PROCEEDINGS OF THE ROYAL SOCIETY | B BIOLOGICAL SCIENCES

Incentives and opportunism: from the carrot to the stick

Christian Hilbe and Karl Sigmund

Supplementary data	"Data Supplement" http://rspb.royalsocietypublishing.org/content/suppl/2010/04/02/rspb.2010.0065.DC1.html
References	**This article cites 35 articles, 5 of which can be accessed free** http://rspb.royalsocietypublishing.org/content/277/1693/2427.full.html#ref-list-1
Subject collections	Articles on similar topics can be found in the following collections evolution (1819 articles)
Email alerting service	Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

To subscribe to *Proc. R. Soc. B* go to: **http://rspb.royalsocietypublishing.org/subscriptions**

PROCEEDINGS OF THE ROYAL SOCIETY B

Incentives and opportunism: from the carrot to the stick

Christian Hilbe[1,*] and Karl Sigmund[1,2]

[1]*Faculty of Mathematics, University of Vienna, Nordbergstrasse 15, A-1090 Vienna, Austria*
[2]*IIASA, A-2361 Laxenburg, Austria*

Cooperation in public good games is greatly promoted by positive and negative incentives. In this paper, we use evolutionary game dynamics to study the evolution of opportunism (the readiness to be swayed by incentives) and the evolution of trust (the propensity to cooperate in the absence of information on the co-players). If both positive and negative incentives are available, evolution leads to a population where defectors are punished and players cooperate, except when they can get away with defection. Rewarding behaviour does not become fixed, but can play an essential role in catalysing the emergence of cooperation, especially if the information level is low.

Keywords: evolutionary game theory; cooperation; reward; punishment; reputation

1. INTRODUCTION

Social dilemmas are obstacles to the evolution of cooperation. Examples such as the Prisoner's Dilemma show that self-interested motives can dictate self-defeating moves, and thus suppress cooperation. Both positive and negative incentives (the carrot and the stick) can induce cooperation in a population of self-regarding agents (e.g. Olson 1965; Ostrom & Walker 2003; Sigmund 2007). The provision of such incentives is costly, however, and therefore raises a second-order social dilemma. This issue has been addressed in many papers, particularly for the case of negative incentives (e.g. Yamagishi 1986; Boyd & Richerson 1992; Fehr & Gächter 2002; Bowles & Gintis 2004; Gardner & West 2004; Walker & Halloran 2004; Nakamaru & Iwasa 2006; Carpenter 2007; Lehmann *et al.* 2007; Sefton *et al.* 2007; Kiyonari & Barclay 2008).

It is easily seen that the efficiency of the two types of incentive relies on contrasting and even complementary circumstances. Indeed, if most players cooperate, then it will be costly to reward them all, while punishing the few defections will be cheap: often, the mere threat of a sanction suffices (Boyd *et al.* 2003; Gächter *et al.* 2008). On the other hand, if most players defect, then punishing them all will be a costly enterprise, while rewarding the few cooperators will be cheap. Obviously, therefore, the best policy for turning a population of defectors into a population of cooperators would be to use the carrot first, and at some later point, the stick.

In the absence of a proper institution to implement such a policy, members of the population can take the job onto themselves. But what is *their* incentive to do so? It pays only if the threat of a punishment, or the promise of a reward, can turn a co-player from a defector into a cooperator. Hence, the co-players must be opportunistic, i.e. prone to be swayed by incentives.

In order to impress a co-player, the threat (or promise) of an incentive must be sufficiently credible. In the following model, we shall assume that the credibility is provided by the players' reputations, i.e. by their history, and thus assume several rounds of the game, not necessarily with the same partner (e.g. Sigmund *et al.* 2001; Fehr & Fischbacher 2003; Barclay 2006). Credibility could alternatively be provided by a verbal commitment, for example. Since mere talk is cheap, however, such commitments need to be convincing; ultimately, they must be backed-up by actions, and hence again rely on reputation. Whether a player obtains information about the co-players' previous actions from direct experience, or by witnessing them at a distance, or hearing about them through gossip, can be left open at this stage. In particular, we do not assume repeated rounds between the same two players, but do not exclude them either. Basically, the carrot or the stick will be applied after the cooperation, or defection, and hence are forms of targeted reciprocation (while conversely, of course, the promise to return good with good and bad with bad, can act as an incentive).

In the following, we present a simple game theoretic model to analyse the evolution of opportunism, and to stress the smooth interplay of positive and negative incentives. The model is based on a previous paper (Sigmund *et al.* 2001; see also Hauert *et al.* 2004), which analyses punishment and reward separately and which *presumes* opportunistic agents. Here, we show how such opportunistic agents evolve via social learning, and how first rewards, then punishment lead to a society dominated by players who cooperate, except when they expect that they can get away with defection. Rewards will not become stably established; but they can play an essential role in the transition to cooperation, especially if the information level is below a specific threshold. Whenever the benefit-to-cost ratio for the reward is larger than 1, the eventual demise of rewarders is surprising, since a homogeneous population of rewarding cooperators would obtain a higher pay-off than a homogeneous population of punishing cooperators. We first analyse the model by means of replicator dynamics, then by means of a stochastic learning

* Author for correspondence (christian.hilbe@univie.ac.at).

model based on the Moran process. Thus, both finite populations and the limiting case of infinite populations will be covered. In the discussion, we study the role of errors, compare our results with experiments and point out the need to consider a wider role for incentives.

2. THE MODEL

Each round of the game consists of two stages—a helping stage and an incentive stage. Individuals in the population are randomly paired. A die decides who plays the role of the (potential) donor, and who is the recipient. In the first stage, donors may transfer a benefit b to their recipients, at their own cost c, or they may refuse to do so. These two alternatives are denoted by **C** (for cooperation) and **D** (for defection), respectively. In the second stage, recipients can reward their donors, or punish them, or refuse to react. If rewarded, donors receive an amount β; if punished, they must part with that amount β; in both cases, recipients must pay an amount γ, since both rewarding and punishing is costly. As usual, we assume that $c < b$, as well as $c < \beta$ and $\gamma < b$. For convenience, the same parameter values β and γ are used for both types of incentives: basically, all that matters are the inequalities. They ensure that donors are better off by choosing **C**, if their recipients use an incentive; and that in the case of rewards, both players have a positive pay-off. But material interests speak against using incentives as they are costly; and in the absence of incentives, helping behaviour will not evolve.

The four possible moves for the second stage will be denoted by **N**, to do nothing; **P**, to punish defection; **R**, to reward cooperation; and **I**, to provide for both types of incentives, i.e. to punish defection *and* to reward cooperation. For the first stage, next to the two unconditional moves **AllC**, to always cooperate, and **AllD**, to always defect, we also consider the opportunistic move: namely to defect except if prodded by an incentive. We shall, however, assume that information about the co-player may be incomplete. Let μ denote the probability to know whether the co-player provides an incentive or not, and set $\bar{\mu} = 1 - \mu$. We consider two types of opportunists, who act differently under uncertainty: players of type $\mathbf{O_C}$ defect only if they know that their co-player provides no incentive, and players of type $\mathbf{O_D}$ defect except if they know that an incentive will be delivered. Hence in the absence of information, $\mathbf{O_C}$ players play **C** and $\mathbf{O_D}$-players **D**. This yields 16 strategies, each given by a pair $[i, j]$, with $i \in M_D := \{\mathbf{AllC}, \mathbf{O_C}, \mathbf{O_D}, \mathbf{AllD}\}$ specifying how the player acts as a donor and $j \in M_R := \{\mathbf{N}, \mathbf{P}, \mathbf{R}, \mathbf{I}\}$ showing how the player acts as a recipient. If player I is donor and player II recipient, the pair (p_I, p_{II}) of their pay-off values is determined by their moves in the corresponding roles. Hence, we can describe these pairs using a 4×4 matrix $(a_{[ij]}, b_{[ij]})$ given by:

This specifies the pay-off values for the corresponding symmetrized game, which is given by a 16×16-matrix. A player using $[i, j]$ against a player using $[k, l]$ is with equal probability in the role of the donor or the recipient, and hence obtains as pay-off $(a_{[i,l]} + b_{[k,j]})/2$. The state of the population $\mathbf{x} = (x_{[i,j]})$ is given by the frequencies of the 16 strategies.

A wealth of possible evolutionary dynamics exists, describing how the frequencies of the strategies change with time under the influence of social learning (Hofbauer & Sigmund 1998). We shall consider only one updating mechanism, but stress that the results hold in many other cases too. For the learning rule, we shall use the familiar Moran-like 'death–birth' process (Nowak 2006): we thus assume that occasionally, players can update their strategy by copying the strategy of a 'model', i.e. a player chosen at random with a probability which is proportional to that player's fitness. This fitness in turn is assumed to be a convex combination $(1 - s)B + sP$, where B is a 'baseline fitness' (the same for all players), P is the pay-off (which depends on the model's strategy, and the state of the population), and $0 \le s \le 1$ measures the 'strength of selection', i.e. the importance of the game for overall fitness. (We shall always assume s to be small enough to avoid negative fitness values). This learning rule corresponds to a Markov process. The rate for switching from the strategy $[k, l]$ to the strategy $[i, j]$ is $(1 - s)B + sP_{[i,j]}$, independent of $[k, l]$.

(a) *Large populations*

The learning rule leads, in the limiting case of an infinitely large population, to the replicator equation for the relative frequencies $x_{[ij]}$: the growth rate of any strategy is given by the difference between its pay-off and the average pay-off in the population (Hofbauer & Sigmund 1998). This yields an ordinary differential equation that can be analysed in a relatively straightforward way, despite being 15-dimensional.

Let us first note that **I** is weakly dominated by **P**, in the sense that **I**-players never do better, and sometimes less well, than **P**-players. Hence, no state where all the strategies are played can be stationary. The population always evolves towards a region where at least one strategy is missing. Furthermore, **AllC** is weakly dominated by $\mathbf{O_C}$, and **AllD** by $\mathbf{O_D}$. This allows to reduce the dynamics to lower-dimensional cases. Of particular relevance are the states where only two strategies are present, and where these two strategies prescribe the same move in one of the two stages of the game. The outcome of such pairwise contests is mostly independent of the parameter values, with three exceptions:

(a) In a homogeneous $\mathbf{O_C}$-population, **R** dominates **N**, if and only if $\mu > \gamma/b$;

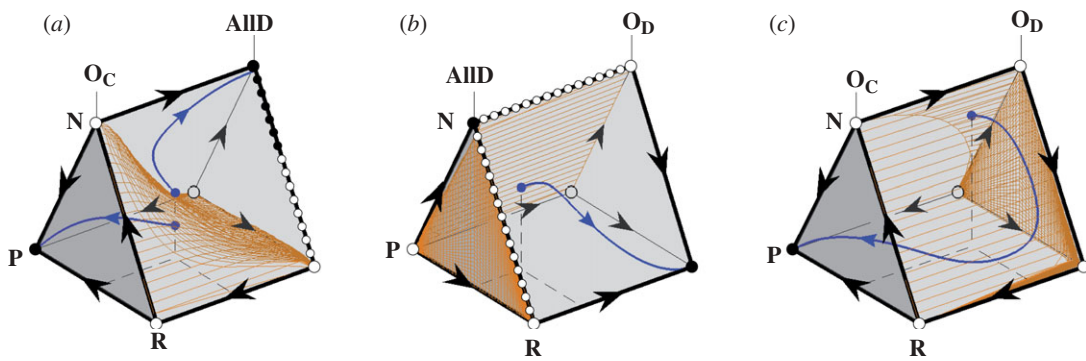| *	N	P	R	I
AllC	$(-c, b)$	$(-c, b)$	$(\beta - c, b - \gamma)$	$(\beta - c, b - \gamma)$
$\mathbf{O_C}$	$(-\bar{\mu}c, \bar{\mu}b)$	$(-c, b)$	$(\beta - c, b - \gamma)$	$(\beta - c, b - \gamma)$
$\mathbf{O_D}$	$(0, 0)$	$(-\mu c - \bar{\mu}\beta, \mu b - \bar{\mu}\gamma)$	$(\mu(\beta - c), \mu(b - \gamma))$	$(-(1 - 2\mu)\beta - \mu c, \mu b - \gamma)$
AllD	$(0, 0)$	$(-\beta, -\gamma)$	$(0, 0)$	$(-\beta, -\gamma)$

Figure 1. Dynamics of a population consisting of (*a*) $\mathbf{O_C}$ and **AllD**, (*b*) **AllD** and $\mathbf{O_D}$, (*c*) $\mathbf{O_C}$ and $\mathbf{O_D}$ (resp.). Black circles represent Nash equilibria, white circles indicate unstable fixed points. The arrows on the edges indicate the direction of the dynamics if only the two strategies corresponding to the endpoints are present. The orange grid is the manifold that separates initial values with different asymptotic behaviour. The blue curves represent the typical dynamics for a given initial population. Parameter values: $b = 4$, $c = 1$, $\beta = \gamma = 2$ and $\mu = 30$ per cent (hence $\gamma/(2b) < \mu < \gamma/(\gamma + b)$).

(b) In a homogeneous $\mathbf{O_D}$-population, **P** dominates **N**, if and only if $\mu > \gamma/(b + \gamma)$;

(c) In a homogeneous $\mathbf{O_D}$-population, **P** dominates **R**, if and only if $\mu > 1/2$.

In each case, it is easy to understand why higher reputation will have the corresponding effect. Owing to our assumption $\gamma < b$, all these thresholds for μ lie in the open interval $(0, 1)$.

One can obtain a good representation of the dynamics by looking at the situations where there are two alternatives for the first stage (namely **AllD** and $\mathbf{O_C}$, or **AllD** and $\mathbf{O_D}$, or $\mathbf{O_C}$ and $\mathbf{O_D}$), and three alternatives **N**, **P** and **R** for the second stage. In each such case, the state space of the population can be visualized by a prism (figure 1). Here, each of its 'square faces' stands for the set of all mixed populations with only four strategies present. For instance, if the population consists only of the four strategies [$\mathbf{O_C}$, **N**], [$\mathbf{O_C}$, **R**], [$\mathbf{O_D}$, **N**], and [$\mathbf{O_D}$, **R**], then the state corresponds to a point in the three-dimensional simplex spanned by the corresponding four monomorphic populations. But since the double ratios $x_{[ij]}x_{[kl]}/x_{[il]}x_{[kj]}$ are invariant under the replicator dynamics (see Hofbauer & Sigmund 1998, pp. 122–125), the state cannot leave the corresponding two-dimensional surface, which may be represented by a square (figure 2).

For several pairs of strategies (such as [$\mathbf{O_C}$, **P**] and [**AllC**, **P**], or [**AllD**, **N** and [$\mathbf{O_D}$, **N**]), all populations which are mixtures of the corresponding two strategies are stationary. There is no selective force favouring one strategy over the other. We shall assume that in this case, small random shocks will cause the state to evolve through neutral drift. This implies that evolution then leads ultimately to [$\mathbf{O_C}$, **P**], and hence to a homogeneous population that stably cooperates in the most efficient way. Indeed, it is easy to see that no other strategy can invade a monomorphic [$\mathbf{O_C}$, **P**]-population through selection. The only flaw is that [**AllC**, **P**] can enter through neutral drift. Nevertheless, [$\mathbf{O_C}$, **P**] is a Nash equilibrium.

But how can [$\mathbf{O_C}$, **P**] get off the ground? Let us first consider what happens if the possibility to play **R**, i.e. to reward a cooperative move, is excluded. The asocial strategy [**AllD**, **N**] is stable. It can at best be invaded through neutral drift by [$\mathbf{O_D}$, **N**]. If $\mu > \gamma/(b + \gamma)$, this can in turn
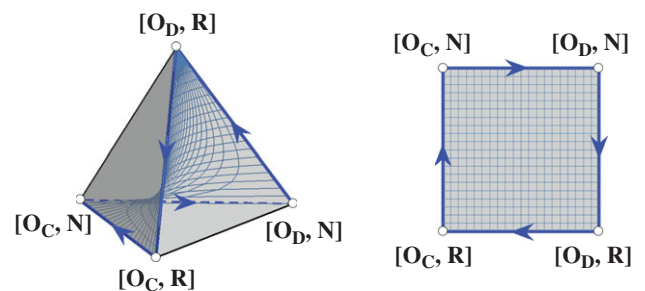

Figure 2. The state space of a game involving the four strategies [$\mathbf{O_C}$, **N**], [$\mathbf{O_C}$, **R**], [$\mathbf{O_D}$, **N**] and [$\mathbf{O_D}$, **R**]. The corners of the three-dimensional simplex correspond to the homogeneous populations using that strategy, the interior points denote mixed populations. For each initial state, the evolution of the system is restricted to a two-dimensional saddle-like manifold that can be represented by a square (right). If $\mu < (\gamma/b)$, the competition between these four strategies is characterized by a rock–paper–scissors-like dynamics, as indicated by the orientation of the edges.

be invaded by [$\mathbf{O_D}$, **P**], which then leads to [$\mathbf{O_C}$, **P**]. If μ is smaller, however, that path is precluded and the population would remain in an uncooperative state. It is in this case that the **R**-alternative plays an essential role. By neutral drift, [**AllD**, **R**] can invade [**AllD**, **N**]. More importantly, [$\mathbf{O_D}$, **R**] dominates [$\mathbf{O_D}$, **N**], [**AllD**, **R**] and [**AllD**, **N**]. From [$\mathbf{O_D}$, **R**], the way to [$\mathbf{O_C}$, **R**] and then to [$\mathbf{O_C}$, **P**] is easy.

The essential step of that evolution occurs in the transition from $\mathbf{O_D}$ to $\mathbf{O_C}$, when players start cooperating by default, i.e. in the absence of information (see the third column in figure 1). If the **R**-alternative is not available, then for small values of μ, the population can be trapped in [$\mathbf{O_D}$, **N**]. But if the **R**-alternative can be used, it can switch from [$\mathbf{O_D}$, **N**] to [$\mathbf{O_D}$, **R**]. In a population where the first move is either $\mathbf{O_D}$ or $\mathbf{O_C}$, and the second move either **N** or **R**, there is a (four-membered) rock–paper–scissors cycle (figure 2): one strategy is superseded by the next. A unique stationary state exists where these four alternatives are used. We show in the electronic supplementary material that for $\mu < \gamma/2b$, this stationary state cannot be invaded by any strategy using **P**. But due to the rock–paper–scissors dynamics, it is inherently unstable. The population will eventually use the strategy

[$\mathbf{O_C}$, $\mathbf{R}$] mostly. There, the strategy [$\mathbf{O_C}$, $\mathbf{P}$] can invade and become fixed.

In the competition between [$\mathbf{O_D}$, $\mathbf{N}$] and [$\mathbf{O_C}$, $\mathbf{P}$], the latter is dominant if and only if $\mu > (c + \gamma)/(c + \gamma + b)$ (a condition which is independent of β). If not, then the competition is bistable, meaning that neither strategy can invade a homogeneous population adopting the other strategy. An equal mixture of both strategies converges to the pro-social strategy [$\mathbf{O_C}$, $\mathbf{P}$] if and only if $\mu(\beta - 2c - 2b - \gamma) < \beta - 2c - \gamma$. In the case $\gamma = \beta$, this simply reduces to $\mu > c/(c + b)$.

We thus obtain a full classification of the replicator dynamics in terms of the parameter μ. The main bifurcation values are $(\gamma/2b) < \gamma/(b + \gamma) < \gamma/b$ and $1/2$. These can be arranged in two ways, depending on whether $b < 2\gamma$ or not. But the basic outcome is the same in both cases (see figure 1 and the electronic supplementary material).

It is possible to modify this model by additionally taking into account the recombination of the traits affecting the first and the second stage of the game. Indeed, recombination does not only occur for genetic transmission of strategies, but also for social learning. A modification of an argument from Gaunersdorfer *et al.* (1991) show that, in this case, the double ratios $x_{[ij]}x_{[kl]}/x_{[il]}x_{[kj]}$ converge to 1, so that the traits for the first and the second stage of the game become statistically independent of each other. Hence, the previous analysis still holds. In Lehmann *et al.* (2007) and Lehmann & Rousset (2009) it is shown, in contrast, that recombination greatly affects the outcome in a lattice and in a finite population model without reputational effects.

(b) *Small mutation rates*
In the case of a finite population of size M, the learning process corresponds to a Markov chain on a state space which consists of frequencies of all the strategies (which sum up to M). The absorbing states correspond to the homogeneous populations: in such a homogeneous population, imitation cannot introduce any change. If we add to the learning process a 'mutation rate' (or more precisely, an exploration rate), by assuming that players can also adopt a strategy by chance, rather than imitation, then the corresponding process is recurrent (a chain of transitions can lead from every state to every other) and it admits a unique stationary distribution. This stationary distribution describes the frequencies of the states in the long run. It is in general laborious to compute, since the number of possible states grows polynomially in M. However, in the limiting case of a very small exploration rate (the so-called adiabatic case), we can assume that the population is mostly in a homogeneous state, and we can compute the transition probabilities between these states (Nowak 2006). This limiting case is based on the assumption that the fate of the mutant (i.e. whether it will be eliminated or fixed in the population) is decided before the next mutation occurs. We can confirm the results from the replicator dynamics. For simplicity, we confine ourselves to the non-dominated strategies $\mathbf{O_C}$, $\mathbf{O_D}$ (resp. $\mathbf{N}$, $\mathbf{P}$ and $\mathbf{R}$); similar results can be obtained by considering the full strategy space.

In the stationary distribution, the population is dominated by the strategy [$\mathbf{O_C}$, $\mathbf{P}$], but for smaller values of μ, it needs the presence of the $\mathbf{R}$-alternative to emerge. This becomes particularly clear if one looks at the transition probabilities (see electronic supplementary material). Except for large values of μ, only the strategy [$\mathbf{O_D}$, $\mathbf{R}$] can invade the asocial [$\mathbf{O_D}$, $\mathbf{N}$] with a fixation probability, which is larger than the neutral fixation probability $1/M$.

If [$\mathbf{O_C}$, $\mathbf{P}$] dominates [$\mathbf{O_D}$, $\mathbf{N}$], or when it fares best in an equal mixture of both strategies, then it needs not the help of $\mathbf{R}$-players to become the most frequent strategy in the long run (i.e. in the stationary distribution). But for smaller values of μ, rewards are essential. In figure 3, it is shown that the existence of rewarding strategies allows the social strategy [$\mathbf{O_C}$, $\mathbf{P}$] to supersede the asocial [$\mathbf{O_D}$, $\mathbf{N}$] even in cases in which the players have hardly any information about their co-players. The time until the system leaves [$\mathbf{O_D}$, $\mathbf{N}$] is greatly reduced if rewarding is available (see figure 4). In the electronic supplementary material it is shown that the state [$\mathbf{O_C}$, $\mathbf{P}$] is usually reached from [$\mathbf{O_C}$, $\mathbf{R}$], while the strategy most likely to invade the asocial [$\mathbf{O_D}$, $\mathbf{N}$] is [$\mathbf{O_D}$, $\mathbf{R}$]. These outcomes are robust, and depend little on the parameter choices. Moreover, they are barely affected by the mutation structure. If, instead of assuming that all mutations are equally likely, we only allow for mutations in the behaviour in one of the two stages (i.e. no recombination between the corresponding traits), the result is very similar. Apparently, if it is impossible to mutate directly from [$\mathbf{O_D}$, $\mathbf{N}$] to [$\mathbf{O_C}$, $\mathbf{P}$], then the detour via [$\mathbf{O_D}$, $\mathbf{P}$] works almost as well.

Even for the limiting case $\mu = 0$ (no reputation effects), the role of rewards is strongly noticeable. Without rewards, the stationary probability of the asocial strategy [$\mathbf{O_D}$, $\mathbf{N}$] is close to 100 per cent; with the possibility of rewards, it is considerably reduced.

3. DISCUSSION
We have analysed a two-person, two-stages game. It is well-known that it corresponds to a simplified version of the ultimatum game (Güth *et al.* 1982), in the punishment case, or of the trust game (Berg *et al.* 1995), in the reward case (De Silva & Sigmund 2009; Sigmund 2010). Similar results also hold for the N-person public good game with reward and punishment (e.g. Hauert *et al.* 2004). However, the many-person game offers a wealth of variants having an interest of their own (as, for instance, when players decide to mete out punishment only if they have a majority on their side; see Boyd *et al.* submitted). In this paper, we have opted for the simplest set-up and considered pairwise interactions only.

In classical economic thought, positive and negative incentives have often been treated on an equal footing, so to speak (Olson 1965). In evolutionary game theory, punishing is studied much more frequently than rewarding. The relevance of positive incentives is sometimes queried, on the ground that helping behaviour makes only sense if there is an asymmetry in resource level between donor and recipient. If A has a high pile of wood, and B has no fuel, A can give some wood away at little cost, and provide a large benefit to B. This is the cooperative act. Where is the positive incentive? It would be absurd to imagine that B rewards A by returning the fuel. But B can reward A by donating some other resource, such as food, or fire, which A is lacking.
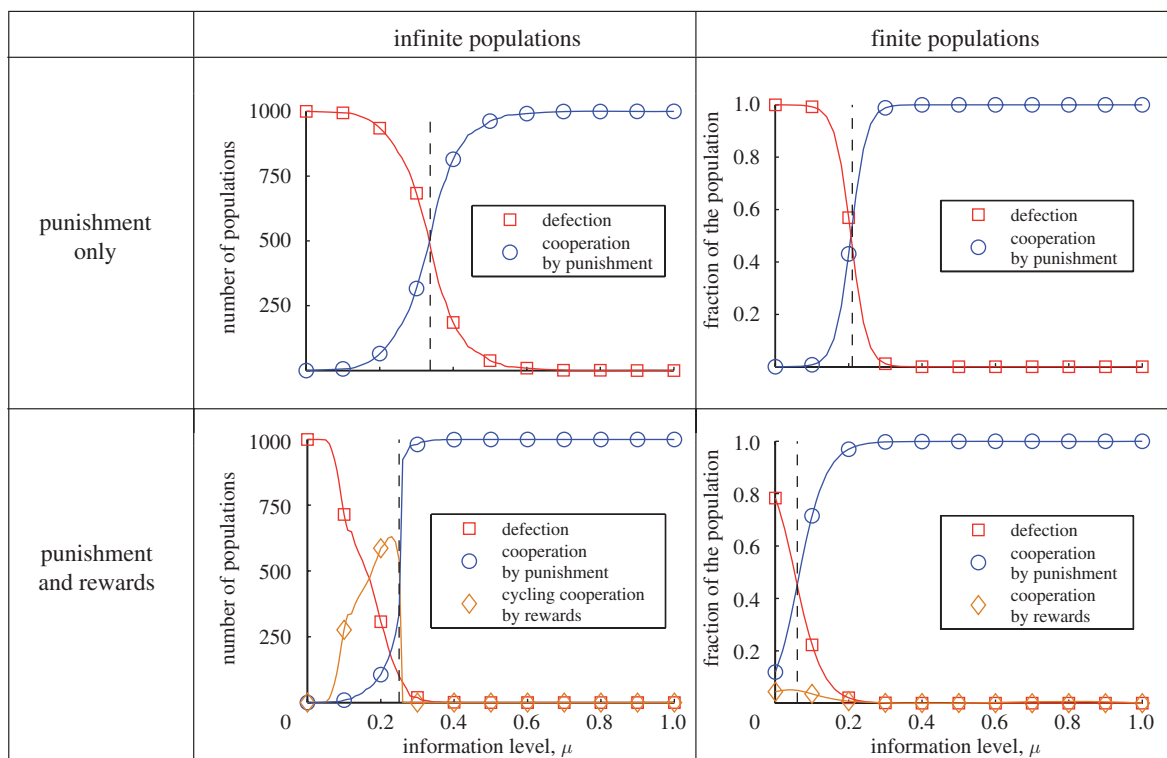
Figure 3. Strategy selection in finite and infinite populations, depending on the information parameter μ. The left column shows the outcome of a simulation of the replicator equation for 1000 randomly chosen initial populations. If only punishment is available to sway opportunistic behaviour, then cooperative outcomes become more likely if μ exceeds roughly 1/3 (in which case $[\mathbf{O_C}, \mathbf{P}]$ becomes fixed). As soon as rewards are also allowed, punishment-enforced cooperation becomes predominant as soon as $\mu > \gamma/2b = 1/4$. Additionally, for smaller values of μ, the population may tend to cycle between the strategies $[\mathbf{O_C}, \mathbf{R}]$, $[\mathbf{O_C}, \mathbf{N}]$, $[\mathbf{O_D}, \mathbf{N}]$ and $[\mathbf{O_D}, \mathbf{R}]$, represented by the orange line in the lower left graph. The right column shows the stationary distribution of strategies in a finite population. Again, without rewards a considerably higher information level μ is necessary to promote punishment-enforced cooperation (either $[\mathbf{O_C}, \mathbf{P}]$ or $[\mathbf{O_D}, \mathbf{P}]$; note that both opportunist strategies become indistinguishable in the limit case of complete information). In finite populations, rewarding strategies act merely as a catalyst for the emergence of punishment; even for small μ, the outcomes $[\mathbf{O_C}, \mathbf{R}]$ (resp. $[\mathbf{O_D}, \mathbf{R}]$) never prevail. Parameter values: $b = 4$, $c = 1$, $\beta = \gamma = 2$. For finite populations, the population size is $M = 100$ and the selection strength $s = 1/10$.
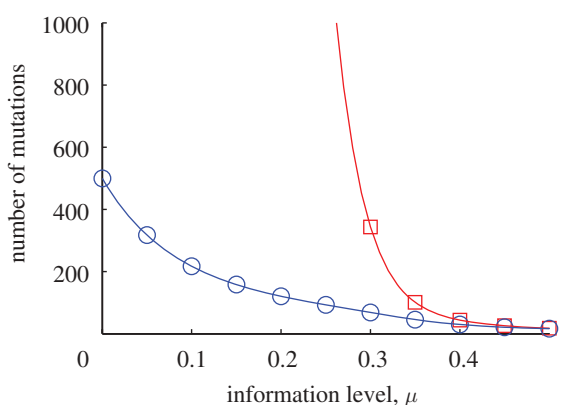

Figure 4. Average number of mutations needed until a population of $[\mathbf{O_D}, \mathbf{N}]$ players is successfully invaded. Adding the possibility of rewards reduces the waiting time considerably (for $\mu = 0\%$ it takes 500 mutations with rewards and almost 500 000 without). As the information level increases, this catalytic effect of rewarding disappears. Parameter values: population size $M = 100$, selection strength $s = 1/10$; $b = 4$, $c = 1$, $\beta = \gamma = 2$. Squares, $\mathbf{P}$ and $\mathbf{N}$ only; circles, $\mathbf{P}$, $\mathbf{N}$, $\mathbf{R}$.

In experimental economics, punishing behaviour has been studied considerably more often than rewarding behaviour (Yamagishi 1986; Fehr & Gächter 2002; Barclay 2006; Dreber *et al.* 2008). In the last few years, there

has been a substantial amount of empirical work on the interplay of the two forms of incentives (Andreoni *et al.* 2003; Rockenbach & Milinski 2006; Sefton *et al.* 2007). The results, with two exceptions to be discussed presently, confirm our theoretical conclusion: punishment is the more lasting factor, but the combination of reward and punishment works best. This outcome is somewhat surprising, because in most experiments, players are anonymous and know that they cannot build up a reputation. One significant exception is the investigation, in Fehr & Fischbacher (2003), of the ultimatum game, which has essentially the same structure as our two-stage game with punishment. In that case, the treatment without information on the co-player's past behaviour yields a noticeably lower level of cooperation than the treatment with information. Nevertheless, even in the no-information treatment, both the level of cooperation (in the form of fair sharing) and of punishment (in the form of rejection of small offers) are remarkably high.

A serious criticism of the model presented in this paper is thus that it does not seem to account for the pro-social behaviour shown by players who know that reputation-building is impossible. We believe that this effect is owing to a maladaptation. Our evolutionary past has not prepared us to expect anonymity. In hunter–gatherer societies and in rural life, it is not often that one can really

be sure of being unobserved. Even in modern life, the long phase of childhood is usually spent under the watchful eyes of parents, educators or age-peers. Ingenious experiments uncover our tendency to over-react to the slightest cues indicating that somebody may be watching (for instance, the mere picture of an eye; see Haley & Fessler (2005) and Bateson *et al.* (2006), or three dots representing eyes and mouth; see Rigdon *et al.* (2009)). The idea of personal deities scrutinizing our behaviour, which seems to be almost universal, is probably a projection of this deep-seated conviction (Johnson & Bering 2006). The concept of conscience was famously described, by Mencken, as 'the inner voice that warns us somebody may be looking' (cf. Levin 2009).

In several experimental papers, however, the role of reputation is very explicit. In Rand *et al.* (2009), players are engaged in 50 rounds of the public goods game with incentives, always with the same three partners. Hence, they know the past actions of their co-players. In this case, we can be sure that $\mu > (\gamma/b)$. Thus, in a homogeneous O_C-population, **R** should dominate **N**. Moreover, as the leverage for both punishment and reward is $1:3$ in this experiment (as in many others), an [O_C, **R**]-population obtains a pay-off $b - c + \beta - \gamma$, which is substantially larger than that of an [O_C, **P**]-population. In the experiment, rewarding performs indeed much better than punishing, and Rand *et al.* conclude that 'positive reciprocity should play a larger role than negative reciprocity in maintaining public cooperation in repeated situations.'

Nevertheless, according to our model, **P**-players ought to invade. This seems counter-intuitive. Punishers do not have to pay for an incentive (since everyone cooperates), but they will nevertheless be rewarded, since they cooperate in the public goods stage. Thus [O_C, **P**] should take over, thereby lowering the average pay-off. By contrast, in the repeated game considered by Rand *et al.*, it is clear that cooperative players who have not been rewarded by their co-player in the previous round will feel cheated, and stop rewarding that co-player. They will not be impressed by the fact that the co-player is still providing an incentive by punishing defectors instead. In other words, in this experiment rewards are not only seen as incentives, but as contributions in their own right, in a Repeated Prisoner's Dilemma game. Players will reciprocate not only for the public goods behaviour, but for the 'mutual reward game' too. In fact, if there had been two players only in the experiment by Rand *et al.*, it would reduce to a repeated Prisoner's Dilemma game with 100 rounds.

This aspect is not covered in our model, where the incentives are only triggered by the behaviour in the public goods stage, but not by previous incentives. In particular, rewarding behaviour cannot be rewarded, and fines do not elicit counter-punishment. This facilitates the analysis of incentives as instruments for promoting cooperation, but it obscures the fact that in real life, incentives have to be viewed as economic exchanges in their own right.

A similar experiment as in Rand *et al.* was studied by Milinski *et al.* (2002), where essentially the public goods rounds alternate with an indirect reciprocity game (see also Panchanathan & Boyd 2006). Helping, in such an indirect reciprocity game, is a form of reward. In

Milinski's experiment, punishment was not allowed, but in Rockenbach & Milinski (2006), both types of incentives could be used. Groups were rearranged between rounds, as players could decide whether to leave or to stay. Players knew each other's past behaviour in the previous public goods rounds and the indirect reciprocity rounds (but not their punishing behaviour). It was thus possible to acquire reputation as a rewarder, but not as a punisher. This treatment usually led to a very cooperative outcome, with punishment focused on the worst cheaters, and a significant interaction between reward and punishment.

In our numerical examples, we have usually assumed $\gamma = \beta$, but stress that this does not affect the basic outcome (see electronic supplementary material for the case $\gamma < \beta$). In most experiments, the leverage of the incentive is assumed to be stronger. Clearly, this encourages the recipients to use incentives (Carpenter 2007; Egas & Riedl 2007; Vyrastekova & van Soest 2008). But it has been shown (Carpenter 2007; Sefton *et al.* 2007) that many are willing to punish exploiters even if it reduces their own account by as much as that of the punished player. In the trust game, it is also usually assumed that the second stage is a zero-sum game. In most of the (relatively few) experiments on rewarding, the leverage is $1:1$ (Walker & Halloran 2004; Sefton *et al.* 2007), in Rockenbach & Milinski and Rand *et al.*, it is $1:3$. In Vyrastekova & van Soest (2008), it is shown that increasing this leverage makes rewarding more efficient. In our view, it is natural to assume a high benefit-to-cost ratio in the first stage (the occasion for a public goods game is precisely the situation when mutual help is needed), but it is less essential that a high leverage also applies in the second stage. Punishment, for instance, can be very costly if the other player retaliates, as seems quite natural to expect (at least in pairwise interactions; in N-person games, sanctions can be inexpensive if the majority punishes a single cheater).

For the sake of simplicity, we have not considered the probability of errors in implementation. But it can be checked in a straightforward manner that the results are essentially unchanged if we assume that with a small probability $\varepsilon > 0$, an intended donation fails (either because of a mistake of the player, or to unfavourable conditions). The other types of error in implementation (namely helping without wanting it) seem considerably less plausible. We note that in a homogeneous [O_C, **P**]-population, usually there is no need to punish co-players, and hence no way of building up a reputation as a punisher. But if errors in implementation occur, there will be opportunities for punishers to reveal their true colours. In Sigmund (2010), it is shown that if there are sufficiently many rounds of the game, occasional errors will provide enough opportunities for building up a reputation.

REFERENCES

Andreoni, J., Harbaugh, W. & Vesterlund, L. 2003 The carrot or the stick: rewards, punishments, and cooperation. *Am. Econ. Rev.* **93**, 893–902.

Barclay, P. 2006 Reputational benefits for altruistic punishment. *Evol. Hum. Behav.* **27**, 325–344. (doi:10.1016/j.evolhumbehav.2006.01.003)

Bateson, M., Nettle, D. & Roberts, G. 2006 Cues of being watched enhance cooperation in a real-world setting. *Biol. Lett.* **2**, 412–414. (doi:10.1098/rsbl.2006.0509)

Berg, J., Dickhaut, J. & McCabe, K. 1995 Trust, reciprocity, and social history. *Games Econ. Behav.* **10**, 122–142. (doi:10.1006/game.1995.1027)

Bowles, S. & Gintis, H. 2004 The evolution of strong reciprocity: cooperation in heterogeneous populations. *Theor. Popul. Biol.* **65**, 17–28. (doi:10.1016/j.tpb.2003.07.001)

Boyd, R. & Richerson, P. J. 1992 Punishment allows the evolution of cooperation (or anything else) in sizeable groups. *Ethol. Sociobiol.* **13**, 171–195. (doi:10.1016/0162-3095(92)90032-Y)

Boyd, R., Gintis, H., Bowles, S. & Richerson, P. J. 2003 The evolution of altruistic punishment *Proc. Natl Acad. Sci. USA* **100**, 3531–3535. (doi:10.1073/pnas.0630443100)

Boyd, R., Gintis, H. & Bowles, S. Submitted. Coordinated contingent punishment is group-beneficial and can proliferate when rare.

Carpenter, J. P. 2007 The demand for punishment. *J. Econ. Behav. Organ.* **62**, 522–542. (doi:10.1016/j.jebo.2005.05.004)

De Silva, H. & Sigmund, K. 2009 Public good games with incentives: the role of reputation. In *Games, groups and the global good* (ed. S. A. Levin), pp. 85–114. New York, NY: Springer.

Dreber, A., Rand, D. G., Fudenberg, D. & Nowak, M. A. 2008 Winners don't punish. *Nature* **452**, 348–351. (doi:10.1038/nature06723)

Egas, M. & Riedl, A. 2007 The economics of altruistic punishment and the maintenance of cooperation. *Proc. R. Soc. B* **275**, 871–878. (doi:10.1098/rspb.2007.1558)

Fehr, E. & Gächter, S. 2002 Altruistic punishment in humans. *Nature* **415**, 137–140. (doi:10.1038/415137a)

Fehr, E. & Fischbacher, U. 2003 The nature of human altruism. *Nature* **425**, 785–791. (doi:10.1038/nature02043)

Gächter, S., Renner, E. & Sefton, M. 2008 The long-run benefits of punishment. *Science* **322**, 1510. (doi:10.1126/science.1164744)

Gardner, A. & West, S. 2004 Cooperation and punishment, especially in humans. *Am. Nat.* **164**, 753–764. (doi:10.1086/425623)

Gaunersdorfer, A., Hofbauer, J. & Sigmund, K. 1991 On the dynamics of asymmetric games. *Theor. Popul. Biol.* **39**, 345–357. (doi:10.1016/0040-5809(91)90028-E)

Güth, W., Schmittberger, R. & Schwarze, B. 1982 An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* **3**, 367–388. (doi:10.1016/0167-2681(82)90011-7)

Haley, K. & Fessler, D. 2005 Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evol. Hum. Behav.* **26**, 245–256. (doi:10.1016/j.evolhumbehav.2005.01.002)

Hauert, C., Haiden, N. & Sigmund, K. 2004 The dynamics of public goods. *Discrete Contin. Dyn. Syst. B* **4**, 575–585.

Hofbauer, J. & Sigmund, K. 1998 *Evolutionary games and population dynamics*. Cambridge, UK: Cambridge University Press.

Johnson, D. & Bering, J. 2006 Hand of God, mind of man: punishment and cognition in the evolution of cooperation. *Evol. Psychol.* **4**, 21 923.

Kiyonari, T. & Barclay, P. 2008 Cooperation in social dilemmas: free riding may be thwarted by second-order

reward rather than by punishment. *J. Pers. Soc. Psychol.* **95**, 826–842.

Lehmann, L. & Rousset, F. 2009 Perturbation expansions of multilocus fixation probabilities for frequency-dependent selection with applications to the Hill–Robertson effect and to the joint evolution of helping and punishment. *Theor. Popul. Biol.* **76**, 35–51. (doi:10.1016/j.tpb.2009.03.006)

Lehmann, L., Rousset, F., Roze, D. & Keller, L. 2007 Strong reciprocity or strong ferocity? A population genetic view of the evolution of altruistic punishment. *Am. Nat.* **170**, 21–36. (doi:10.1086/518568)

Levin, S. A. (ed.) 2009 Preface to *games, groups, and the global good*. New York, NY: Springer.

Milinski, M., Semmann, D. & Krambeck, H. J. 2002 Reputation helps solve the 'tragedy of the commons'. *Nature* **415**, 424–426. (doi:10.1038/415424a)

Nakamaru, M. & Iwasa, Y. 2006 The coevolution of altruism and punishment: role of the selfish punisher. *J. Theor. Biol.* **240**, 475–488. (doi:10.1016/j.jtbi.2005.10.011)

Nowak, M. A. 2006 *Evolutionary dynamics*. Cambridge, MA: Harvard University Press.

Nowak, M. A., Page, K. & Sigmund, K. 2000 Fairness versus reason in the ultimatum game. *Science* **289**, 1773–1775. (doi:10.1126/science.289.5485.1773)

Olson, M. 1965 *The logic of collective action*. Cambridge, MA: Harvard University Press.

Ostrom, E. & Walker, J. 2003 *Trust and reciprocity: interdisciplinary lessons from experimental research*. New York, NY: Russel Sage Foundation.

Panchanathan, K. & Boyd, R. 2006 Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature* **432**, 499–502. (doi:10.1038/nature02978)

Rand, D. G., Dreber, A., Fudenberg, D., Ellingsen, T. & Nowak, M. A. 2009 Positive interactions promote public cooperation. *Science* **325**, 1272–1275. (doi:10.1126/science.1177418)

Rigdon, M., Ishii, K., Watabe, M. & Kitayama, S. 2009 Minimal social cues in the dictator game. *J. Econ. Psychol.* **30**, 358–367. (doi:10.1016/j.joep.2009.02.002)

Rockenbach, B. & Milinski, M. 2006 The efficient interaction of indirect reciprocity and costly punishment. *Nature* **444**, 718–723. (doi:10.1038/nature05229)

Sefton, M., Shupp, R. & Walker, J. M. 2007 The effects of rewards and sanctions in provision of public goods. *Econ. Inquiry* **45**, 671–690. (doi:10.1111/j.1465-7295.2007.00051.x)

Sigmund, K. 2007 Punish or perish? Retaliation and collaboration among humans. *Trends Ecol. Evol.* **22**, 593–600. (doi:10.1016/j.tree.2007.06.012)

Sigmund, K. 2010 *The calculus of selfishness*. Princeton NJ: Princeton University Press.

Sigmund, K., Hauert, C. & Nowak, M. A. 2001 Reward and punishment. *Proc. Natl Acad. Sci. USA* **98**, 10 757–10 762.

Vyrastekova, J. & van Soest, D. P. 2008 On the (in)effectiveness of rewards in sustaining cooperation. *Exp. Econ.* **11**, 53–65. (doi:10.1007/s10683-006-9153-x)

Walker, J. M. & Halloran, W. A. 2004 Rewards and sanctions and the provision of public goods in one-shot settings. *Exp. Econ.* **7**, 235–247. (doi:10.1023/B:EXEC.0000040559.08652.51)

Yamagishi, T. 1986 The provision of a sanctioning system as a public good. *J. Pers. Soc. Psychol.* **51**, 110–116.