Journal of Economic Behavior & Organization
Vol. 51 (2003) 217–247

JOURNAL OF Economic Behavior & Organization

www.elsevier.com/locate/econbase

Decentralized management of common property resources: experiments with a centuries-old institution

Marco Casari [a], Charles R. Plott [b,*]

[a] Department of Economics, 410 Arps Hall, Ohio State University, Columbus, OH 43210, USA
[b] Division of the Humanities and Social Sciences, 228-77, California Institute of Technology, Pasadena, CA 91125, USA

Received 16 October 2000; accepted 17 October 2001

Abstract

For centuries, villages in the Alps employed a special system for managing their common properties. Individual users could inspect other users at their own cost and impose a predetermined sanction (a fine) when a free rider was discovered. The fine was paid to the user who found a violator.

Experiments with the institutions demonstrate that this mechanism considerably improves efficiency of resource use. The classical model of identical selfish agents does not capture the data as well as a model with heterogeneous and linear other-regarding preferences. Altruism and especially potentially dysfunctional behavior, such as spite and mistakes, play important positive roles.
© 2002 Elsevier Science B.V. All rights reserved.

JEL classification: C70; C91; Q2

Keywords: Experiments; Common dilemmas; Environment

1. Introduction

For six centuries a special institution for managing common property resources was long and enduringly employed by villages in the Italian Alps. It began to emerge at the beginning of the 13th century for application to common forests and pastures and remained in place until it was forcefully removed by Napoleon in 1805 (Casari, 2000). The experiments and theory reported here are the first attempts to study this institution. The research reflects an

* Corresponding author. Fax: +1-626-793-8580.
E-mail addresses: marco@apollo.econ.ohio-state.edu (M. Casari), cplott@hss.caltech.edu (C.R. Plott).

0167-2681/02/$ – see front matter © 2002 Elsevier Science B.V. All rights reserved.
PII: S0167-2681(02)00098-7

attempt to pinpoint the reasons for the success of the system when compared to other systems with similar institutional features. Part of the research involves the study of experiments reported by others and the paradoxes one can see in their data.

The decentralized management system negotiated by the Italian villages had a very simple structure. The population of a village developed a contract among themselves, subject to the approval of the regional government, called 'Carte di Regola', where they described a system for monitoring and sanctioning those who are discovered violating or exceeding patterns of use that the villagers agreed upon in the contract. The 'Carte di Regola' specified in advance the conditions under which a sanction could be inflicted on a person found in violation of the contract and the amount of the fine. The village court would sentence people who used the common resource above an established limit to pay a fine proportional to the severity of the damage inflicted to the community. Any villager could report a violation but he usually incurred a cost in the form of a monitoring effort to discover the violator and additional costs to bring him to court. A share of such a fine usually went to the person who discovered the violator in order to give an incentive to monitor. The questions posed for study here are related to how this particular management system performs and why.

Experimental results have demonstrated that unregulated use of a common-pool resource, such as a common pasture or fishery, generates inefficient levels of use. However, the experimental literature contains a fundamental paradox, which we will call the spite/altruist paradox that is in need of replication and explanation. On the one hand, WGO (Walker et al., 1990) report group overuse of the resource at levels that go beyond what a pure free-riding Nash equilibrium model would predict. That is, the individuals choose to exploit the resource beyond what the self-interested Nash model would predict, as if individuals actually wanted to harm the group in a manifestation of spitefulness. On the other hand, in a public goods environment, which is essentially a transposition of a common-pool environment, Andreoni (1995) and Isaac et al. (1994) report behavior of an opposite nature. They report cooperation levels above the Nash equilibrium in public goods environments while WGO report cooperation levels below the Nash equilibrium in the common-pool environments. It is as if people are altruistic when faced with public goods provision and the opposite ("spiteful") when using common-pool resources. The spite/altruist paradox is the suggestion that cooperative behavior is the opposite within two institutions that are theoretically similar.[1]

The existence of the paradox suggested a need for a replication of results. The paradox could be due to subject pools, incentives, experimental protocols or features of the experimental environment including the experimenters themselves. Any attempt to explore the paradox should remove those variables from the list of possibilities. Replication helps to do that.

The experimental results reported in this paper are first that the 'Carte di Regola' institution is surprisingly successful in raising efficiency in the use of a common-pool resource. In addition, the patterns of results previously reported in the literature are replicated. As it turns out, the spite/altruist paradox can be explained by a model where agents have heterogeneous, linear other-regarding preferences. Such a model allows for selfish, altruistic and also spiteful agents. The interaction among these different types of agents can explain

[1] See also Ledyard, 1995, for a survey of cooperation in public good experiments and Palfrey and Prisbey, 1997.

a wide array of regularities related to these and other, similar experiments. The model provides insights about how the institution relies on the heterogeneity of preferences that exist in the population to efficiently coordinate decisions into the enforcement of management plans. In particular, the presence of spiteful individuals, who enjoy harming others through the imposition of sanctions, can serve a socially useful purpose. In addition, the model explains how the 'Carte di Regola' provides a greater opportunity for altruism to benefit the group.

The structure of the paper is as follows. Section 2 outlines aspects of the institutions under consideration. The basic common-pool resource environment is introduced in Section 3. The classical game theoretic model is outlined in Section 4. The stylized version of the 'Carte di Regola' monitoring and sanctioning institution is described in Section 5. These two sections also present the Nash equilibrium according to the classical model, where all agents are selfish and this fact is common knowledge. The linear other-regarding agent model is outlined in Section 6 along with the predictions for the various treatments with and without sanctions. The experimental procedures are described in some detail in Section 7. Baseline experiments without sanctions are reported in Section 8. Here, the results of WGO are replicated and extended. The results of experiments with sanctions can be found in Sections 9 and 10. Two treatments of the sanctioning mechanism are studied: a weak sanction institution that should have no effect on the outcome according to the classical model and instead greatly improves efficiency; a strong sanction institution that has the goal of bringing the group to the socially optimal outcome and turns out to fall slightly short of the target. A discussion in Section 11 demonstrates how the models extend themselves to related issues and data. Section 12 contains reflections about the nature of the institution and speculates about why it had such remarkable success. Section 13 is a summary of conclusions.

2. Management systems and sanctioning institutions

A variety of systems for the management of common property resources and public goods provision can be found in the literature. All systems are structured around the assumption that unless guided by specially crafted institutions the users of the resource hold the potential for over-exploitation and sub-optimal use or, in the case of public goods, a failure to pay for the provision of the public good. To an untrained eye the differences among the different systems might seem insignificant, but through the lens provided by economic theory and by game theory, the differences are substantial. This paper is concerned only with provisions of sanctions that are associated with the more decentralized systems. Table 1 lists major variables and the papers in which investigations have been reported.[2]

The variables fall into two major classes. The first is related to the cost of monitoring and what might be revealed as a result of monitoring. The presumption is that the actions of individual users are not necessarily freely observable and that management institutions might differ accordingly. The 'Carte di Regola' system studied here was crafted to address

[2] Some studies have been not included in the review either because they have an external sanctioning authority (Beckenkamp and Ostmann, 1999; Cardenas et al., 2000) or because the experimental design is for other reasons too different from ours (Yamagishi, 1988; McCusker and Carnevale, 1995). See Ostrom, 1990, for field examples.

Table 1
Monitoring and sanctioning institutions

	This study	Moir (1999)	Ostrom et al. (1994)	Fehr and Gächter (2000)
Environment	Common-pool resource	Common-pool resource	Common-pool resource	Public goods provision
Monitoring (search)				
Monitoring fee	Fixed fee for each request	Variable fee for each request	None	None
Are all individual use levels (investment levels) revealed?	No, only if somebody in the group requests it	No, only if the agent requests it	All use levels are public; all agent histories are public	All use levels are public; no individual history is available
Sanctioning				
Targeted agent				
Amount of the fine	In a fixed proportion of overuse	Subjective choice of inspector (variable upper bound)	Subjective choice of inspector (fixed upper bound)	Subjective choice of inspector (up to 100 percent of period earnings)
Condition for inflicting the fine	If overuse occurred	If overuse occurred	Subjective	Subjective
Multiple fines on the same action	No	Yes	Yes	Yes
Identity of targeted agent	Publicly known after fine	Publicly known after fine	Unclear[a]	Known only to targeted and inspecting agent
Inspecting agent				
Fee (cost of administering the fine)	Included in monitoring fee	Proportional to the amount of the sanction	Proportional to the amount of the sanction	More than proportional to the amount of the sanction
Who receives the fine	Inspector	Experimenter	Experimenter	Experimenter
Limits to requests of sanctions per period	None	Limited to the budget of the inspector	Each agent is limited to a single request	Unclear[a]
Identity of requesting agent	Not revealed	Not revealed	Not revealed	Not revealed[a]

[a] Monitoring is always perfect (i.e. there is truthful revelation of the action). In Ostrom et al., 1992 and Fehr and Gächter, 2000, agents can sanction each other but there is really no monitoring device, since the individual actions are automatically revealed to everybody at the end of each period. Moir (1999) introduces two distinct decisions, first to monitor an agent and then to eventually sanction her. We have compacted them in a single decision: to inspect an agent or not. An inspection uncovers another agent's action and automatically inflicts a sanction if some conditions are met. This feature was not explicitly described in the papers.

common property resources and the assumption is that for any one agent there is a cost of observing the levels at which various other agents use the resource.

The second class of variables is related to the nature of rewards and punishments involved in detecting an agent who overuses the resource, the sanctioning institution. Some variables are related to whether or not a sanction/fine is levied on the targeted agent, including the amount of the sanction or sanctions and whether the sanction was public. Other variables are related to the position of the inspecting agent including the financial costs and benefits of conducting an inspection. In particular, the amount of the sanction and indeed whether a sanction is levied or not is determined according to a known rule as opposed to being determined by the inspector.

A study of the Table 1 will reveal that the 'Carte di Regola' differs from other systems in two important ways. First, the fines resulting from social sanctions involve a transfer of income from the inspected to the inspector as opposed to a loss of system wealth. Of course, the cost of inspections is a loss to the system. By contrast, successful fines in the other systems are a dead weight loss to participants. Second, the punishment level cannot be changed by the inspector and is set to fit the crime. In particular, the targeted agent pays a fine only if use of the resource is more than a publicly known level and the fine is proportional to the excessive use.

The theory and experiments reported below are focused on the 'Carte di Regola'. The focus is on the efficiency properties of the system and on models that are put forward to help explain the behavior of the systems. Experiments with no sanctions establish an important baseline for comparison with other studies and for measuring the impact of the 'Carte di Regola'. Then by varying the level of sanctions the model with heterogeneous agents can be explored and compared to the classical model that has been used to capture the performance of other institutional arrangements.

3. The laboratory common property resource environment

The environment studied here is similar to common-pool resource environments studied elsewhere in the literature (Appendix A). A group of N agents interacts in the use of a common-pool resource. In general, agents have preferences for any benefits they receive from the availability of the common-pool resource as well as the cost to them of their efforts to harvest it. However, for purposes of developing the models in terms of preferences that will be found later in experimental environments, each agent i is characterized by preferences of a special form where π_i can be interpreted as a personal monetary value of resource use net, of any per unit cost α resulting from a level of effort x_i:

$$U_i(\pi_i) \tag{1}$$

where

$$\pi_i = \frac{x_i}{X} f(X) + \alpha(\omega - x_i) \tag{2}$$

where $x_i \in [0, \hat{x}]$ is interpreted as the level of "effort" expended by agent i in the use or harvesting of the common-pool resource, $X = \sum_1^N x_i$ the total use levels of the group of

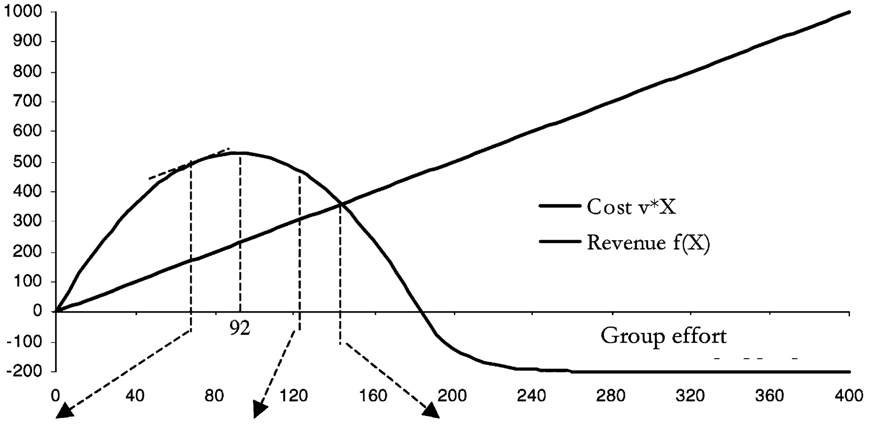

OPTIMAL USE	*NASH EQUILIBRIUM*	OPEN ACCESS	
72	128	144	Group use X
9	16	18	Symmetric individual use
100%	39.5%	0%	Efficiency level *(% of maximum rent)*

Fig. 1. Technology of the common-pool resource.

N agents, $f(X)$ the quantity of a common-pool resource, interpreted as a group revenue, and x_i/X the fraction of the total group effort that is expended by agent i. If $x_i = 0$, then the agent gets none of the common-pool resource. The "initial endowment" of agent i is indicated by a parameter ω and it is the fraction of revenue that is returned to i. If $x_i = 0$, the agent gets $\alpha\omega$. The group revenues $f(X)$ are non-linear in the group effort and first increase in X up to a maximal point and then decrease as illustrated in Fig. 1 and defined by Eq. (3).

$$f(X) = \begin{cases} aX - bX^2, & \text{if } X \leq 184 \\ 200[\mathrm{e}^{-0.0575(X-184)} - 1], & \text{if } X > 184 \end{cases} \tag{3}$$

The dynamic of renewable resources is generally modeled with a parabola (Clark, 1976; Gordon, 1954) as it is done in the first piece of $f(X)$. For high level of efforts, $f(X)$ has a lower bound at -200.

The parameters used throughout the experiments are $N = 8, \hat{x} = 50, \alpha = 5/2, a = 23/2$ and $b = 1/16$. The parameter ω has no effect on the incentive structure of the game theoretic models and in the experiment can be viewed as a fixed bonus to subjects for participation.

The resource environment can be viewed as the physical environment, including preferences and the relationships among effort and the magnitude of the resource pool. How those interact depends on institutional arrangements and the behavior that takes place within those institutions. The following sections address models of those relationships.

4. The classical model

Behavior within the context of a common-pool resource environment can be understood in the context of principles of game theory. The structure of the environment in the absence of any intervening institutions leads to a game theory model that we will call the 'classical model'. That is, a description of the physical environment is simultaneously a description of the institutional environment. Individuals are assumed to maximize $U_i(\pi_i)$ as defined in Eqs. (1) and (2), subject to the strategy set $x_i \in [0, \hat{x}]$ and the relationships found in Eq. (3).

It is important to notice that the utility function defined by Eq. (1) has no parameters related to the income of others or risk aversion. Furthermore, it is assumed in the classical model that all agents have the same utility function.

The Nash equilibrium of the game is easy to compute. From the first-order conditions to maximize earnings $\partial\pi_i/\partial x_i = 0$, we derive the best response functions $x_i^* = ((a-\alpha)/2b) - (1/2)X_{-i}$, where $X_{-i} = \sum_{j\neq i}^{N} x_j$, which is a linear function of the use level of everybody else (Fig. 2). As all agents have identical incentives and preferences, at the Nash equilibrium the group outcome is $X^* = (N/(N+1))((a-\alpha)/b)$, when $(Na - 184(N+1)b) > N\alpha$. In this game, the Nash equilibrium is unique and symmetric.

Given the parameter values $N = 8$, $a = 23/2$, $\alpha = 5/2$ and $b = 1/16$, the Nash equilibrium outcome is $X = 128$, which corresponds to an individual effort level of $x_i = 16\forall i$.

The efficiency of an outcome is defined in reference to the group earnings minus the endowment money $\Pi' = \Pi - N\alpha\omega$. A group outcome is normalized using the maximum earnings that the group could reach ($\bar{\Pi}' = 324$). This efficiency index scores 100 percent when the resource is used at the socially optimal level and 39.5 percent at Nash equilibrium. The efficiency of the Nash equilibrium goes down as group size goes up because there is less incentive to take into account the strategic interaction among the agents.

The last observation is of particular interest because it leads to another model that we shall call 'open access'. When the number of users goes to infinity, agents completely ignore the strategic interactions among them. If the number of agents is finite but agents are poorly informed about the consequences of the actions they take (in the sense that agents

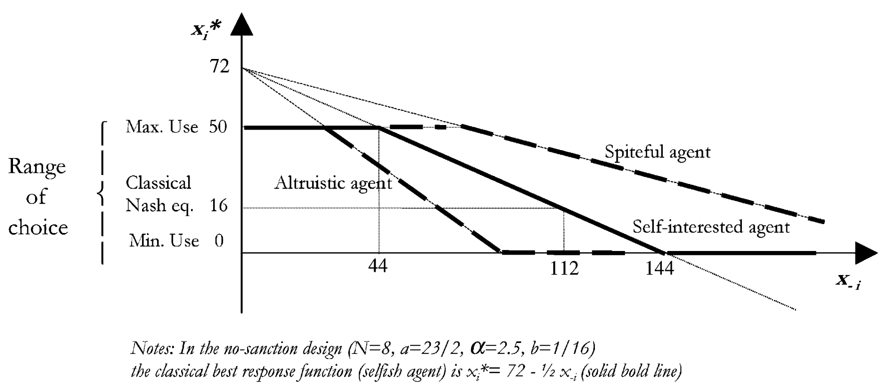

Fig. 2. Best response functions, no sanction treatment.

believe that $\partial[(x_i/X)f(X)]/\partial x_i = 1$), then the model leads to exactly the same behavior. The solution at the open access equilibrium has efficiency at zero percent, which implies a complete destruction ($\Pi' = 0$) of the potentially positive incomes that the group could have made out of the common-pool resource. The efficiency is zero percent because agents use the resource up to a point where average costs equal average benefits.[3] Given the parameters adopted, the socially optimal outcome is at $X = 72$ and could be obtained if all the eight agents in the group choose $x_i = 9$ while the open access outcome is at $X = 144$ ($x_i = 18$; Fig. 1). If the group uses the resource above the open access level, group earnings are negative and efficiency can be as negative as −321 percent (for $X = 400$).

Proposition 1A (Resource use without sanctions). *Without a sanctioning institution, the classical Nash equilibrium outcome has an efficiency of* 39.5 *percent* [*group appropriation* $X = 128$]. *All the agents use the resource at an identical rate of* $x_i = 16$.

5. The 'Carte di Regola' system under the classical model

The basic features of the 'Carte di Regola' mechanism for monitoring and sanctioning are captured by a simple game where any agent i in the group has the option of selecting other individuals $j \neq i$ for inspection after he has privately decided his own exploitation level of the common-pool resource. At a unitary cost k, the inspector can view the decision of any individual. If the inspected individual has exploited the resource excessively, relative to a publicly known amount λ, a fine s_j is imposed and paid to the inspector:

$$s_i = \begin{cases} 0, & \text{if } x_j \leq \lambda \\ h(x_j - \lambda), & x_j > \lambda \end{cases}$$

The parameter h is the unitary fine for each extra unit of effort and measures the stiffness of the punishment. Agent i makes a profit when the fee k paid to carry out the inspection is more than compensated by the transfer s_j from agent j to agent i, namely when $r_{ij} = (s_j - k) > 0$. As the transfer s_j is proportional to the use of agent j in excess of the "legal" threshold λ, a profit is made when agent j uses the resource more than $\tilde{x} \equiv (k/h) + \lambda$.

Considering both the use and the inspection decisions, the payoff of agent i is: $\pi_i = (x_i/X)f(X) + \alpha(\omega - x_i) - I_i s_i + \sum_{j\neq i} I_{ij} r_{ij}$, where $I_{ij} = 1$ indicates that agent i inspected agent j and $I_{ij} = 0$ indicates otherwise; moreover, $I_i = 1$ if $\sum_{j\neq i} I_{ji} \geq 1$ and $I_i = 0$ otherwise.

The game model takes place sequentially in two steps. In step 1, agents decide the use of the common-pool resource. In step 2, agents make inspection decisions by selecting a subset of agents for which he/she wants to pay for having inspected. Anyone can inspect anyone else so at this stage there are 2^7 actions available to each agent. An inspection involves at the same time information discovery as well as punishment. Before requesting an inspection, agents know only the total group use. After the inspection phase, the use levels of the inspected individuals become public information. Neither the number of inspections, nor the identities

[3] To use a market analogy, the three situations correspond to a monopoly, Cournot oligopoly and perfect competition situation.

of the inspectors are public information. Another feature of the mechanism is that there is no accumulation of sanctions when more than one agent request to inspect the same person. If multiple inspectors exist, one inspector is randomly selected out of the requesting agents.

The threat of sanctions reduces the incentives for high use levels of the common-pool resource because they increase the cost of any effort above the "legal" limit, $x_i > \lambda$. The consequence is a downward shift in the best response function of a targeted agent (Fig. 2). The degree of the shift depends on the perceived probability p_i, that an agent has of being inspected. If such probability is positive, the agent's best response is to use the resource less than in the corresponding no inspection case.

The symmetric Nash equilibrium outcome is a pair (X^*, p^*) that jointly satisfies the following two conditions, one for each of the steps of the game model:

$$(X^*, p^*) \text{ that solves } \begin{cases} X^* = \dfrac{N}{N+1}\left(\dfrac{a-(\alpha+hp^*)}{b}\right) \\ p^* = \begin{cases} 0, & \text{if } E[X^*] \leq N\tilde{x} \\ 1, & \text{if } E[X^*] > N\tilde{x} \end{cases} \end{cases}$$

Notice the crucial role that is played in the inspection decisions by the assumption that agents are identical and by the assumption that it is common knowledge (so in equilibrium $x_i = x^*$ and $p_i = p^* \forall i$).

As will be explained in the later sections, experiments will be performed with two levels of sanctions—a weak sanction and a strong sanction design. Both will be analyzed under the assumptions of the classical model. In the strong sanction design, the unitary fine h is four times higher than in the weak sanction treatment and the definition of excessive resource use λ is stricter (lower).

The weak sanctions ($k = 7$, $\lambda = 9$, $h = 1$) are designed to have no effect on the Nash equilibrium of the group outcome of the classical model. The weak sanction parameters are set so that no inspection is strictly profitable when the individual use of resource is at or is lower than $\tilde{x} = 128/N = 16$.[4] Thus, in classical Nash equilibrium the total group use does not change from the no sanction design level and there are no inspections, $(X^*, p^*) = (128, 0)$. As explained, even without the threat of sanctions, selfish agents have no incentive to use the resource more than $X = 128$. Instead, when $X > 128$ the inspection of each one of the agents will be profitable. If all the agents expect to be inspected with probability one, though, the group resource use drops to $X = 113.7$ (66.4 percent efficiency).

Proposition 2A (Resource use with weak sanctions). *The introduction of weak sanctions does not change the classical Nash equilibrium level* $[X = 128]$.

Proposition 3A (Inspections under weak sanctions). *When weak sanctions are introduced, at the classical Nash equilibrium inspections pay zero However, if a slight psychological cost exists, no inspections are requested.*

[4] The equilibrium $X^* = 128$, $p^* = 0$ is slightly altered when the agents have a "trembling hand" in their inspecting decisions. If a subject inspects "by accident" and this kind of event is common knowledge the equilibrium will be below $X = 128$. We believe that this point affects neither our basic results nor our conclusions.

Strong sanctions ($k = 7, \lambda = 7, h = 4$) are designed to move the equilibrium away from the inefficient outcome of the no sanction treatment to a fully efficient outcome ($X = 72$).[5] The following properties are immediate.

Proposition 4A (Resource use with strong sanctions). *The introduction of strong sanctions moves the classical Nash equilibrium outcome very close to the socially optimal level* [*above* 99 *percent efficiency*, $X = 71.1$].

Proposition 5A (Inspections under strong sanctions). *When strong sanctions are introduced, at the classical Nash equilibrium all agents inspect everybody.*

In the symmetric Nash equilibrium under strong sanctions, $(X^*, p^*) = (71.1, 1)$, all the agents are inspected and the group efficiency is at 99.97 percent. Inspecting an agent is profitable when $x_i > 70/N = 8.75$. A slight discrepancy between the total group use X^* in equilibrium and the social optimal value was preferred to assigning non-integer numbers to the parameter values. The difference in terms of efficiency is, however, negligible.

6. The linear other-regarding model

This section outlines a simple, linear other-regarding model where an agent's utility depends not only on personal earnings but also on the earnings of the other people in the group and computes the Nash equilibrium for appropriate parameters. Many other-regarding models can be found in the literature (Krebs, 1970; Rabin, 1993; Ito et al., 1995; Chan et al., 1997; Levine, 1998; Bolton and Ockenfels, 2000; Saijo, in press; Fehr and Schmidt, 1999).The version of heterogeneous, linear other-regarding model adopted here intends to capture in a parsimonious way a specific motivation for economic behavior. Let, π_i be defined as in Eq. (2) and let $\Pi_{-i} = \sum_{j \neq i} \pi_j$. Then the assumptions of the heterogeneous, linear other-regarding model are:

$$\text{other-regarding preferences,} \quad U_i(\pi_i, \Pi_{-i}) = \pi_i + \gamma_i \Pi_{-i} \tag{4}$$

$$\text{range of preferences,} \quad \gamma_i \in [-1, +1] \tag{5}$$

$$\text{heterogeneity,} \quad \exists_{i,k} \text{ with } i \neq k \text{ such that; } \gamma_i \neq \gamma_k \tag{6}$$

In this model, agent *i* is willing to give up US$ 1 of personal earnings (π_i) in order to see the other people's earnings (Π_{-i}) changed by $1/\gamma_i$ dollars. The classical model is a special case when all agents are selfish ($\gamma_i = 0, \forall i$). Agent *i* is called altruistic if other-regarding parameter of agent *i* is $\gamma_i > 0$ and if $\gamma_i < 0$ then agent *i* is called spiteful. A spiteful agent finds enjoyment in decreasing the earnings of others and is, therefore, willing to pay in order for that to happen. Although not crucial for the conclusions, we fix by assumption boundaries to the degree of altruism and spite, $\gamma_i \in [-1, +1]$, such that nobody chooses to

[5] The inspection fee *k* was not changed because it relates to the difficulty of observing other people's actions, which is a technological parameter generally outside the control of the mechanism designer.

pay more than US$ 1 to modify the group earnings by less than US$ 1.[6] The definition of spite adopted is similar to Saijo (in press) but different from the concept of envy suggested by Mui (1995). The model does not incorporate any reciprocity nor equity nor fairness considerations.[7] As in the classical model, agents are assumed to be risk neutral and the preferences of all the agents, common knowledge.

The rest of this section is devoted to the computation of the Nash equilibrium of the heterogeneous, linear other-regarding model in the three levels of sanctions, no sanctions, weak sanctions and strong sanctions.

When there are no sanctions, the best response function with linear other-regarding preferences is $x_i^* = ((a - v)/2b) - ((1 + \gamma_i)/2)X_{-i}$. There are both symmetric and asymmetric equilibria but all outcomes are within the values of the two "extreme" equilibria computed assuming $\gamma_i = -1, \forall i$ or $\gamma_i = 1, \forall i$.[8] For instance, if all agents are fully altruistic ($\gamma_i = 1$) the equilibrium is at the socially optimal outcome $X^* = 72$. If all agents are moderately altruistic ($\gamma_i = 1/7$) then $X^* = 115.2$, while if they are moderately spiteful ($\gamma_i = -1/7$) then $X^* = 144$.[9]

With heterogeneous preferences the individual use levels are heterogeneous. In particular, the lower the other-regarding parameter γ_i, the higher is the individual use x_i. In other words, spiteful agents use the resource more than selfish agents and selfish agents use it more than altruistic ones.

Proposition 1B (Resource use without sanctions). *Without a sanctioning institution, the Nash equilibrium outcome with heterogeneous, linear other-regarding agents has an efficiency ranging in* [−321 *percent*, 100 *percent*]— *or use levels in* [72, 400]— *depending on the preference structure of the agents. In general, individual agents use the resource at different rates, with spiteful agents using it more than altruistic agents.*

When preferences in a group are heterogeneous there are two consequences for inspecting decisions. First, because agents use the resource at different rates (Proposition 1B) some actions could be inspected for profit even in the weak sanction treatment. In particular, if there are at least two types of agents and at least one agent is altruistic, then there exists an agent that uses the resource above $x_i = 16$. Second, the decision to inspect depends on the preferences of the inspector. In particular, spiteful agents are willing to request also non-profitable inspections. The utility of agent i from inspecting agent j is $V_i = U_i(\pi_i, \Pi_{-i}) + (s_j - k) - \gamma_i s_j$. A spiteful agent finds enjoyment not only from the cash flow $(s_j - k)$ but also from decreasing the income of agent j by s_j. On the other hand, an altruistic agent does not consider all the money $(s_j - k)$ as a gain since the sanction s_j

[6] We assume that the vector of other-regarding parameters γ is such that the individual response function is within the interval $x_i \in [0, 50]$. This assumption might further restrict the range of γ to a subset of the $[-1, +1]$ interval.

[7] For an reciprocity model applied to a common-pool resource game, see Falk et al., 2000.

[8] The same outcome X^* can be the equilibrium result of more than one vector of agent preferences but given a vector of preferences, an individual shift in preferences has a predictable change on the outcome, $\partial X^*/\partial \gamma_i < 0$.

[9] An interesting case is when the preferences in the group are symmetrically heterogeneous or, in other words, when for every altruistic agent i with $\gamma_i > 0$ there is a spiteful agent k with $\gamma_k = -\gamma_i$. The Nash equilibrium with symmetrically heterogeneous agents is in general more efficient than the classical Nash equilibrium. For instance, when half of the agents $\gamma_i = 1/7$ and the other half $\gamma_k = -1/7$ the outcome is $X^* = 126$, $x_i = 0$ and $x_{k_i} = 31.5$.

has been subtracted from somebody else that he cares about. A completely altruistic agent ($\gamma_i = 1$) never inspects, while a complete spiteful one inspects when $x_j > 12.5$ in the weak sanction treatment and when $x_j > 7.9$ in the strong sanction treatment. When facing the same use pattern, spiteful agents are more aggressive inspectors than altruistic agents because they inspect for lower values of x_j. This propensity is captured by the fact that an agent of type γ_i inspects the actions of *j* if $\gamma_i < -(k - s_j)/s_j$.

Sanctions also affect level of use decisions. When a sanction is imposed, agents of every type lower the resource use because the sanction increases the unitary cost of an appropriation effort. Under the threat of sanctions, the best response function of an agent *i* with linear other-regarding preferences is $x_i^* = ((a - v - p(1 + \gamma_i)h)/2b) - ((1 + \gamma_i)/2)x_{-i}$, if $x_j > \lambda$. The utility of agent *i* when inspected is $V_i = U_i(\pi_i, \Pi_{-i}) - s_j + \gamma_i(s_j - k)$ and, given the nature of the 'Carte di Regola' institution, sanctions induce a spiteful agent to lower his appropriation level proportionally more than any other type. In fact, spiteful types are doubly troubled while paying a sanction: for a start they earn less money and in addition somebody else is going to be better off. Altruistic types are the least affected because a portion of the sanction is transferred to other people in the group they care about. Only fully altruistic agents, though, are not affected by sanctions in their appropriation decisions.

How do sanctions change the relative appropriation levels of altruistic versus spiteful agents? When there are no sanctions, altruistic agents use the resource less than spiteful agents. As explained above, sanctions induce a more than proportional reduction in the level of use by spiteful people than altruistic people (Proposition 1B). The effect of relatively light sanctions—as it is the case of the weak sanction treatment—is to reduce inequalities in use levels within the group, although spiteful agents still use the resource more than altruistic agents (Proposition 3B). When sanctions are sufficiently heavy—as in the strong sanction design—spiteful agents' concern for the transfers of money induced by the sanction dominates the incentives coming from their own earnings. The reduction in the use levels across the different types is such that altruistic agents use the resource more than spiteful agents. This reversal of the solution can be verified by substituting the parameter values into the best response function (Proposition 5B). In a strong sanction environment, when all agents are fully spiteful the Nash equilibrium outcome is $X^* = 64$ while is $X^* = 72$ when all agents are fully altruistic, which corresponds to an efficiency range [98.77 percent, 100 percent].

Finally, the predictions of the model with heterogeneous, linear other-regarding preferences—which hold under some mild regularity conditions on preferences—are listed below.[10]

[10] Because of the reversal in behavior of spiteful agents from the weak to the strong sanction treatment, a stringent condition on preferences is required to obtain the prediction that in equilibrium all agents are inspected with probability one. In order to have inspections with weak sanctions it is sufficient if two or more agents are not altruistic. A sufficient condition for all agents to be inspected in the strong sanction treatment is that in addition to the above, the most spiteful agent is "not too far apart" from the next. More formally, when agents are ranked low to high other-regarding parameters $\gamma_{(1)}, \gamma_{(2)}, \dots, \gamma_{(8)}$, then $\gamma_{(1)}, \gamma_{(2)} \le 0$ and $|\gamma_{(1)} - \gamma_{(2)}| < 0.25$. This condition is satisfied in three of the four no sanction experiments. The 7 April experiment satisfies a different sufficient condition: No agent is very spiteful ($\gamma_I > 0.45 \forall i$) and at least two agents are not altruistic ($\exists_{i,j}$: γ_i, $\gamma_i \le 0$). The above statements are based on the estimation described in point (e) in Section 8. There is a much milder condition that ensures that at least 87.5 percent (i.e. 7/8) of the action is inspected (it suffices that at least one agent is not strongly altruistic, $\gamma_{(1)} < 0.08$).

Proposition 2B (Resource use with weak sanctions). *When agents are heterogeneous and other-regarding, the introduction of weak sanctions improves efficiency upon the Nash equilibrium level without sanctions.*

If two or more agents are selfish or spiteful (*sufficient condition*), *the improvement is strict.*

Proposition 3B (Inspections under weak sanctions). *With weak sanctions and agents with heterogeneous, linear other-regarding preferences*:

(i) *there are inspections when two or more agents are not altruistic* (*sufficient condition*);
(ii) *the heaviest users* (*the most spiteful agents*) *are more aggressive inspectors than lightest users (the most altruistic agents*) *and also purposively request non-profitable inspections*.

Proposition 4B (Resource use with strong sanctions). *When agents are heterogeneous and linear other-regarding, the Nash equilibrium outcome with strong sanctions has an efficiency above* 98.5 *percent* ($X^* \in [64, 72]$) *under some regularity conditions on preferences*.

Proposition 5B (Inspections under strong sanctions). *With strong sanctions and agents with heterogeneous*, *linear other-regarding preferences*:

(i) *under some regularity conditions on preferences all agents are inspected*;
(ii) *lightest users* (*the most spiteful agents*) *are more aggressive inspectors than heaviest users* (*the most altruistic agents*).

7. Experimental procedures

A total of 56 subjects were recruited from the campus of the California Institute of Technology for a total of 10 experimental sessions. The different treatments are outlined in Table 2. There are three different sanctioning designs: no sanction, weak sanction and strong sanction. Within each treatment, half of the experiments were conducted with inexperienced subjects and the other half with experienced subjects.

There were eight subjects in each experimental session. All subjects were seated at terminals, separated by partitions and assigned identification numbers. No communication was allowed.[11] Instructions were read aloud to everyone. The experiments were run on networked personal computers using dedicated software for Netscape.

Procedures and rewards of the original study were changed as illustrated in Appendix A. In particular, the marginal monetary incentives were increased three- or four-fold compared to WGO and the minimum safe earning was reduced. In order to make it easier to understand the rules, the instructions were rewritten, special software developed and the action space was rescaled.

[11] Rocco and Warglien (1996) conducted a common-pool resource experiment allowing email exchanges without significantly improving cooperation.

Table 2
Summary table for use decisions

	No sanction				Weak sanction				Strong sanction	
Experiments	1	2	3	4	5	6	7	8	9	10
Date	12 February 1998	8 September 1998	7 April 1998	9 September 1998	25 February 1998	24 August 1998	8 April 1998	25 August 1998	31 August 1998	1 September 1998
Sanctions	No	No	No	No	Yes	Yes	Yes	Yes	Yes*	Yes*
Experience	No	No	Yes	Yes$^+$	No	No	Yes	Yes$^+$	No	Yes$^+$
Number of rounds	32ˆ	32	33	32	27	27	27	27	27	27
Period endowment (tokens)	4	4	4	0	4	4	4	0	0	0
Conversion rate (US$ per franc)	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.03	0.03
Group use										
Average	124.7	134.2	133.1	133.3	125.6	115.4	110.8	110.0	91.8	79.1
Classical Nash equilibrium	128	128	128	128	128	128	128	128	71.11	71.11
Minimum	100	85.5	121.5	87	87.3	87	92.5	90	74	69
Maximum	154	161	149.5	167	186.8	160	154.7	137	126	95.5
S.D.	14.84	15.49	6.01	15.46	23.13	20.57	12.15	10.20	13.22	6.50
S.D. first/second half	4.52	1.68	2.59	0.79	0.92	0.83	2.57	1.46	2.18	0.85
Group efficiency (percent of maximum rent)										
Average rent (1)	42.29	20.97	27.36	23.00	34.67	55.74	68.20	70.15	89.73	98.24
First 25 periods	42.48	23.74	27.13	22.27	32.43	52.71	68.94	70.75	89.29	98.14
Last two periods (after announce)	54.23	1.28	28.21	41.06	62.68	93.58	58.95	62.62	95.29	99.47
Fees (2)	–	–	–	–	13.04	8.24	6.64	7.68	17.20	17.04
Fines	–	–	–	–	14.01	8.90	6.49	6.19	43.39	28.44
Net average rent [(1) − (2)]	–	–	–	–	21.63	47.49	61.56	62.47	72.53	81.20
Individual average use levels										
Lowest average user	8.8	10.6	9.9	9.0	11.7	10.2	11.0	11.1	8.9	7.1
Highest average user	32.5	39.2	44.8	35.2	20.0	17.3	18.4	15.3	14.1	13.0
Rank correlation first/ second half	0.917	0.902	0.961	0.966	0.907	0.892	0.926	0.872	0.823	0.858

The experiments were done at the California Institute of Technology Sanctions: "No" is a no sanction experiment; "Yes" means that a monitoring and sanctioning device was added to the no sanction experiment; "Yes*" indicates a different set of sanctioning parameters. Experience: "No" means that no subject has ever participated in any one of these experiments before; "Yes" means that all the subjects have already participated in this type of experiment (on 9 December 1997 at the earliest); "Yes$^+$" means that all the subjects have participated the day before in this type of experiment with the same group of people. Number of rounds: number of effective rounds of interaction, which excludes two practice rounds; (ˆ) in column 2, a paper copy of the "return from investment" table was handed to the subjects between the 10th and the 11th round instead of before the first round. During the whole experiment the table was projected on the wall of the room. Period endowment: The endowment ω indicates the number of tokens given each period toeach subject.

The problem was presented as an abstract decision-making situation where there was an opportunity to earn money by "investing" in a market. The use level was chosen without knowing the choices of the other subjects. Use levels were expressed in "tokens" and payoffs were in terms of "francs" (an artificial laboratory currency with a publicly known dollar-exchange rate) and in dollars.

An experiment lasted from 1 to 2:20 h including the preliminaries. Individual earnings ranged from US$ 5.80 to 53.10. Each subject was paid privately in cash immediately following the experiment.

Within each period of the experiment, there was just one step in the no sanction treatment and two steps in the sanctioning treatments. During step 1, the computer prompted a request for a number of tokens that the subject wished to put in the market. A subject could digit any real number between 0 and 50. After everybody completed the input, the total group use and gross group return were displayed. In the no sanction treatment, subjects could also see their individual period payoff (your share of gross, cost of tokens, period payoff), while in the sanctioning treatments this part was postponed until the end of step 2. Step 2 gave a chance to inspect other subjects. By clicking on a box next to the subject identification number, a subject could ask to uncover the use level of any number of subjects from 0 to 7.

The period payoff was computed and explained in terms of its three components: result of use decisions, result of inspections asked and notices of the eventual charge for an inspection targeting the subject. A full record of the past decisions could always be seen, including personal individual uses and cumulative payoffs, total group uses and gross group returns and the individual use levels of inspected agents uncovered by anybody in the group.

To ensure that the rules were well understood we adopted the following procedure. First, the rules were publicly explained in detail and with examples. Second, a quiz was given. All the correct answers were read aloud after completion of the quiz and the ones where mistakes were noticed in the answers were further explained. Third, two practice periods were run, to help the subjects familiarize themselves with the rules of the experiment and with the software. After the two practice rounds, a number of periods from 27 to 33 were run. Subjects were not told the number of rounds that were to take place. At the end of the third-before-the-last period, an announcement was made that the experiment was going to end in two periods. After the experiment was over, a questionnaire was submitted to the subjects asking for the strategy they followed.

8. Results of no sanction experiments

The experimental results are compared with the predictions of the classical model (1A) and of the linear, other-regarding agent model (1B). The data demonstrate that the predictions of the classical model in the no sanction environment are subject to systematic errors. The linear other-regarding agent model does better.

8.1. Result 1A

Without a sanctioning institution the resource is overused relative to the classical Nash equilibrium (i.e. with homogeneous, selfish agents). People cooperate less than expected according to that model and are worse off than the model predicts.

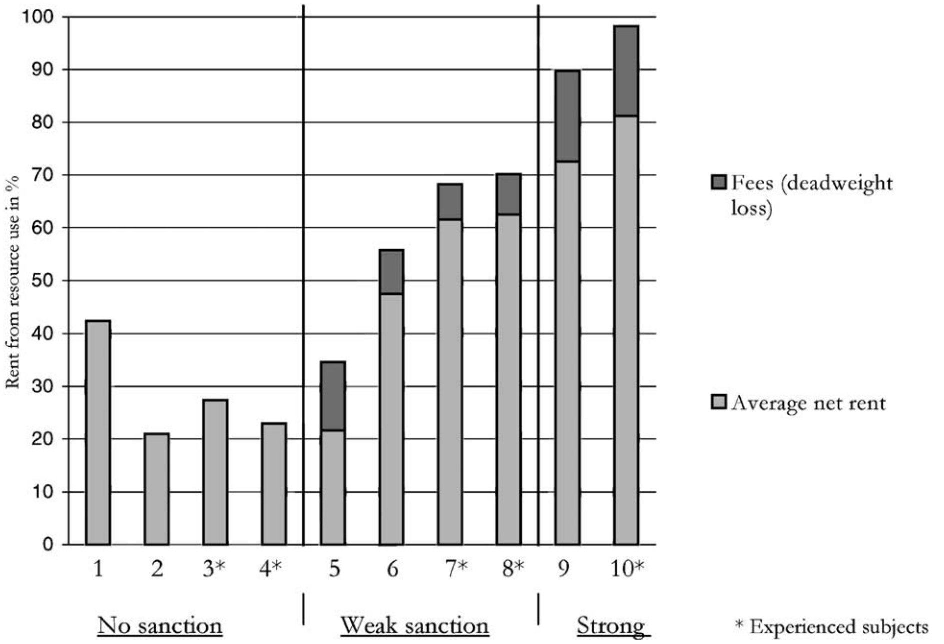

Fig. 3. Average efficiency by experiment.

8.1.1. Support

In terms of efficiency the groups scored 28.4 percent of the maximum possible net return, a value that is lower than the classical Nash equilibrium level of 39.5 percent. The overall average of the group use for the four experiments was 131.3, which is statistically different from 128 at a 0.01 level[12] (see Table 2 for details). The group use varied considerably across periods, ranging from a minimum of 85.5 to a maximum of 167 tokens.

Neither subject experience nor time effects alter the main conclusion that the group use is persistently above the one-shot classical Nash equilibrium level. In particular, experienced subjects do not perform better than inexperienced subjects do. Differences in efficiencies actually favor inexperienced subjects (25.2 percent versus 31.6 percent; Fig. 3).

Moreover, there is no indication of collusion among users. In fact, if a repeated interaction effect is present, the pattern in the total group use should be:

(a) a convergence to the one-shot Nash equilibrium from below, i.e. in the range $X \in [72, 128]$;
(b) an eventual jump to the one-shot Nash equilibrium level after the end-of-experiment announcement has been made.

[12] The symmetric Nash equilibrium value $X = 128$ was never recorded in any of the 129 rounds in which the appropriation decisions were taken. The open access level is $X = 144$.

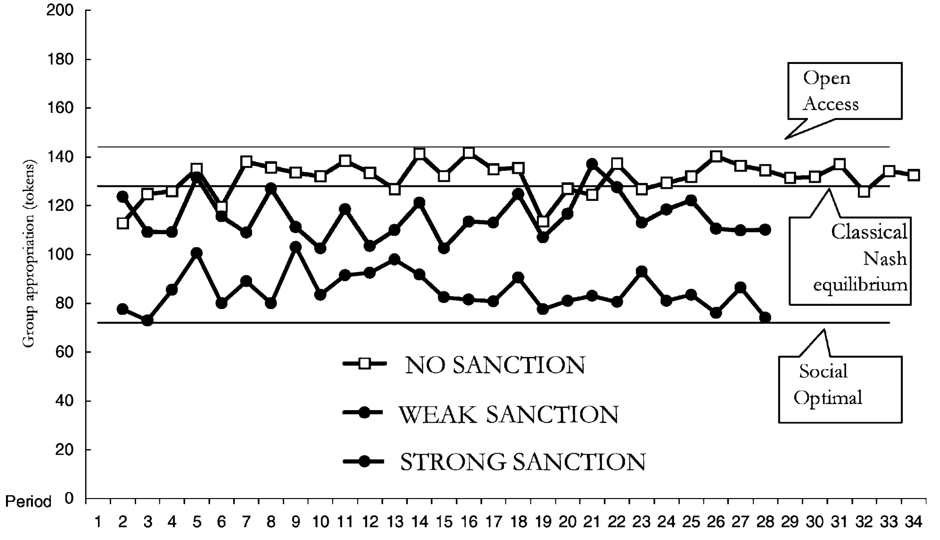

Fig. 4. Average group use by experimental treatment.

A comparison between the first half, second half, and after announcement period averages[13] show no statistical differences at 0.01 significance level. As an overall average, the values are 131.37 in the first half, 131.39 in the second half and 130.31 after the announcement (Fig. 4).[14] The volatility of the group use level decreases over time in three out of four experiments (see variance comparisons in Table 2) but it mostly reflects oscillations around the same average.

In conclusion, the observed level of group use can be explained by the heterogeneous, linear other-regarding agent model, given an appropriate pattern of group preferences that is biased toward spite.[15]

8.2. Result 1B

In a sanction-free environment, a model of heterogeneous, linear other-regarding agents is compatible with the patterns of individual resource use better than the classical model.

[13] The subjects were not told the length of an experiment. In no sanction experiments, the first half includes periods 1–15, second half 16–30 (or 16–31) and after announcement 31–32 (or 32–33). In sanction treatment, the first half includes periods 1–12, second half 13–25 and after announcement 26–27.

[14] The presence of a repeated interaction effect and of time effects was also evaluated using the Ashenfelter–El Gamal model, which is described in Noussair et al., 1995. The asymptote for the no sanction experiments is 134.0, which is statistically different from the equilibrium level of 128 but not significantly different from the overall average group use of 131.3 at a 0.05 level. This result confirms once more that the overuse of the resource persists and does not tend to die out. The convergence to the asymptote starts from below for all the experiments, which supports point (a).

[15] Consider for example a group with three types of agents: two are moderately altruistic $\gamma_i = 1/21$, four self-interested agents and two quite spiteful ones $\gamma_i = -1/4$. The group appropriation is $X^* = 132$ with individual appropriations x_i of 6, 12 and 36, respectively.

Furthermore, some clues exist in the data suggesting that it is an appropriate modification to the classical model. Specifically, overuse and underuse are properties of individuals. About 37 percent of the agents are other-regarding and most of them are spiteful.

8.2.1. Support

The patterns of individual use do not conform to the one-shot classical Nash equilibrium prediction of $x_i = 16 \forall i$. Individual actions are very dispersed in the action space (a) and this variability is due to individual heterogeneity (b). Individual heterogeneity is not a consequence of confusion (c) but is consistent over time (d) and is due to other-regarding preferences (e).

(a) Only 2.5 percent of all actions are $x_i = 16$. The actions within a 25 percent bandwidth around the prediction (i.e. in the interval [14, 18]) account for 15.7 percent of all the actions and the rest are not symmetrically distributed around that value: about 61 percent are below and 23.3 percent are above. The mean is 16.4 and the standard deviation (S.D.) is 10.00.

(b) A brief look at the individual average use levels makes clear that heterogeneity is a trait of agents and that only a few agents were accountable for a systematic overuse (Fig. 5). We can reject the hypothesis that the agent average use is at the individual symmetric Nash equilibrium ($x_i = 16$) for 28 out of 32 agents at 0.05 level. Within each experiment there are at least four different types of agents whose use behavior is statistically different at 0.05 level. The presence of different types of individuals is a common finding in the experimental literature (Brandts and Schram, 2001; Van Winden et al., in press).

(c) There are reasons to believe that the differences in individual behavior are not due to confusion. The experimental design was not simple and a possible explanation of heterogeneity is that the "heavy users" might have been confused subjects who did not

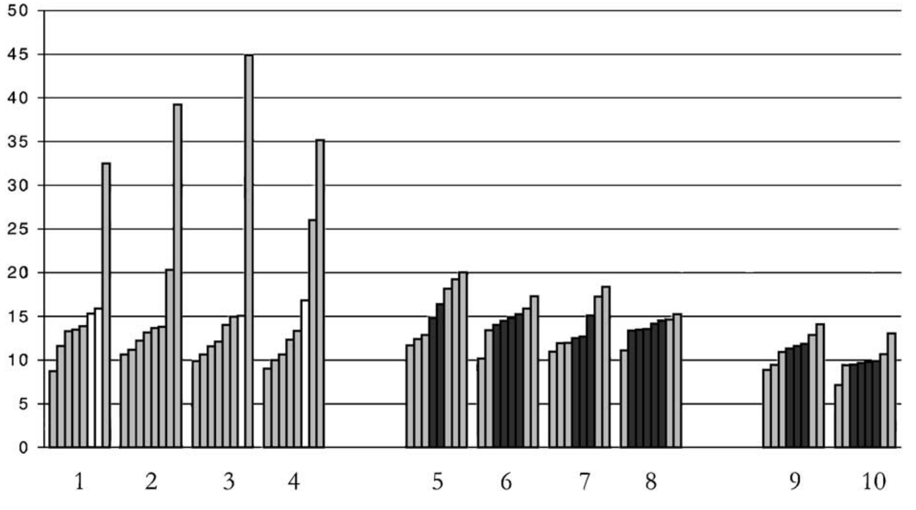

Fig. 5. Agent average use levels.

properly understand the incentive structure of the experiment.[16] The evidence from the quiz completed by each subject before the experiment does not show any support for this option. We have assigned a score to each quiz taken, which is 1 if all the answers are correct, 0.5 if some answers are not perfect but it is clear that the subject overall understood the rules and 0 if there are substantial and repeated mistakes. The four highest users score an average of 0.92 against a general average of the 32 subjects of 0.89. In other words, the heavy users are not less skilled than average.

(d) There is a remarkable consistency over time in the individual use patterns, which indicates that the differences across agents are purposive rather than random. Consider the ranking of agents by individual average use levels in the first and second half of the experiment. A rank correlation computed with an OLS regression without a constant term informs about the existence of any monotonic relation between the two rankings. A 1-value coefficient denotes a perfect positive correlation in agents' behavior over time. When all the no sanction experiments are pooled together, the estimated coefficient is 0.936 (number of observations is 32, $R^2 = 0.88$; see Table 2 for single experiment regressions). This test supports the view that over time agents are consistently heterogeneous.

(e) The estimation of the individual best responses for the linear other-regarding agent model on the experimental data leads to consistent results. Since the difference between the classical and other-regarding model is just in the slope of the best response function (cf. (2) and (2″); see also Fig. 2), the regressions assume a correct value for the intercept. The estimation is done under the assumption that the agents expect the others to act in period t as they did in period $(t-1)$: $x_{i,t} = 72 - ((1+\gamma_i)/2)X_{-i,t-1} + \varepsilon_i$. All the 32 agent-specific estimated values of the other-regarding parameter γ_i fall into the allowed interval $[-1, +1]$. The γ_i estimates range from a minimum of -0.40 to a maximum of 0.08. About 37 percent of the agents have a parameter γ_i significantly different from zero at a 0.05 level, and our model classifies them as either altruistic when γ_i is positive (2 agents) or spiteful when γ_i is negative (10 agents).

To sum up, in the experiments group efficiency of 28.4 percent is below the classical Nash equilibrium level of 39.5 percent. WGO reported an average negative efficiency (−3.2 percent) but at a closer analysis the difference between the two studies turns out to occur in the earlier rounds and dies out over time.[17] Moreover, individual actions are widely heterogeneous. WGO reports that in the 48 percent of the rounds not a single agent used 16 tokens. In our experiments the figure is 90 percent.[18]

[16] If the heavy investors are confused subjects, however, it is unclear why we do not find them in the experiment with the weak sanction treatment (Fig. 5). Such experimental design is more complex than the no sanction design, although the threshold level for sanctioning gives a vague clue about the equilibrium level and the monetary incentive against high appropriation levels are higher.

[17] The convergence values estimated with the Ashenfelter–El Gamal model are statistically indistinguishable (131.97 WGO and 133.78 ours), although our data reject the Nash equilibrium value of $X = 128$ at a 0.05 level where WGO data are more noisy. Average group efficiency is computed using WGO's three experiments and the first 20 rounds of the four no sanction experiments in this paper. The 0.95 confidence interval of the Ashenfelter–El Gamal asymptotes are [124.38, 139.56] for WGO and [129.33, 138.22] for ours.

[18] Part of the increase observed in our experiments might be due to the rescaling of the action space and to the opportunity to invest any real and not only integer number.

The predominance of spiteful agents over altruistic ones can account both for the overuse at the group level and for the observed pattern of individual actions. A model relying on homogeneous, selfish agents cannot explain either one of the two regularities.

9. Results of weak sanction experiments

This section describes the outcome of four experiments run under the weak sanction treatment and in particular it focuses on the inspection decisions (Section 9.2) and their effects on use decisions (Section 9.1).

9.1. Result 2

With the introduction of weak sanctions, group efficiency improves substantially. Resource use efficiency moves from below the classical Nash equilibrium to above the classical Nash equilibrium. These results are not predicted by the classical model (Proposition 2A) but they are consistent with the heterogeneous, linear other-regarding agent model (Proposition 2B).

9.1.1. Support

In other words, an institution—such as weak sanctioning—that according to the classical model should have no effect on behavior has instead a significant impact on the outcome. The classical model predicts an efficiency of 39.5 percent for both the no sanction and weak sanction treatment. The gross efficiency level with weak sanctions is 57.2 percent, which is roughly double the efficiency without sanctions (28.4 percent). When a correction is made for the differences in length among experiments, the situation does not change (28.9 percent versus 56.2 percent for the first 25 periods only).

The net efficiency—computed by subtracting the inspection fees (8.9 percent)—is 48.3 percent, which is about 20 points above the no sanction level (Tables 3 and 4). The inspection fee represents the cost of using the inspection mechanism and is a deadweight loss for the group. Instead, the fines are a plain transfer from an agent to another. Moreover, efficiency improves with experience: the net efficiency is on average 62 percent for experienced subjects versus 34.6 percent for inexperienced ones (Fig. 3).

Table 3
Inspection decisions

	Weak sanction			Strong sanction		
	No	Yes	Total	No	Yes	Total
Number of actions						
Balance if action is inspected						
Negative, $s_i - k < 0$	372	276	648	4	141	145
Zero, $s_i - k = 0$	20	37	57	0	0	0
Positive, $s_i - k > 0$	27	132	159	0	287	287
Total	419 (48.5 percent)	445 (51.5 percent)	864 (100 percent)	4 (0.9 percent)	428 (99.1 percent)	432 (100 percent)

Table 4
Spiteful agents inspect more than altruistic agents

OLS regression	Weak sanctions		Strong sanctions	
	Coefficient	*P*-value	Coefficient	*P*-value
Dependent variable	Total number of requests of inspections per period			
Sample size (without median users)	486		243	
Independent variables				
Highest users (dummy variable)	0.34	0.015	−1.23	0.000
Period use of the other agents	0.06	0.000	0.03	0.159
Constant	−5.42	0.000	1.58	0.366

The classical model predicts insignificant coefficients for the highest users dummy variable. See Fig. 5 to identify the median users whose actions were excluded from the regressions.

The total group use is substantially lower for sanction experiments than for no sanction ones. As an overall average, group use drops from 131.3 to 115.5 tokens (statistically different at 0.01 level). The aggregate use is statistically different from both the classical Nash equilibrium and the socially optimal level (0.01 level). The classical Nash equilibrium $X = 128$ was recorded in 2 out of 108 periods. When considering the classical model, the overall group use average is not statistically different from the one-probability inspection prediction ($X = 113.7$) (0.05 level). The group use across periods ranged from a minimum of 87 to a maximum of 186.8, which is wider than the same range for no sanction experiments.[19]

The comparison of the results in the no sanction and sanction environments is carried on under the assumption that agents were drawn from a population with identical preference patterns. Even if the agents were not the same in the different experiments, we think that the conclusions drawn under the above assumption are reasonable.

9.2. *Result 3*

With the introduction of weak sanctions:

(i) about half of the actions are inspected;
(ii) the highest users are more aggressive inspectors than the lowest users.

These results are not predicted by the classical model (Proposition 3A) but they are consistent with the heterogeneous, linear other-regarding agent model (Proposition 3B).

9.2.1. *Support*

In the weak sanction treatment 51.5 percent of the actions were inspected (Table 3). The classical model cannot explain either the magnitude of efficiency improvement nor why those inspections were requested in the first place. The introduction of different types of

[19] Similar results come from the estimation of the Ashenfelter–El Gamal model explained above. The ordinary least-squared asymptote of $X = 114.8$ is not significantly different from the one-probability inspection level (0.05 level) while it is different from the zero-probability level.

agents can account for the data. Under the classical model, if agents have a probability 0.5 of being inspected the Nash equilibrium is $X = 120.9$ with a 53.9 percent efficiency. That seems to be a good fit of the data. When agents are experienced, however, the effect on efficiency is stronger than that (62 percent efficiency with 41.4 percent of actions inspected). The reason for the higher than expected increase in efficiency relative to the share of actions inspected comes from the fact that the agents that have been identified as the heavy user types face a much higher probability of being inspected than the others.

Why are so many inspections requested? Because inspectors are aware that agents are heterogeneous in their effort decisions and because some agents are willing to request unprofitable inspections. As the first point has been already documented this section will focus on the second point.

According to the heterogeneous, linear other-regarding agent model, the "heavy user" types are spiteful agents. Those agents will purposely request some unprofitable inspections. In particular, the data support the prediction that spiteful agents are more aggressive inspectors than altruistic agents.

Agents were divided into three groups according to their average use level in the experiment. The inspecting behaviors of high versus low users was compared, keeping out of the analyses the group of median users among whom there were no significant differences in individual use at a 0.05 level. Relatively spiteful agents requested on average more inspections per period than relatively altruistic agents did, when controlling for the resource use by all the other agents in the group. This conclusion is based on the sign and significance of the coefficient of the dummy variable for highest users in Table 4 (positive for weak sanctions, negative for strong sanctions).[20]

10. Results of strong sanction experiments

10.1. Result 4

Strong sanctions have the effect of increasing resource use efficiency as predicted by both the classical and the heterogeneous, linear other-regarding agent models.

Efficiency levels fall short of the Nash equilibrium of both models. Experienced subjects tend to be closer to the equilibrium.

10.1.1. Support

In the strong sanction experiments, the total group use was on average 85.1. This level was significantly (0.01 level) higher than the outcome predicted by both the standard model (71.1) and the heterogeneous, linear other-regarding agent model ([64, 72]).[21] The group use across periods ranged from a minimum of 69 to a maximum of 126. The efficiency level is very high, 94 percent, but still sub-optimal and lower than the target level. When the inspection fees (17.1 percent) are subtracted, the net rent is 76.9 percent (Tables 3

[20] Regressions for each single experiment confirm this general conclusion with the exception experiment 5 where the highest investors dummy is not significant at 0.10 level.

[21] This conclusion does not change when the Ashenfelter–El Gamal model is estimated. The ordinary least-squared asymptote is 86.1 and none of the predicted values are in its 95 percent confidence interval.

and 4). Sub-optimality might be due to the inexperience of subjects, since there is a significant improvement in the group efficiency when subjects are experienced (gross rent 98.2 percent versus 89.7 percent).[22]

10.2. Result 5

In the strong sanction environment:

(i) more than 99 percent of all actions are inspected;
(ii) the lowest users are more aggressive inspectors than the highest users.

Result (ii) is not predicted by the classical model (Proposition 5A) but is consistent with the heterogeneous, linear other-regarding agent model (Proposition 5B).

10.2.1. Support

(i) About 99.1 percent of the actions were inspected, a value definitely close to the 100 percent predicted. Although almost all the agents were inspected every period, not all the agents requested to inspect everybody every period, as predicted by the classical model. On average an agent requested less than four inspections per period instead of seven (Table 3).[23]
(ii) Relatively spiteful agents (lower users) are more aggressive inspectors than relatively altruistic agents (higher users) as can be seen from Table 4. In the strong sanction treatment, the highest user variable is negatively correlated and significantly so with the number of inspections.

11. The spite/altruist paradox and implications of other-regardingness

One might argue that the model with heterogeneous, linear other-regarding agents fits the experimental data from the common-pool resource experiments better than the classical model because of the flexibility given by N additional parameters (one for each agent). Of course that is a concern, but additional credibility is given to the model—despite this argument—by its ability to provide insights about three additional perplexing aspects of behavior uncovered by these and other experiments.

The first phenomenon is the correlation between use decisions and inspections decisions at the individual level. As was illustrated by point (ii) in Sections 9.2 and 10.2, the highest users are the most aggressive inspectors in the weak sanction experiments while the opposite is true for the strong sanction experiments. The phenomena cannot be explained by the classical model. However, this regularity has an elegant explanation within the heterogeneous, linear

[22] The inequality in the average use across agents is substantially lower than in the no sanction environment (Table 2). The average S.D. of the agent period earnings from appropriation is of 5.6 francs, and 1.6 francs fines are subtracted.

[23] The large inspection balance is a surprise. In equilibrium with $(X^*, p^*) = (71.1, 1)$, the inspection balance is 1.4 percent of the maximum rent. The result is a balance of 18.8 percent, more than 10 times higher than what was predicted. The reason of such “success” was not mainly in the exceptional ability in discovering high investors but in the high average value of total group appropriation.

other-regarding model. On one hand, spiteful agents are always more aggressive inspectors than altruistic agents. On the other hand, the use levels of spiteful agents are highly sensitive to the level of sanctions inflicted on themselves because part of the sanction they pay makes someone else better off. The increased income of another person, the person who made the inspection, is something that the spiteful do not like. Such sensitivity goes to the point that in strong sanction experiments spiteful agents use the resource less than altruistic agents.

The second phenomenon is the spite/altruism paradox in the no sanction design. When agents are heterogeneous and other-regarding this paradox can be explained by looking at the boundary conditions of the individual "strategy space" and noticing that they are different in common-pool environments and public goods environments. Consider the following three designs: (A) the individual use level strategy space is [0, 50]; (B) the individual strategy space is [0, 20]; (C) the strategy space is [0, 16]. All three designs have the same Nash equilibrium, $x_i^* = 16$, and differ only in the strategy space. From the point of view of the classical model, designs A and B simply supply agents with options that are irrelevant to their actions and there is no substantive difference with C. Common-pool environments are typically like A or B while public goods environments with an equilibrium of zero contributions by selfish agents are like C. In a common-pool environment the spiteful individuals have great latitude for harming others, while in the public good environment, where they cannot take away amounts of the public good provided by others, they have a relatively limited strategy space. Altruistic agents, on the other hand, are not placed in such asymmetric conditions between the two environments and thus the actions of altruists have a disproportionate influence on the outcome and efficiency levels in public goods experiments.

In the common-pool experiments, changes in behavior as a result of changes in the strategy space were first reported by WGO. An increase in efficiency is expected going from design A to B, although smaller than going from A to C, because the best response of some agents could be in [20, 50], either due to the high spite of some or the high altruism of others. This "surprising" efficiency improvement was observed and reported by WGO. Moreover, while individual data are not available from the WGO study, such data that are available suggest that some individual choices of use in design B were at the maximum possible.[24] On the public goods experiment side of the issue, data frequently reflect unexpectedly high levels of contribution (altruism) but typically the experimental design does not allow elements of spite to be expressed because the Nash equilibrium is at zero levels which leaves no room for spiteful behavior. When high levels of spite have relatively great possibility for expression in relation to altruistic attitudes in a public goods environment, then the aggregate outcome shows lower levels of cooperation.

The possible influence of spite in the context of the heterogeneous, linear other-regarding model can be seen in the data reported by Isaac and Walker (1998). A switch in behavior from efficiency above the Nash equilibrium to efficiency below the Nash equilibrium takes place as the strategy space is kept the same and the equilibrium is moved away from the boundary of zero contributions to the public good toward the Pareto optimal levels.[25] As this takes place the range of spiteful actions that can be expressed is expanded relative to

[24] In the experiments corresponding to our design B ("restricted strategy space"), the observed modal strategic response of individuals was to use the resource to the maximum of their ability (Ostrom et al., 1994, p. 121).

[25] This change is similar to going from case C to B or from C to A.

the range of possible altruistic actions and as a result the system efficiency falls as spiteful actions emerge.

A third pattern of phenomena relates to a slightly different sanction institution studied by Ostrom et al. (1992). Consider a system where an agent pays a cost in order to inflict a sanction on some other person but receives no monetary reward for doing so, similar to the system of vigilantes discussed in Section 12. The model of a finitely repeated classical game predicts that no such sanctions should be requested. Yet, in the experiments conducted by WGO such sanctions were frequently observed. Furthermore, when the amount of sanction per unit of cost was increased, the number of sanctions increased even though the result was a net decrease in system efficiency. Such behavior can be understood in terms of the heterogeneous, linear other-regarding model. Spiteful agents enjoy harming others and the more so, the more intense is the punishment for a given cost. Phenomena do exist in experiments that are difficult to explain without resort to either attitudes of fairness or repeated interactions. In the WGO experiments, the fines were not connected to guilt or innocence. Even an altruistic person would be forced to pay a fine if someone decided to inspect. In such a world, the spiteful people would not care who they inspect and altruists would have no incentive to inspect. Yet, people inspect others and, more importantly, they tend to target those who are heavy users of the common property resource. Without a modification the spiteful model cannot explain such phenomenon. On the other hand, preferences reflecting fairness or envy would lead to inspecting the richest people since that would lead to a more even distribution of the income.

Explanations other than the heterogeneous, other-regarding model have been considered to account for the patterns of data reported in this paper. Specifically, the open access model can explain some of the results better than the classical model but does not perform as well as the other-regarding model. In the experiments without sanctions, the group use is in between the classical model prediction ($X^{NE} = 128$) and the open access model prediction ($X^{OA} = 144$) levels but statistically different from both. The introduction of weak sanctions raises the efficiency of 27.3 points, somewhat less than predicted by the open access model (39.5) and remarkably more than the no change predicted by the classical model. The actual level of cooperation achieved with weak sanctions is considerably higher than predicted by both models (57.2 percent versus 39.5 percent). The open access model is closer than the classical model in explaining the results with strong sanctions ($X^{OA} = 80$, $X^{NE} = 71.1$, actual $X = 85.1$). Neither model accounts for the observed individual heterogeneity both in use and inspection. Finally, the notes of the subjects after the experimental sessions show a widespread concern about the strategic interaction of one agent with the rest of the group.

12. Insights about successful institutions

Reflection on the nature of the ‘Carte di Regola’ institution reveals five features that may have been important to its success. These are features suggested by theory but are not fully captured by the model that we employ. Nevertheless, because of their novelty they seem to be worthy of mention.

First, the success of the ‘Carte di Regola’ system appears to be related to its ability to use the heterogeneity of preferences to socially advantageous ends. Notice first the "institutional invisible hand", the incentive features that channel attitudes that might normally be considered as socially dysfunctional, such as spiteful preferences, into socially useful purposes. People with spiteful preferences choose to monitor and sanction at a monetary loss. But when their preferences are considered as part of system efficiency, they are the ones who can perform the function most efficiently and are channeled into the activity for which they have a comparative advantage.

One might think that the 'Carte di Regola' is similar to a system of vigilantes but there are important differences. In the model, spiteful people do not care whom they hurt; they just enjoy hurting others, so it is important to direct and constrain them. The 'Carte di Regola' directs them by reserving the judgment of guilt for the courts, as opposed to the vigilantes, who would be happy to judge anyone guilty. The court convicts a person only when the guilt is consistent with social purposes. The magnitude of punishment is also reserved for the courts in the 'Carte di Regola' system, while in a vigilante system the inspector is allowed to judge and determine punishment. So, the 'Carte di Regola' constrains what the spiteful can do to the guilty. Thus, there are important differences (Ostrom et al., 1992).

Second, the 'Carte di Regola' also channels arbitrary or random behavior toward useful ends. Such behavior might ordinarily be regarded as dysfunctional from the point of view of economic efficiency. Mistaken inspections or impulsively random inspections are costly to the inspector and thus involve efficiency losses, but the fact that inspections take place has consequences for those who are excessive users of the common-pool resource by increasing the likelihood that a sanction is imposed. Thus, random inspection behavior that would appear irrational helps preserve the commons.

Third, one might think that the 'Carta di Regola' would evolve in response to incentives to bribe inspectors and thus become ineffective as a management system. Instead, the 'Carte di Regola' has a sort of "bribery proof" or "institutional change immunity" feature that could prevent infection from a system of bribery that could undermine its effectiveness. The immunity resides with the court's transfer of a portion of the fine to the inspector and, especially, on the existence of multiple potential inspectors. While an agent who detects a violator could be better off accepting a bribe from the violator than reporting the event to the court, the violator faces the possibility that other agents detect him after he has already bribed one agent. Indeed, there may be nothing to prevent an inspector from extracting more than one bribe from the same violator by acting in collusion with other inspectors. When the level of the bribe relative to the sanction and the number of potential inspectors are sufficiently high, the best choice of a discovered violator is to pay the fine instead of bribing. From this analysis one would predict that anonymity (the identity of the inspector not being revealed) played no special role in the experiment.

Fourth, the 'Carte di Regola' system has a type of "repetition independent" efficiency. It does not depend on repeated play for the creation of efficient management of the commons. One can easily imagine an enforcement system in which those who bear the cost of enforcement do so because of the benefits derived from repeated play. In such systems, the willingness to bear the cost of inspection is balanced against the benefits of modified behavior in repeated interactions. Notice that is not the case in the 'Carte di Regola' system in which repeated interactions have no central role.

Finally, the 'Carte di Regola' also exhibits a type of "fairness independence". The effectiveness of the enforcement system does not depend upon any agent being motivated by a sense of fairness. Such special preferences need not be in the population for it to work. In fact, the 'Carte di Regola' does not depend on other-regardingness at all for proper functioning but it does channel people with such attitudes to useful purposes should they exist.

13. Conclusions

The primary purpose of the research reported here was to study the effectiveness of a special decentralized system of sanctioning rules applied to the problem of managing common property resources. The rules were fashioned after an ancient method of managing renewable resources in the Italian Alps called 'Carte di Regola', where people could inspect one another, inflict punishments and be rewarded for doing so according to well defined legal proceedings.

The paper reports three different types of results. Results of the first type are related to the performance as a system for managing the commons. Results of the second type are related to the relative accuracy of alternative models of system behavior that help us understand why the system performs as it does, including both macro and micro observations. Results of the third type reflect insights about the nature of the success of these particular institutions that are not clearly anticipated by theory.

The overriding result is that the 'Carte di Regola' greatly improves the efficiency of resource use over a system that carries no sanctions. Importantly, the improvement is not only in terms of gross efficiency but also net efficiency, where the costs of administering the system (the inspection fees) are deducted. Under a weak sanction treatment there was a spectacular improvement in gross efficiency from 28.4 to 57.2 percent. Once the inspecting costs are considered, (net) efficiency remains very high (48.3 percent). Group behavior in the strong sanction environment shows large improvements from 28.4 to 94 percent gross or 76.9 percent net but is not at the optimal level.

The behavior of the system is largely understandable in terms of theory that explains previously observed inaccuracies and paradoxes. Two major changes in the classical game theoretic model account for the improvements. First, agents were assumed to have a capacity to be other-regarding and second, the nature of the other-regarding capacity could differ from person to person. The other-regarding feature is captured by a parameter of the utility function that defines the weight placed on the monetary income of others in comparison with one's own monetary income. An individual is called altruistic if there is a gain in utility when the income of others increases while his own income remains constant. A spiteful agent loses utility when the income of others increases while his own income remains constant. About one-third of the agents are other-regarding to various degrees, either altruistic or spiteful, as measured by the model.

The heterogeneous, linear other-regarding model predicts observed behavior that the classical model does not predict. For example, while both the classical and heterogeneous, linear other-regarding agent models predict the observed increases in efficiency due to strong sanctions (Sections 9.1 and 10.1), the increases in system efficiency resulting from weak sanctions are not predicted by the classical model but are predicted by the heterogeneous other-regarding model. The patterns of resource use and agents' decisions to inspect the use levels of others also hold evidence of differential model accuracy. The prediction of the heterogeneous, linear other-regarding agent model is that spiteful agents are more aggressive inspectors than altruistic agents under all conditions (weak and strong sanctions). In fact, the spiteful agents are even willing to request unprofitable inspections. On the other hand, the level of use of the resource by the spiteful agents relative to altruistic agents reverses from relative heavy user to lowest level user as the treatment is changed from weak to strong sanctions. This seemingly perverse relationship, the flip in the relative behavior, is reflected in the data from the experiments (Sections 9.2 and 10.2).

A spite/altruist paradox emerges from the experimental literature because reported individual behavior appears to be altruistic in public goods experiments while being spiteful in the theoretically similar common-pool experiments. Given the results reported here, the paradox is potentially explained by the existence of both spiteful and altruistic agents in conjunction with a change in the strategy spaces available to agents in the public goods environment as compared with the common-pool environment. In the latter the spiteful have greater latitude for action and the consequences of that latitude are manifest in system behavior.

The 'Carte di Regola' seems to have been a successful system for managing common-pool resources. The research provides insights into the features that might account for its success. First, the system holds a potential for increasing efficiency. Second, the system seems to have a type of resistance to infection of bribes. One might think that this legal arrangement could easily evolve to a system of widespread bribery. On the contrary, within the 'Carte di Regola' rules the person caught in violation could readily prefer to pay the sanction as imposed by the court to the payment of a smaller bribe. Third, the 'Carte di Regola' system channels possibly harmful human tendencies such as mistakes and spitefulness to useful social ends. Costly, random or mistaken decisions by some agents to inspect others serve as a deterrent to those who would otherwise exploit the commons and thus does not result in a total loss or waste of resources as do some types of economic mistakes. Spiteful attitudes that might ordinarily be considered as dysfunctional are channeled in a different way. Spiteful agents are those who find enjoyment in decreasing the earnings of others and who willingly inspect others at a loss just for the thrill of inflicting damage. However, as it turns out such people provide a public good by bearing the monitoring cost of the system. The genius of the 'Carte di Regola' is that it allows agents to specialize into activities that they like while controlling potentially dysfunctional behavior. The control and efficient direction of dysfunctional aspects is something that seems to be lacking in other decentralized management systems such as a system of vigilantes.

Acknowledgements

The financial support of the National Science Foundation and the Laboratory of Experimental Economics and Political Science is gratefully acknowledged. We are also deeply appreciative of the collaboration of Stephen Van Hooser who developed the software program and helped extensively with the experiments. Many thanks to Giangiacomo Bravo, Anthony Kwasnica, David Messick, Robert Moir, William Morrison, Alvaro Gonzalez Staffa, Leslie Title, James Walker, Roberto Weber and an anonymous referee for their suggestions.

This paper has benefited from comments from participants at the ESA meeting in Lake Tahoe, Nevada, University of Trento, Italy, and IASCP conference at Indiana University, Bloomington, University of Pittsburgh and CERGE-EI in Prague.

Appendix A

A.1. Comparison with WGO experimental design

There are differences between our no sanction design and WGO, both in the incentive structure (1–4) and in the way the information is conveyed (A–D):

1. The range of the choice variable x_i has been rescaled from [0, 25] to [0, 50] to increase the perceived distance between Nash equilibrium and open access level. Moreover, the socially optimal use level corresponds to an integer individual use (9 tokens versus 4.5 tokens in WGO). In the experimental design of WGO, a change in the endowment ω is tied to the upper limit in the individual use of the common-pool resource because $\omega - x_i \geq 0$. In our experimental design the endowment level ω is chosen for the sole purpose of ensuring some minimum earnings to the agents and the range of choice is $0 \leq x_i \leq \hat{x}$. Our design has a continuous decision space (i.e. any real number) that avoids the problems of asymmetric equilibria that WGO had.[26]
2. The conversion rate franc/dollar has been increased four times from US$ 0.01 to 0.04 (or US$ 0.03) per 1 franc in order to maintain a higher effort level by the participants in the experiment. As a result, the difference in terms of individual earnings between the social optimum and open access points has increased from US$ 0.405 to 1.62 (these figures and all the earnings are expressed in dollars per person per period).
3. The minimum safe earning level has been decreased. If nothing is used in the "risky" market 2, the original earnings are US$ 1.25. In our setting, a zero use ($x_i = 0$) yields a period return of US$ 0.4 (when $\omega = 4$). The change implies a downward shift in the payoff but does not affect the incentive structure. The reason of the change is to limit the maximum earnings that would have otherwise been too high given the new conversion rate (point 2).
4. The gross group return has been modified in the interval [184, 400]. The function *f*(*X*) now has a lower bound at −200 francs that is much higher than it originally was. A group use of $X = 184$ has an efficiency of −142 percent in both settings. At $X = 400$, the difference in dollar earnings is small (US$ −5.6 instead of US$ −6.75). The reason for the change is to limit the maximum loss given the new conversion rate in case people "go crazy" (see point 2). In the experiments conducted, we observed a use level above 184 just in one round out of 291.

(A) *Graph instead of formula*: WGO provided the subjects with the analytical expression of the gross group return from market 2. The expression may have been used to compute the equilibrium. We replaced it with a plot of the function *f*(*X*) because we think that

[26] Thank you to T. Saijo for sending us a memo, "Stability in common-pool resource experiments: validity of Nash response", March 1994 on this point.

the graph would improve the understanding of the basic underlying phenomenon. The graph in the instructions is similar to the upper part of Fig. 1, once the cost line is removed.

(B) *Detailed table*: In order to compute the equilibrium, subjects could use a very detailed table of gross group return. The table gives the gross group return and the return per token used for 100 values of the total group use, compared to the 10 values given by WGO. All theoretical equilibrium points are listed in the table given to the subjects. Our table does not supply the marginal returns, which instead WGO provided.

(C) *Different software*: The software was run on Netscape and was written specifically for this application. It includes a calculator to compute the cost of tokens, the gross group return and the individual share of gross for every possible real level of use in the admissible range.

(D) In WGO, market 1 represents the opportunity cost of the use and yields a constant return. The way in which it was presented to the subjects in this study is as a direct cost of the use. You can order the tokens to use and pay a constant unit price for them. This change may make the decision of the agents easier by suggesting a comparison of the price of tokens with the return from the market.

References

Andreoni, J., 1995. Cooperation in public goods experiments: kindness or confusion? American Economic Review 4, 891–904.

Beckenkamp, M., Ostmann, A., 1999. Missing the target? Sanctioning as an ambiguous structural solution. In: Foddy, M., Smithson, M., Schneider, S., Hogg, M. (Eds.), Resolving Social Dilemmas. Taylor & Francis, Philadelphia, pp. 165–180.

Bolton, G.E., Ockenfels, A., 2000. ERC: a theory of equity, reciprocity and competition. American Economic Review 90 (1), 166–193.

Brandts, J., Schram, A., 2001. Cooperation and noise in public goods experiments: applying the contribution function approach. Journal of Public Economics 79 (2), 399–427.

Cardenas, J.C., Stranlund, J., Willis, C., 2000. Local environmental control and institutional crowding-out. World Development 28 (10), 1719–1733.

Casari, M., 2000. Emergence of endogenous legal institutions: the rural charters in northern Italy. Social Science working paper no. 1105, California Institute of Technology, Pasadena, November.

Chan, K.S., Godby, R., Mestelman, S., Muller, R.A., 1997. Equity theory and the voluntary provision of public goods. Journal of Economic Behavior and Organization 32, 349–364.

Clark, Colin, 1976. Mathematical Bioeconomics: The Optimal Management of Renewable Resources. Wiley, New York.

Falk, A., Fehr, E., Fischbacher, U., 2000. Understanding behavior in common-pool resource games. Working paper, Institute for Empirical Research in Economics, University of Zurich.

Fehr, E., Gächter, S., 2000. Cooperation and punishment in public goods experiments. American Economic Review 90 (4), 980–994.

Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition and cooperation. Quarterly Journal of Economics 114 (3), 817–868.

Gordon, H.S., 1954. The economic theory of a common property resource: the fishery. Journal of Political Economy 62, 124–142.

Isaac, R.M., Walker, J.M., 1998. Nash as an organizing principle in the voluntary provision of public goods: experimental evidence. Experimental Economics 1, 191–206.

Isaac, R.M., Walker, J.M., Williams, A.W., 1994. Group size and the voluntary provision of public goods: experimental evidence utilizing large groups. Journal of Public Economics 54 (1), 1–36.

Ito, M., Saijo, T., Une, M., 1995. The tragedy of the commons revisited: identifying behavioral principles. Journal of Economic Behavior and Organization 28, 311–335.
Krebs, D.L., 1970. Altruism: an examination of the concept and a review of the literature. Psychological Bulletin 73 (4), 258–302.
Ledyard, J., 1995. Public goods: a survey of experimental research. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, pp. 111–194.
Levine, D.K., 1998. Modeling altruism and spitefulness in experiments. Review of Economic Dynamics 1, 593–622.
McCusker, C., Carnevale, P.J., 1995. Framing in resource dilemmas: loss aversion and the moderating effects of sanctions. Organizational Behavior and Human Decision Processes 61 (2), 190–201.
Moir, R., 1999. Spies and swords: behavior in environments with costly monitoring and sanctioning. Manuscript, Department of Economics, University of New Brunswick, Canada.
Mui, V.-L., 1995. The economics of envy. Journal of Economic Behavior and Organization 26, 311–336.
Noussair, C.N., Plott, C.R., Riezman, R.G., 1995. An experimental investigation of the pattern of international trade. American Economic Review 85 (3), 462–491.
Ostrom, E., 1990. Governing the Commons: The Evolution of Institutions for Collective Action. Cambridge University Press, New York.
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., 1994. Rules, Games and Common-Pool Resources. University of Michigan, Ann Arbor.
Palfrey, T.R., Prisbey, J.E., 1997. Anomalous behavior in public goods experiments: how much and why? American Economic Review 87 (5), 829–846.
Rabin, M., 1993. Incorporating fairness into game theory and economics. American Economic Review 83 (5), 1281–1302.
Rocco, E., Warglien, M. 1996. Computer mediated communication and the emergence of "electronic opportunism", Working paper 1996-01, University of Trento.
Saijo, T. Spiteful behavior in voluntary contribution mechanism experiments. In: Plott, C.R., Smith, V.L. (Eds.), Handbook of Results in Experimental Economics. North-Holland, Amsterdam, in preparation.
Van Winden, F., van Dijk, F., Sonnemans, J. Intrinsic motivation in a public good environment. In: Plott, C.R., Smith, V.L. (Eds.), Handbook of Results in Experimental Economics. North-Holland, Amsterdam, in preparation.
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.
Yamagishi, T., 1988. Seriousness of social dilemmas and the provision of a sanctioning system. Social Psychology Quarterly 51, 32–34.