtain a sample of independent observations. In our setting, this procedure yields a sample of only three observations per treatment. Hence, in order to obtain statistical




Fig. 2. Average group effort by treatment.

Table 4

Summary of group statistics by treatment.

Treatment	Mean20	Mean10	Effy20	Effy10	Gini20	Gini10
Incomplete	123.55 (14.36)	127.80 (12.51)	0.448 (0.28)	0.370 (0.25)	0.207 (0.06)	0.213 (0.06)
Complete	134.98 (18.56)	137.20 (12.21)	0.169 (0.44)	0.152 (0.32)	0.119 (0.05)	0.133 (0.04)

Note: Standard deviations in parentheses.

Table 5

One-sample tests for theoretical benchmarks (two-tailed), 20 rounds.

Treatment Null hypothesis	Incomplete		Complete	
	Table-lookup's <i>t</i>	Likelihood-ratio	Table-lookup's <i>t</i>	Likelihood-ratio
Social optimum	25.24***	11.33***	34.76***	13.71***
Nash equilibrium	−2.18	2.55	3.85*	5.23**
Midpoint	−6.10**	7.54***	−0.56	0.25
Open access	−10.01**(+)	9.52***	−4.98**	7.17***

* Significant at the 2% level.

* Significance at the 10% level.

** Significance at the 5% level.

*** Significance at the 1% level.

significance, we averaged observations across groups (sessions) by round for each treatment, yielding a time-series sample of 20 observations per treatment. However, non-independence within samples is expected given the theoretical discussion in Section 3. Consequently, to formalize the comparison between treatments we applied the one-sample tests of the mean proposed by *Zwiers and von Storch (1995)*, which account for serial correlation and are well suited for small samples: (i) the likelihood-ratio (henceforth LR) test under an AR(1) model, and (ii) the table-lookup (henceforth TL) test, which offers protection against spurious reject decisions by superior small sample properties.¹⁴ These tests were run, for each treatment, on data from the entire experiment (20 rounds) and the last 10 rounds.

Table 4 suggests that the mean group effort is higher when subject-specific performance information is provided. According to (two-tailed) LR tests¹⁵, the mean group average effort in treatment 'complete' is significantly different than in treatment 'incomplete' at the 0.1% level for the entire experiment and at the 5% level for the last 10 rounds. Moreover, the mean group average effort in treatment 'complete' is higher than in treatment 'incomplete' since (one-sided) TL tests are significant at the 1% level for the entire experiment, and almost at the 10% level for the last 10 rounds.¹⁶ Based on our findings, we can state that the provision of information about individual efforts and payoffs of others significantly increases average group effort levels compared to the case in which only the individual's own efforts and payoffs and group use levels are known to subjects.

The hypothesis stating that there is no significant difference between the mean group average effort in treatment 'incomplete' and the group NE cannot be rejected at the 10% level for the entire experiment, and remarkably, even at the 90% level for the last 10 rounds according to (two-sided) LR tests. These results are supported by TL tests at the 20% level. Regarding treatment 'complete', as suggested by *Table 4* one cannot reject the hypothesis that the mean group effort equals the *midpoint* between the NE and OA outcomes according to (two-sided) LR tests, even at the 60% level for the entire experiment and last 10 rounds. Again, these results are supported by (two-sided) TL tests at the 20% level. Since the data was tested against theoretical benchmarks, these findings complement those presented above and provide indirect evidence for the learning theories used in developing the statements in *Hypothesis 1 and 2* (in particular, for the employment of 'mixed learning' in treatment 'incomplete', and along with 'imitate the best' rule in treatment 'complete'). Results for (two-sided) one-sample hypothesis testing of the PE, NE, midpoint, and OA outcomes for the entire experiment are summarized in *Table 5*.¹⁷

The statistical analysis presented above, which aims to develop inferences at the treatment level, can be extended to the session level. Since sessions are independent, the three time-series per treatment are also independent. *Table 6* presents a summary of average group effort levels per session. For the entire experiment, leaving aside session 1 (or alternatively session 6), it can be noted that the average group effort point-estimates by session are lower in treatment 'incomplete' than in treatment 'complete'. Likewise, the average group effort point-estimates for the last 10 rounds are systematically lower in treatment 'incomplete' compared with treatment 'complete'.

Table 7 summarizes results for (two-sided) one-sample LR hypothesis testing of the PE, NE, midpoint, and OA outcomes by session.¹⁸ Consider first the sample of 20 rounds. In treatment 'incomplete' the mean group effort equals the NE in sessions 1 and 2 at the 30% and 10% levels respectively, while in session 3 differs at the 5% level. Regarding treatment 'complete', in session

¹⁴ The Student's *t*-test is commonly used with time-series data coming from experiments. However, it is well known that the effect of serial correlation in this test is to make comparisons of means liberal (i.e. the actual significance level is smaller than the specified level). For details, see, e.g., *Zwiers and von Storch (1995)*.

¹⁵ From a statistical point of view paired tests are adequate in our context since the comparison between treatments is based on *common* base times (i.e. rounds). Thus, we use the one-sample version of the LR and TL tests.

¹⁶ More precisely, the *t* statistic equals 2.32 and the (linearly interpolated) appropriate critical value for the TL tests equals 2.43 for a one-sided test conducted at the 10% significance level.

¹⁷ The tests used here are based on the assumption of AR(1) processes. Skewness–kurtosis and Shapiro–Francia tests support the hypothesis of Gaussian errors for the samples of 20 rounds (last 10 rounds) at the 70% level (80%), 90% level (15%), and 45% level (50%) for the paired, 'incomplete', and 'complete' time-series respectively. The point-estimates of the AR coefficients are −0.164, 0.305, and 0.193 (20 rounds) and −0.208, 0.004, and 0.186 (last 10 rounds) respectively.

¹⁸ For simplicity, we report results only for the LR test. The Gaussian assumption cannot be rejected for sessions 1–5 according to skewness–kurtosis and Shapiro–Francia tests at least at the 35% level for the 20 rounds and last 10 rounds. In session 6 the hypothesis is accepted (rejected) at the 10% (1%) level for the last 10 rounds (20 rounds). In this case the LR test was conducted using robust standard errors. The AR coefficients are for the entire experiment (last 10 rounds) 0.325 (0.185), 0.543 (−0.052), 0.176 (0.179), −0.041 (−0.588), 0.061 (0.441), and 0.689 (−0.344) for sessions 1–6, respectively.

Table 6

Summary of group effort decisions by session.

Treatment	Incomplete			Complete		
	1	2	3	4	5	6
Complete experiment	131.95 (12.49)	118.85 (13.83)	119.85 (13.45)	136.25 (11.75)	145.2 (17.26)	123.5 (19.59)
Last 10 rounds	132.50 (10.92)	128.3 (8.38)	122.6 (16.13)	133.5 (10.62)	141.6 (15.92)	136.5 (8.76)

Note: Standard deviations in parentheses.

Table 7

One-sample LR tests for theoretical benchmarks by session (two-tailed).

Treatment	Incomplete			Complete		
	1	2	3	4	5	6
Null hypothesis	Complete experiment (sample size 20)					
Social optimum	10.62***	6.75***	11.68***	11.68***	13.63***	2.22
Nash equilibrium	0.82	2.26	4.08**	5.96**(+)	7.73***	0.30
Midpoint	1.22	4.31**	8.15***	0.01	3.86**	0.89
Open access	5.36**	5.52**(+)	10.18***	6.17**(+)	0.10	1.44
Null hypothesis	Last 10 rounds (sample size 10)					
Social optimum	6.17**(+)	11.34***	6.44**(+)	19.56***	3.94**	14.96***
Nash equilibrium	0.35	0.01	0.97	6.17**(+)	0.44	7.81***
Midpoint	1.10	4.67**	3.40*	3.23*	0.00	0.01
Open access	3.96**	7.70***	5.00**	11.63***	0.60	6.50**(+)

+ Significant at the 2% level.

* Significance at the 10% level.

** Significance at the 5% level.

*** Significance at the 1% level.

4 the mean group equals the midpoint even at the 90% level, while in session 5 differs at the 5% level but equals the OA outcome at the 70% level. Moreover, in session 6 there is insufficient evidence to reject any of the theoretical benchmarks at standard levels. Taking into account the last 10 rounds, the mean group effort equals the NE in each session of treatment 'incomplete' at least at the 30% level. In treatment 'complete' the mean differs from the midpoint in session 4 at the 10% level but not at the 5% level, while differs from the NE at the 2% level. This suggests that the mean group effort in session 4 is closer to the midpoint benchmark. In session 6 the midpoint outcome cannot be rejected even at the 90% level. Finally, there is no evidence to reject the NE, midpoint, nor the OA outcomes at standard levels in session 5. By comparing the statistics in Table 7 at the group NE and midpoint outcomes it can be noted that the results at the session level generically agree with those at the treatment level. Moreover, they also suggest that within-session heterogeneity in mean group effort levels could be significant (e.g., consider the results from session 5 for the last 10 rounds).¹⁹

Consistent with the findings about the mean group effort levels at the treatment level, the mean efficiency index differs between treatments since the (two-sided) LR statistic is significant at the 0.1% and 5% levels for the entire experiment and last 10 rounds, respectively. Likewise, the appropriate (one-tailed) TL statistic is significant at the 1% level for the 20 rounds, and almost at the 10% level for the last 10 rounds.²⁰ Consequently, the *detriment in economic efficiency* because of the higher level of exploitation in treatment 'complete' is substantial. In addition, efforts and payoffs are *more symmetric* across subjects in treatment 'complete' since the (two-sided) LR test for the mean difference in Gini coefficients is significant at the 1% level for the 20 and last 10 rounds, while the (one-sided) TL statistic is significant at the 1% and 5% levels for the 20 and last 10 rounds, respectively. These results together with those about the ordering of the mean average group effort across treatments strongly suggest that apparently subjects tried to match the highest efforts (note that in the CPR game the subject that sets the highest effort obtains the highest payoff if the group effort is below the OA outcome) when the information needed to do so is available, even if that behavior resulted in lower individual payoffs.²¹

The main findings can be summarized in the following three statements:

Result 1. *More information about individual efforts and profits of others yields higher average aggregate effort levels, i.e., average aggregate effort levels are systematically lower in treatment 'incomplete' than in 'complete', which is consistent with Hypothesis 1.*

¹⁹ It should be noted that comparisons at the treatment level consider the mean of group average effort levels, while at the session level these take into account the mean of group effort levels. Since we are interested in assessing the main impact of subject-specific performance information, in this work we deal mostly with inferences about mean outcomes at the treatment level.

²⁰ The critical value for a one-sided test at the 10% level equals 2.61, while the obtained *t*-statistic equals 2.32.

²¹ Both skewness–kurtosis and Shapiro–Francia tests support the assumption of Gaussian residuals for the efficiency paired time-series and Gini index paired time-series at the 70% and 50% level respectively, for both the 20 rounds and last 10 rounds. The point-estimates of the AR coefficients for the efficiency time-series are –0.096 and –0.146, and for the Gini index time-series are 0.287 and 0.132, for the 20 rounds and last 10 rounds respectively.

Result 2. In treatment ‘incomplete’ the average group effort level matches the NE, while in treatment ‘complete’ the corresponding effort level matches the midpoint between the NE and OA outcomes, which is consistent with Hypothesis 2.

Result 3. A more equal distribution of efforts and payoffs arises when information about individual efforts and profits of others is available, a finding that, in conjunction with Results 1 and 2, suggests that subjects do imitate successful subjects to some extent.

The theoretical predictions for ‘imitate the best’ and ‘follow the exemplary’ learning rules are based on the notion of stochastic stability, and hence, they are valid only for the long run. For an experiment lasting 20 rounds, one should not attempt to observe a stochastically stable state, as consistent by the findings of average Gini coefficients higher than zero in both treatments.²² Moreover, we found that the average aggregate effort in treatment ‘complete’ is significantly lower than the OA outcome, suggesting that if subjects to some extent employ imitation of the successful subject, the latter may not be the only learning mechanism employed by subjects. However, at the aggregate level, our findings qualitatively support the theoretical results for ‘imitate the best’, since a change in the aggregate effort towards the OA outcome was found. Note that at the aggregate level our results also support the theoretical predictions for the ‘mixed learning’ rule. Whether these findings are valid at the individual level is a question that will be addressed in the following subsection.

As can be expected after the discussion above, there have not been successful attempts to establish the PE effort outcome in any group (session). In the 20 rounds for all sessions and treatments, the group effort level was observed below the midpoint between the group NE and the group PE outcomes in only five out of 120 rounds. For the 480 individual effort observations for all sessions in a treatment, in treatment ‘incomplete’ 10.63% were less or equal the individual PE level, while the corresponding percentage in treatment ‘complete’ was only 4.58%. The latter suggests that cooperators were rapidly exploited by other subjects, discouraging any attempt to establish the PE aggregate effort level.

This observation is summarized in the following result.

Result 4. Despite some attempts to reach the PE effort level, these were exploited by other subjects, even in treatment ‘complete’ where the information required in order to play the ‘follow the exemplary’ rule was readily provided.

Since experimental data on human subjects usually have an important noise component, which is consistent with heterogeneous individuals whose presence is a common situation in the experimental literature (Van Winden, Van Dijk, & Sonnemans, 1998), Results 1–3 are quite remarkable. The analysis at the aggregate level indirectly provided interesting insights that to some extent shed light on the explanatory performance of the proposed learning and imitation theories used in the construction of Hypothesis 1 and 2. However, a closer analysis of the effort decisions is needed to develop an explanation of why the provision of subject-specific performance information yields higher aggregate resource extraction levels. This question is addressed by the estimation of a regression model at the individual level.

5.2. Individual behavior

Given that the target of inference is the population of subjects and the interest is in analyzing potential subject heterogeneity in the employment of the learning rules, we estimated a random-effects model for analyzing subjects’ effort adjustment decisions (see, e.g., Huck et al., 1999).

We define the adjustment predicted for subject i in round t by rule r as

$$y_{i,r}^t = a_{i,r}^t - x_i^{t-1}, \quad (18)$$

where $a_{i,r}^t$ denotes the effort level predicted for subject i in round t by learning rule r , which takes values in {‘best response’, ‘imitate the average’, ‘imitate the best’, ‘follow the exemplary’}. Let $\bar{y}_{i,r}$ be the subject-mean for subject i by learning rule r across rounds,

$$\bar{y}_{i,r} = \frac{1}{19} \sum_{t=2}^{20} y_{i,r}^t. \quad (19)$$

The purpose of these definitions is clarified in what follows.

In general, in cross-sectional time-series (panel) data, there are two types of effects: (1) time-series, or within-subject effects, which measure the expected change in the response of subject i , if some i ’s covariate x is increased by 1; and (2) cross-sectional, or between effects, which measures the expected difference between the responses of subjects i and j if they differ in some covariate x by 1. Standard random-effects models assume that within and between effects are equal. If this is true, these two effects can be pooled to obtain a more efficient common estimator. However, the estimated within effect may differ from the estimated between effect due to omitted subject-specific explanatory variables (for example, gender). In this case, the within effect estimator is inconsistent (see e.g., Skrondal & Rabe-Hesketh, 2004). To address this problem, we follow Mundlak (1978), and include subject-means (19) in the model. This addition allows differences in the estimators of within and between effects, which ensure consistent estimation of within effects, but do not ensure consistent estimation of

²² Formally, from (one-sided) TL tests the null hypothesis in which the mean Gini coefficient equals zero is rejected at the 1% level of statistical significance.

between effects (see, e.g., Skrondal & Rabe-Hesketh, 2004). Nevertheless, following the experimental literature, we only concerned with the within effect estimations of the examined rules (see, e.g., Huck et al., 1999).²³ Therefore, when appropriate, we set within and between effects as equal, in order to get a more efficient estimator; otherwise we allow for different effects. The specific procedure is resumed below.

Consider the following general level-1 model which accounts for different within and between effects of the covariates,

$$\begin{aligned} x_i^t - x_i^{t-1} = & \beta_{0i} + \beta_{1i}^1 (y_{i,BR}^t - \bar{y}_{i,BR}) + \beta_{2i}^2 \bar{y}_{i,BR} + \beta_{2i}^1 (y_{i,IA}^t - \bar{y}_{i,IA}) + \beta_{2i}^2 \bar{y}_{i,IA} \\ & + \beta_{3i}^1 (y_{i,IB}^t - \bar{y}_{i,IB}) + \beta_{3i}^2 \bar{y}_{i,IB} + \beta_{4i}^1 (y_{i,FE}^t - \bar{y}_{i,FE}) + \beta_{4i}^2 \bar{y}_{i,FE} + \varepsilon_i^t, \end{aligned} \quad (20)$$

where the intercept β_{0i} and slopes $\beta_{(\cdot)i}$ are subject-specific coefficients. The coefficient β_{ri}^1 of the mean-centered covariate $(y_{i,r}^t - \bar{y}_{i,r})$ and the coefficient β_{ri}^2 of the subject-mean $\bar{y}_{i,r}$ are the within and between effects of learning rule r , respectively. The level-2 model, which accounts for potential heterogeneity across subjects in the employment of learning rules, has the intercept and coefficients of within effects of covariates as responses,

$$\beta_{0i} = \gamma_{00} + \zeta_{0i}, \beta_{1i}^1 = \gamma_{11} + \zeta_{1i}, \beta_{2i}^1 = \gamma_{22} + \zeta_{2i}, \beta_{3i}^1 = \gamma_{33} + \zeta_{3i}, \text{ and } \beta_{4i}^1 = \gamma_{44} + \zeta_{4i}, \quad (21)$$

where the disturbances $\zeta_{(\cdot)i}$ have a multivariate normal distribution with zero mean and covariance matrix ψ given the covariates, and ε_i^t is a normally distributed residual term with zero mean and subject-heteroskedastic variance θ_i . The coefficients which were determined separately for each treatment are the γ_{00} , γ_{11} , γ_{22} , γ_{33} , γ_{44} and β_{1i}^2 , β_{2i}^2 , β_{3i}^2 , β_{4i}^2 , along with the elements of the covariance matrix ψ and the subject-heteroskedastic variances θ_i .

Different within and between effects were assumed in (20) since Hausman tests indicated misspecification of the standard random-effects models²⁴ (p -value < 0.000). Therefore, following Mundlak (1978) we assumed equal within and between effects for learning rules whose within and between effects were equal at the 10% level in (20). Given this yardstick, we allowed different effects for 'best response' and 'imitate the average' in treatment 'incomplete'; and for 'imitate the average', 'imitate the best', and 'follow the exemplary' in treatment 'complete'; in the latter treatment we imposed the restriction $\beta_{1i}^1 = \beta_{2i}^1$ for the 'best response' rule. Moreover, a variance-inflation factor (henceforth VIF) and condition number (henceforth CN) analysis revealed multicollinearity problems in treatment 'complete', due to high correlations among subject-means.²⁵ Since these variables are of no interest for our purposes beyond having a correctly specified model, we removed the subject-means of 'imitate the best' and 'follow the exemplary' by setting $\beta_{3i}^2 = \beta_{4i}^2 = 0$.²⁶ After allowing different within and between effects for the appropriate learning rules, Hausman tests were no longer significant even at the 90% level.

In treatment 'incomplete' we imposed the restriction $\beta_{3i}^1 = \beta_{3i}^2 = \beta_{4i}^1 = \beta_{4i}^2 = 0$ since subjects had no information to 'imitate the best' or 'follow the exemplary'. Indeed, after estimating (20), (21) including these covariates, we found that they were not significant at any standard statistical levels. Besides, in treatment 'complete' we found some degree of collinearity between the mean-centered covariates 'imitate the average' and 'following the exemplary'. Nevertheless, the mean VIF and CN were not large enough to be of concern.²⁷ Hence, we kept both covariates in order to avoid potential specification errors.²⁸

We estimated the full model by maximum-likelihood²⁹ and progressively obtained a more compact model regarding the random-effects through LR tests.³⁰ The random-slope variance was not significant for 'best response' at the 15% level in treatment 'incomplete' and even at the 40% level in treatment 'complete'. It was also not significant for 'imitate the best' at the 10% level and for 'follow the exemplary' at the 40% level. Moreover, the random-intercept variance in both treatments was not significant, even at the 40% level. Nevertheless, the random-slope for 'imitate the average' was highly significant in both treatments (p -value < 0.000 versus an ordinary linear-regression model with subject-heteroskedastic residuals). For the sake of simplicity, we present only the estimated coefficients for this reduced model, which are summarized in Table 8.³¹

The coefficients (within effects) yield an indication of the relative importance of the learning rules in explaining subjects' adjustments. In treatment 'incomplete' the 'imitate the average' and 'best response' rules emerge as important variables. All

²³ Huck et al. (1999) performed a regression at the individual level including subject dummies for intercepts and slopes to control for individual differences in learning behavior. Thus, their formulation is a fixed-effects model that provides estimates for within-subject effects of covariates.

²⁴ Note that the standard random-effects model results from letting $\beta_{1r}^1 = \beta_{1r}^2$ for all rules r in (20).

²⁵ The VIF's of the subject-means were 18.94, 23.57, and 11.76, yielding a mean VIF of 9.74. The CN was 21.74, which is quite high (a common rule of thumb is that multicollinearity is of concern if the mean VIF and CN are equal or greater than 10 and 15 respectively). Pearson correlations among subject-means were above 0.90.

²⁶ After removing these covariates the mean VIF and CN fell to 2.78 and 4.77 respectively. Alternatively, we could have removed any other pair of subject-means. The results of the estimation appear almost the same.

²⁷ See footnote 26. The VIF's of the mean-centered covariates for 'imitate the average' and 'follow the exemplary' are equal to 4.92 and 4.63 respectively. The Pearson correlation of these variables equals 0.88.

²⁸ Nevertheless, the results remain almost the same after removing the mean-centered 'follow the exemplary', except that the coefficient of 'imitate the average' (mean-centered) drops to 0.580, as expected from results in Table 8.

²⁹ The model was estimated using GLLAMM (see Skrondal & Rabe-Hesketh, 2004).

³⁰ Here the "naive" p -values are already corrected by division by two (see, e.g., Skrondal & Rabe-Hesketh, 2004).

³¹ In both treatments subject-heteroskedastic level-1 variance (whose estimates are omitted for simplicity) was significant according to LR tests (p -value < 0.000). Both skewness–kurtosis and Shapiro–Francia statistics are significant at the 5% level in treatment 'incomplete', while almost at 10% and 5% respectively in treatment 'complete'. Hence, we use robust standard errors based on the Huber–White (sandwich) estimator. Cross-sectional dependence was not significant even at 10% and 20% in treatments 'incomplete' and 'complete' according to Pesaran's and Friedman's statistics (De Hoyos & Sarafidis, 2006). Serial correlation was not significant even at 70% and 20% levels in treatments 'incomplete' and 'complete' according to the Wooldridge test.

Table 8
Maximum likelihood estimates for experimental data, robust standard errors.

Treatment Parameter	Incomplete		Complete	
	Coefficient	Std. error	Coefficient	Std. error
<i>Fixed part</i>				
γ_{11} [response]	0.159***	(0.024)	0.058***	(0.007)
γ_{22} [average]	0.489***	(0.062)	0.645***	(0.087)
γ_{33} [best]	NE	NE	0.128***	(0.031)
γ_{44} [exemplary]	NE	NE	-0.075**	(0.038)
γ_{00} [cons]	0.030	(0.077)	0.310***	(0.025)
β_{1i}^2 [mn response]	0.018	(0.021)	NE	NE
β_{2i}^2 [mn average]	-0.041**	(0.018)	-0.056***	(0.021)
<i>Random part</i>				
$\sqrt{\psi_{22}}$ [average]	0.237***		0.302***	
Observations	456		456	
# Subjects	24		24	
Obs per subject	19		19	
Log-likelihood	-1179.426		-1151.941	

Note: *Significance at the 10% level.

** Significance at the 5% level.

*** Significance at the 1% level.

within effects have the expected sign. Interestingly, 'imitate the average' seems to be employed on average more frequently by subjects since (one-tailed) Wald tests indicates that the null hypothesis, in which both coefficients are equal is rejected, at any standard significance level (p -value < 0.000). Since the same information is needed to apply these rules, this finding may be explained by the differential in cognitive efforts.³² Thus, subjects on average adjust efforts mixedly as *best response learners* who are driven by absolute payoffs and as *imitators* of what others do. These findings provide support to the mixed-learning process that converges theoretically to the NE as stated in Proposition 1, and is consistent with Result 2 of the aggregate behavior analysis.

In treatment 'complete' the 'best response', 'imitate the average' and 'imitate the best' learning rules are significant at any standard levels (p -value < 0.000), while the 'following the exemplary' rule is significant at the 5% level.³³ Regarding the importance of explaining subjects' adjustments, (one-tailed) Wald tests indicate that the within effect of 'imitate the average' is significantly higher than the corresponding effect of 'imitate the best', at any standard levels (p -value < 0.000) while the within effect of 'imitate the best' is higher than the effect of 'best response' at the 2% level. As in treatment 'incomplete', this may be explained by the difference in cognitive efforts needed to apply these learning rules. Although the estimated coefficient of 'following the exemplary' is considerably less significant than the other within effects, its negative value may provide evidence that adjustment is on average against the rule.

Consequently, in treatment 'complete' subjects on average seem to *complement* imitation of the average and belief learning with imitation of the subject that obtained the highest payoffs when deciding adjustments. Hence, regarding the 'imitate the best' component, there are two basic alternative explanations for our findings: (1) that subjects do imitate on average the most successful one when information to do so is available, and (2) that subjects adjust, trying to maximize *relative* instead of *absolute* payoffs, which can also be interpreted as spiteful behavior.³⁴ Either way, since on average the explanatory power of 'imitate the best' is higher than that of 'best response', it is expected an average aggregate effort from NE towards the OA outcome, consistent with Results 1 and 2. Moreover, since in 'imitate the best' subjects mimic the same pattern of behavior, a lower average Gini coefficient is expected to emerge, consistently with Result 3.

We summarize our main findings in the following statements:

Result 5. *The mixed-learning process significantly explains subjects' adjustments on average in treatments 'incomplete' and 'complete'. Moreover, the 'imitate the average' learning rule provides the highest explanatory power on average for effort adjustments in both treatments.*

Result 6. *The 'imitate the best' rule significantly explains subjects' adjustments on average in addition to the mixed-learning process in treatment 'complete', suggesting that subjects on average complement mixed learning with imitation when subject-specific information is provided, which is consistent with Hypothesis 3.*

³² Unlike other related studies (e.g. Huck et al., 1999, 2002) subjects in our experiments did not have a calculator available to help them determine the best response against the aggregate effort of the remaining players. Of course, subjects had all the information for making the calculations themselves. In this way, our experiment is similar to treatment Q in Offerman et al. (2002).

³³ Regarding the collinearity between 'imitate the average' and 'following the exemplary', the point-estimated sum of coefficients equals 0.570 and is significant at any standard values (p -value < 0.000).

³⁴ It can be shown that if all subjects maximize relative payoffs, then the OA profile $x^0 = (18, \dots, 18)$ is a NE of the game G. For a proof in the context of Cournot oligopoly see, e.g., Domènesh and Vriend (2003).

Including other covariates in (20) did not add any explanatory power to our model. We included an adapted version of ‘fictitious play’ (see, e.g., Fudenberg & Levine, 1999), in which subjects play a best response to the average of aggregate efforts of others across previous rounds, but it was not significant at any standard levels in both treatments. Round dummies (two-way fixed effects) were jointly not significant according to Wald tests in both treatments at the 10% level. The modeling of round-heteroskedastic residuals did not prove successful according to LR tests, even at the 20% level, which seems reasonable since our experiments lasted only 20 rounds. It should also be noted that, as shown in Section 3, the pure ‘imitate the average’ process theoretically converges to the average of all initial efforts. Since a difference between the distribution of initial efforts across treatments may interfere with the findings above (for example, if subjects started choosing effort levels closer to the open access outcome in treatment ‘complete’) we formally tested the equality of the distribution of initial efforts. According to Mann–Whitney U tests, the null hypothesis in which both distributions are equal cannot be rejected even at the 70% level. Hence, the difference between average group effort levels, efficiency levels, and Gini coefficients across treatments cannot be explained by differences in the distribution of initial efforts.

Two contrasting hypotheses are consistent with the finding of significant coefficients for ‘imitate the average’ and ‘best response’ learning rules in treatment ‘incomplete’, and for these rules and ‘imitate the best’ in treatment ‘complete’; namely that either *all* subjects in each treatment strictly mix the learning rules (e.g., a mix of ‘imitate the best’ and ‘best response’ in treatment ‘incomplete’), or that *some* agents purely adopt some learning rules instead of others when deciding their effort adjustments (e.g., some subjects purely ‘imitate the average’ and others are purely *belief learners*). In what follows we address this question formally by examining the estimated random part of (20), (21).

Table 8 indicates that according to LR tests, the variance slope (expressed as standard deviation) for ‘imitate the average’ is significant at standard levels (p -value < 0.000) in both treatments. This suggests that there is a population of *heterogeneous* ‘imitators of the average’, thus yielding a heterogeneous mixed-process learning across subjects. To interpret this result, note that it is (roughly) expected that 95% of subject slopes for ‘imitate the average’ lies in the range of 0.024–0.953 in treatment ‘incomplete’ and in the range of 0.053–1.237 in treatment ‘complete’.³⁵ These ranges are fairly wide, indicating that such heterogeneity has practical significance. While it seems that subjects strongly differ in the adoption of ‘imitate the average’, the employment of the other learning rules seems to be *homogeneous* across subjects, since the slope variances for these rules, as stated before, are not significant, even at the 10% level. The relevant interpretation is that *all* subjects are tempted to mimic the most successful subject to the same degree when subject-specific performance information is available, yielding higher levels of exploitation resource that emerge when this information is not available.

These findings are summarized in the following statement.

Result 7. *In both treatments the experimental data suggest that subjects significantly differ in the way they ‘imitate the average’, while in treatment ‘complete’ they seem to be homogeneously influenced by ‘imitate the best’, suggesting that their learning modes are equally sensitive to the information about others’ performances.*

Finally, even if all three examined learning and imitation theories in Table 8 have some predictive power in explaining effort decisions, subjects adjust on average only incompletely. None of the basic theoretical learning and imitation rules discussed in this work are able to explain the observed effort adjustments by themselves, thus revealing the richness and complexity of the learning process. Moreover, note that while in treatment ‘incomplete’ the intercept is not significant at any standard level, in treatment ‘complete’ the estimated intercept is positive, significant at the 1% level, and very substantial. This indicates that there are omitted factors other than the basic learning rules that explain adjustments toward higher effort levels. Nevertheless, our experimental findings strongly suggest that a simple and plausible (but incomplete) explanation for the significantly higher average aggregate effort levels in treatment ‘complete’ other than NE is that subjects *imitate the successful subject* even if individual, absolute payoffs are reduced in the long run by doing so.

6. Concluding remarks

In a series of experiments we investigated the behavioral consequences of the provision of subject-specific performance information in the context of a CPR game. We ran two treatments, one with ‘incomplete’ information and another with ‘complete’ information. In the former, subjects were informed only about the aggregate extraction effort level of the group, and in the latter they were also informed about the individual effort levels and payoffs of each subject. At the aggregate level, in the ‘incomplete’ information treatment we found that group exploitation efforts equal the aggregate NE of the constituent CPR game, thus yielding an inefficient use of the resource from the group perspective, while in the ‘complete’ information treatment we found an even greater overexploitation level. Our results are in stark contrast and complement those reported in the experimental work of Apesteguia (2006) regarding the effects of the provision of information about the payoff structure of the CPR game, since the results of his work show that aggregate behavior is not significantly different between the two treatments, converging in both cases to the NE effort level.

In order to motivate our hypotheses and explain our experimental results at the individual level, we make use of the theory of learning in games, which goes beyond standard non-cooperative game theory, allowing us to explore the three basic

³⁵ These ranges are obtained from the calculation $0.489 \pm 1.96 \times 0.237$ in treatment ‘incomplete’ and from $0.645 \pm 1.96 \times 0.302$ in treatment ‘complete’.

benchmarks in the commons context: NE, PE, and OA optimum outcomes. We use several learning and imitation theoretical models, which are based on contrasting assumptions about the level of rationality and the information available to subjects, namely: 'best response', 'imitate the average', mix of 'best response' and 'imitate the average', 'imitate the best' and 'follow the exemplary' learning rules.

According to our findings, learning emerges as a composite process in which different components coexist. In treatment 'incomplete' best response dynamics and imitation of the average rules seem to be important drivers of subjects' behavior, which is consistent with a finding of an average group effort level close to NE. By contrast, subjects in treatment 'complete', in addition to the mentioned learning and imitation rules, tend also to mimic the most successful behavior. This kind of behavior resembles a preference of subjects to 'beat the opponent', or alternatively it can imply that agents are boundedly rational. Either way, this behavior yields outcomes from the NE towards the OA outcome, consistently with our experimental results, so it provides a reasonable explanation for the higher average group effort levels in treatment 'complete'. We also found some evidence that subjects adjust efforts against 'follow the exemplary', and thus against reaching more cooperative outcomes. Moreover, the results of the random-effects model suggest that subjects are heterogeneous in the sense that they differ in the way their effort levels are explained by imitation of the average aggregate effort of others. However, the impact of the provision of subject-specific information on the learning mode employed by subjects is homogeneous, that is, all subjects are influenced by imitation of the most successful subject in the same degree when the information to do so is available. This is an interesting result that deserves more attention in future experimental research.

Finally, in contrast to Apesteguia's basic result, it can be inferred from our results that the provision of information, in this case subject-specific performance information, at the individual level does make a difference and can indeed worsen the tragedy of the commons.

Acknowledgements

The authors gratefully acknowledge comments by Matias Braun, Eduardo Gonzalez, Eduardo Moreno and Marcelo Villena. The authors also wish to thank the editor, Andreas Ortmann, and two anonymous referees for their useful comments. The usual caveat applies. Mauricio Villena thanks Fondecyt Grant No. 1101002 for financial support.

References

- Alós-Ferrer, C., & Ania, A. B. (2005). The evolutionary stability of perfectly competitive behavior. *Economic Theory*, 26(3), 497–516.
- Alós-Ferrer, C., & Schlag, K. H. (2007). Imitation and social learning. In P. Anand, P. Pattanaik, & C. Puppe (Eds.), *Handbook of rational and social choice*. Oxford University Press.
- Apesteguia, J. (2006). Does information matter in the commons?: Experimental evidence. *Journal of Economic Behavior & Organization*, 60(1), 55–69.
- Bosch-Domènech, A., & Vriend, N. J. (2003). Imitation of successful behavior in Cournot markets. *The Economic Journal*, 113(487), 495–524. 04.
- Bromley, D. W. (Ed.). (1992). *Making the commons work: Theory, practice, and policy*. San Francisco, California: International Center for Self Governance (ICSG) Press.
- Cardenas, J.-C. (2000). How do groups solve local commons dilemmas? Lessons from experimental economics in the field. *Environment, Development and Sustainability*, 2, 305–322.
- Cardenas, J.-C. (2003). Real wealth and experimental cooperation: Evidence from field experiments. *Journal of Development Economics*, 70, 263–289.
- Cardenas, J.-C., Stranlund, J. K., & Willis, C. E. (2000). Local environmental control and institutional crowding-out. *World Development*, 28, 1719–1733.
- Casari, M., & Plott, C. (2003). Decentralized management of common property resources: Experiments with a centuries-old institution. *Journal of Economic Behavior & Organization*, 51(2), 217–247.
- Dasgupta, P., & Heal, G. (1979). *Economic theory and exhaustible resources*. Cambridge: Cambridge University Press.
- De Hoyos, R. E., & Sarafidis, V. (2006). Testing for cross-sectional dependence in panel-data models. *Stata Journal, StataCorp LP*, 6(4), 482–496.
- Fudenberg, D., & Levine, D. (1999). *The theory of learning in games*. The MIT Press.
- Gordon, H. S. (1954). The economic theory of a common-property resource: The fishery. *Journal of Political Economy*, 62, 124–142.
- Hamilton, W. (1970). Selfish and spiteful behavior in an evolutionary model. *Nature*, 228, 1218–1225.
- Huck, S., Normann, H.-T., & Oechssler, J. (1999). Learning in Cournot oligopoly – An experiment. *The Economic Journal*, 109, C80–C95.
- Huck, S., Normann, H.-T., & Oechssler, J. (2002). Stability of the Cournot process – Experimental evidence. *International Journal of Game Theory*, 31(1), 123–136.
- Levine, D. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1, 593–622.
- Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica, Econometric Society*, 46(1), 69–85.
- Offerman, T., Potters, J., & Sonnemans, J. (2002). Imitation and belief learning in an oligopoly experiment. *The Review of Economic Studies*, 69, 973–997.
- Ostrom, E. (1990). *Governing the commons, in the evolution of institutions for collective action*. New York: Cambridge University Press.
- Ostrom, E. (2006). The value-added of laboratory experiments for the study of institutions and common-pool resources. *Journal of Economic Behavior & Organization*, 61(2), 149–163.
- Ostrom, E., & Walker, J. (1991). Communication in a commons: Cooperation without external enforcement. In T. R. Palfrey (Ed.), *Laboratory research in political economy* (pp. 287–322). Ann Arbor: University of Michigan Press.
- Ostrom, E., Walker, J., & Gardner, R. (1992). Covenants with and without a sword: Self-governance is possible. *American Political Science Review*, 86, 404–417.
- Ostrom, E., Gardner, R., & Walker, J. M. (1994). *Rules, games, and common-pool resources*. Ann Arbor: The University of Michigan Press.
- Plott, C. R., & Smith, V. L. (Eds.). (2004). *Handbook of experimental economics results*. New York: Elsevier.
- Rodriguez-Sickert, C., Guzmán, R. A., & Cárdenas, J. C. (2008). Institutions influence preferences: Evidence from a common pool resource experiment. *Journal of Economic Behavior & Organization*, 67(1), 215–227.
- Salazar, C., & Villena, M. G. (2005). Evolución de Preferencias en un Juego de Explotación de un Recurso de Propiedad Común: Evidencia Experimental. *Estudios de Economía*, 32(2), 159–186.
- Sally, D. (1995). Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and Society*, 7, 58–92.
- Schlag, Karl H. (1999). Which one should I imitate? *Journal of Mathematical Economics*, 31(4), 493–522.
- Seabright, P. (1993). Managing local commons: Theoretical issues in incentive design. *Journal of Economic Perspectives*, 7(Fall), 116–134.
- Skrondal, A., & Rabe-Hesketh, S. (2004). *Generalized latent variable modeling: multilevel, longitudinal and structural equation models*. Boca Raton, FL: Chapman & Hall/CRC.
- Stevenson, G. G. (1991). *Common property economics: A general theory and land use applications*. Cambridge: Cambridge University Press.

- Van Winden, F., Van Dijk, F., & Sonnemans, J. (1998). Intrinsic motivation in a public good environment. In C. R. Plott & V. L. Smith (Eds.), *Handbook of results in experimental economics*. North Holland: Elsevier Press.
- Vega-Redondo, F. (1997). The evolution of walrasian behaviour. *Econometrica*, 65(2), 375–384.
- Walker, J. M., Gardner, R., & Ostrom, E. (1990). Rent dissipation in a limited-access common-pool resource: Experimental evidence. *Journal of Environmental Economics and Management*, 19, 203–211.
- Walker, J. M., Gardner, R., Herr, A., & Ostrom, E. (2000). Collective choice in the commons: Experimental results on proposed allocation rules and votes. *The Economic Journal*, 110, 212–234.
- Zwiers, F. W., & von Storch, H. (1995). Taking serial correlation into account in tests of the mean. *Journal of Climate*, 8, 336–351.