ARTICLE IN PRESS

Available online at www.sciencedirect.com

ScienceDirect

Journal of Theoretical Biology ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Journal of Theoretical Biology

www.elsevier.com/locate/yjtbi

Selfish punishment: Altruism can be maintained by competition among cheaters

Omar Tonsi Eldakar[a,*], Dene Leo Farrell[b], David Sloan Wilson[a,c]

[a]*Department of Biological Sciences, Binghamton University, Binghamton, NY 13902-6000, USA*
[b]*Department of Bioengineering, Binghamton University, Binghamton, NY 13902-6000, USA*
[c]*Department of Anthropology, Binghamton University, Binghamton, NY 13902-6000, USA*

Received 22 August 2006; received in revised form 21 June 2007; accepted 23 July 2007

Abstract

Altruistic punishment refers to a class of behaviors that deters cheating at a cost to the punisher, making it a form of second-order altruism. Usually, it is assumed that the punishers are themselves "solid citizens" who refrain from cheating. We show in a simulation model that altruism and punishment paradoxically become negatively correlated, leading to a form of selfish punishment. Examples of selfish punishment can be found in organisms as diverse as wasps, birds, and humans.

Keywords: Altruism; Selfish; Punishment; Cooperation; Public goods

---

1. Introduction

Altruism is famously difficult to evolve because of potential exploitation by cheaters. Punishment can potentially deter cheating, but it often requires time, energy and risk. The term altruistic punishment refers to a class of behaviors that deters cheating at the expense of the punisher, qualifying as a form of second-order altruism in comparison to first-order altruists who do not punish (Gintis, 2000a; Fehr and Gachter, 2002; Fehr et al., 2002; Boyd et al., 2003; Fehr and Fischbacher, 2003; Gintis et al., 2003; Fehr, 2004). The cost of punishing cheaters, along with the cost of being cheated, make it difficult to explain altruistic punishment as an evolutionary stable strategy (Bowles and Gintis, 2002, 2004; Fehr, 2004). This report suggests another way that altruism can be maintained; by cheaters who punish other cheaters.

The concept of selfish punishment was suggested to us by an empirical study on humans showing that individuals most likely to punish cheaters were also most tempted to cheat (Eldakar et al., 2006). This seems hypocritical in moral terms but makes sense as a behavioral strategy because cheaters decrease the fitness of everyone in their groups, including other cheaters. A negative correlation between punishment and altruism might exist if cheaters have an even greater incentive than altruists to get rid of other cheaters. A few theoretical models have addressed this possibility (e.g. Nakamaru and Iwasa, 2006; Sigmund et al., 2001), but it needs to be explored more fully, especially in the context of the public goods games used by experimental economists to study the dynamics of cooperation in human social groups (e.g. Fehr and Gachter, 2000, 2002; Gintis et al., 2003).

Our model suggests that when the propensity to cooperate and the propensity to punish are modelled as independent traits, a negative correlation between altruism and punishment robustly evolves, although the size of the correlation varies with parameters such as group size, duration of the group, and the cost of punishing others.

---

*Corresponding author. Department of Biological Sciences, Binghamton University, Binghamton, NY 13902-6000, USA. Fax: +1 607 777 6521.

E-mail addresses: oeldakar@gmail.com (O.T. Eldakar), dfarrell1@binghamton.edu (D.L. Farrell), dwilson@binghamton.edu (D.S. Wilson).

0022-5193/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jtbi.2007.07.024

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

2. Methods

The program was implemented in Mathematica and is available from the authors upon request. We composed an N-person evolutionary game theory model that emulates one of the standard public goods games in experimental economics (Fehr and Gachter, 2000). The model begins with an infinite population of individuals that vary in their propensity for altruism (A) and punishment (P). These traits are modelled as two variables that initially vary uniformly and independently between 0 and 1 at 0.1 increments. A large number (T) of groups of size N are formed at random. Members of each group play multiple rounds (R) of a two-phase public goods game. During phase 1, each individual is given an endowment (E) and allowed to contribute a proportion to a central fund, which is doubled and distributed equally to all members of the group. The remainder is retained by the individual at its initial value. Each individual's (i) altruism trait determines the proportion of its endowment that it contributes (A_i) and withholds ($1-A_i$). Individual payoffs at the end of phase 1 can be represented by Eq. (1), where the profits earned for a given individual (pay_i) is calculated by adding the individual's share of the public fund $2E(\sum_{j=1}^{N} A_j)/N$ to the portion of the endowment the individual selfishly withheld from group donation $E(1-A_i)$. Individuals maximize their own payoff by withholding all of their endowment ($A=0$), but this strategy minimizes the payoff for the group, resulting in the classical prisoner's dilemma situation:

$$pay_i = E(1-A_i) + \frac{2E\left(\sum\limits_{j=1}^{N} A_j\right)}{N}. \tag{1}$$

During phase 2, individuals are allowed to contribute resources to detect and punish those who were stingy during phase 1 (the cheaters). Each individual is assumed to know the total contribution of other group members but not the contribution of each individual. Investing in punishment results in a probability that the least altruistic member of the group (other than oneself) will be detected and excluded from subsequent rounds of the game, to be replaced by another individual drawn randomly from the same population as the original members. This is biologically reasonable if we assume that not everyone can get into groups and that the remainder forms a waiting list for replacements. The fact that the replacements play fewer rounds than the original members is immaterial because they still contribute to fitness differentials in the total population, based on how they play the game during the remaining rounds.

The amount that an individual invests in punishment is based on three factors, as shown in Eq. (2). The first term (P_i) represents the individual's static punishment trait. The second term $(\sum_{j=1,j\neq i}^{N-1} 1-Aj)/(N-1)$ represents the average amount of cheating that took place among other members of the group. The third term (C) represents the amount required to detect the worst cheater with certainty:

$$punC_i = P_i \frac{\left(\sum\limits_{j=1,j\neq i}^{N-1} 1-Aj\right)}{N-1} C. \tag{2}$$

A maximum of two individuals can be removed during any particular round of the game; the worst cheater, based on the efforts of the other group members, or the next worst cheater, based on the efforts of the worst cheater. The probability that the worst cheater will not be detected by a given member of the group i is

$$esc_i = 1 - P_i \frac{\left(\sum\limits_{j=1,j\neq i}^{N-1} 1-Aj\right)}{N-1}. \tag{3}$$

The probability that the worst cheater will be detected and removed by any member of the group is

$$rem_{all} = \left(1-\prod_{i=1}^{n-1} esc_i\right) D, \tag{4}$$

where the worst cheater is not included in the calculation. The term D gives the probability that the cheater can be removed, once detected. When $D=1$, the probability of detection is equal to the probability of removal. When $D=0$ then removal is impossible, regardless of detection. The probability that the next worst cheater is removed is similar to Eq. (4), with only the worst cheater included in the calculation. The idea of the worst cheater removing the second worst cheater makes sense for two reasons. First, by removing the second worst cheater, the worst cheater reduces the amount of cheating perceived by the group, effectively weakening the strength of punishment (decreasing the middle term in Eq. (2)). Second, despite the punishment efforts of the group there remains uncertainty that the worst cheater will be banished. Therefore, by removing the next worst cheater, the worst cheater increases its likelihood of remaining in later rounds. The other ($N-2$) members are safe during a particular round of the game but can become vulnerable if replacements make them one of the worst two cheaters in subsequent rounds.

After the game is played for a number of rounds (R) within each group and for a large number (T) of randomly formed groups, each individual is assigned a fitness based on its total earnings and a baseline fitness value (B), representing the fact that fitness is not determined entirely by the interactions that take place during the game. Fitness is then summed for each strategy-type accounting for both abundance and fitness, deriving a cumulative fitness value for each combination of altruism and punishment (121 types). These cumulative sums are then normalized to sum 1, representing the frequency of each strategy type following asexual reproduction in direct proportion to fitness, which become the new frequencies of the 121 types in the infinite population for the next round of group

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

Table 1
List of model parameters and default values

Variable	Baseline value	Definition
A	(0–1) at 0.1 increments	Proportion of endowment allocated to the group fund
C	40	Maximum cost of punishment
D	0.5	Efficiency of removal of a cheater upon detection
E	50	Resources allocated to each player at the beginning of each round of play
M	10^{-4}	Mutation rate
N	4	Group size
P	(0–1) at 0.1 increments	Propensity to punish
R	6	Number of rounds played per generation
T	10^4	Number of groups per simulation

formation. It should be noted that asexual reproduction is interpreted loosely in terms of the replicator dynamic of evolutionary game theory, which includes any process that causes the most successful strategies to increase in frequency in the population (Gintis, 2000b).

Mutations in the altruism and punishment traits were assumed to occur with a frequency of (M) and took place during the asexual reproduction stage. In one set of simulations, a type was assumed to mutate into any other type with equal probability, resulting in all types potentially being present in the population at a low frequency. In another set of simulations, mutations were assumed to deviate by a value of ± 0.1, which means that a given type could be completely absent from the population.

Simulations were run with 2 alternate scenarios of initial population frequencies. The first began with all combinations of the altruism (A) and punishment (P) traits in equal proportions. The second began with the population fixed for $A = 0$ and $P = 0$ to see if altruism and punishment could evolve from mutation frequencies.

To summarize (1) Groups are most productive when everyone invests their entire endowment, (2) in the absence of punishment, individuals are most productive within each group when they withhold their own investment; (3) punishment can cause cheating to become disadvantageous; (4) punishment is costly for the punisher; and (5) the altruism and punishment traits are initially uncorrelated. Our prediction is that a negative correlation (cheaters more likely to punish) will develop on the basis of the model dynamics. The parameters and their default values are listed in Table 1.

3. Results

3.1. Initial conditions

A representative simulation run with the initial population consisting of an even distribution of all possible combinations of altruism and punishment is shown in Fig. 1 (see legend for parameter values). The most stingy individuals are quickly removed from the population by punishment (generation 20), followed by the elimination of most punishers due to the cost of punishment (generation 50). Variation in both the altruism and punishment

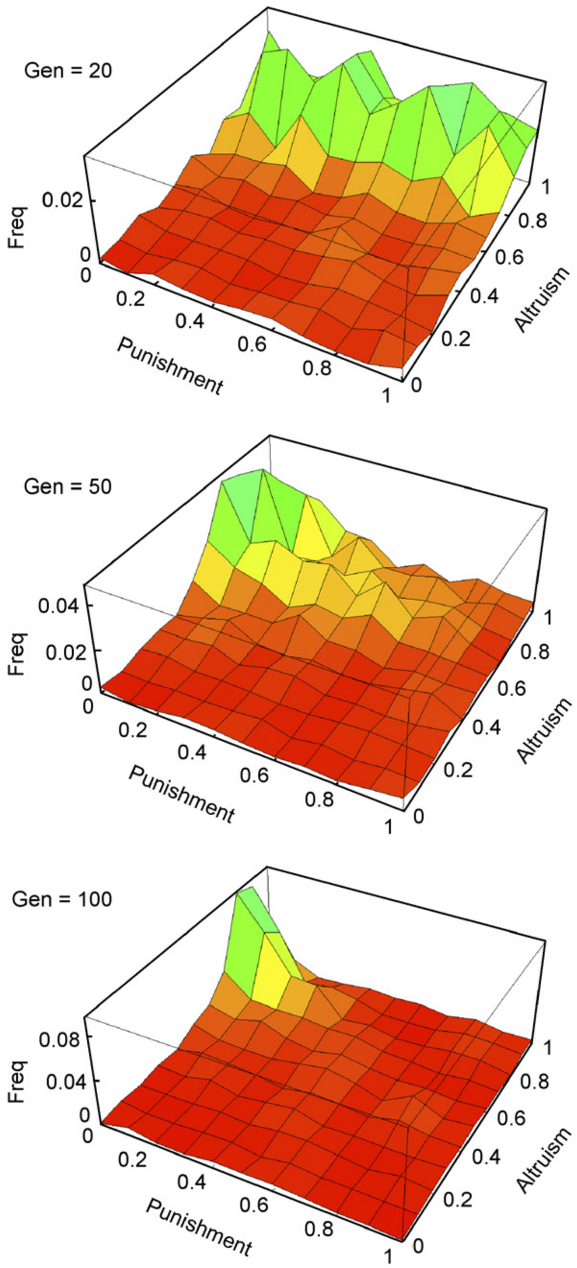

Fig. 1. Selected generations of a simulation run started from an initial population composed of even population distribution of all strategies. Topography figures illustrate the phenotypic distribution of the population, demonstrating the reduction of cheaters (generation 20), then the reduction of altruistic punishers (generation 50), and a negative correlation between altruism and punishment at equilibrium (generation 100). This run consisted of 1000 groups of $N = 4$ individuals created at random from the total population every "generation". Within each group, the game was played for $R = 6$ rounds. The maximum cost of detecting and excluding a cheater was $C = 80\%$ of one's endowment and the efficiency of removal of a cheater upon detection was $D = 0.5$.

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

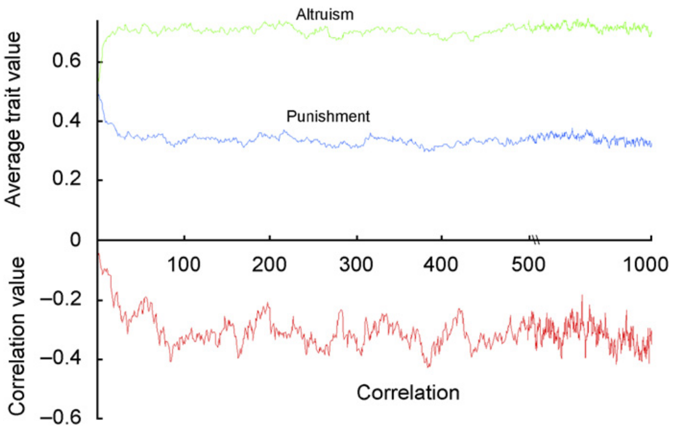

Fig. 2. Time-series graph of a selected simulation showing the average trait values of the population for altruism and punishment over 1000 generations along with the emerging negative correlation between the two traits. Parameter values are the same as Fig. 1.

trait is maintained at equilibrium, with a negative correlation between the two traits representing a stable equilibrium of altruistic non-punishers and selfish punishers, as we predicted (generation 100). Fig. 2 shows that the equilibrium is maintained over the long term, although coupled oscillations between the frequencies of the two traits and their covariance take place over shorter time scales. Unlike altruistic punishers, selfish punishers possess the ability to recoup the cost of punishment through their exploitation of altruists within groups.

Fig. 3 shows a comparable run in which the initial population consists entirely of selfish non-punishers ($A = 0$, $P = 0$). Remarkably, the same equilibrium is established, although a large number of generations is required. If the capacity for punishment is eliminated by setting D to zero, altruism does not evolve from this starting point. To see how selfish punishment promotes the evolution of altruism, consider a single selfish punisher in a given group. By expelling the most selfish individuals, which are replaced by randomly chosen members of the total population, the punisher increases the average degree of altruism within the group. Altruists now benefit from each other and the selfish punisher recovers the cost of punishment by exploiting the altruists during subsequent rounds. Of course, this will only work if there is a sufficient frequency of altruists in the total population. Although the simulation run begins with $A = 0$, a mutation rate of $M = 10^{-4}$ results in a selection–mutation balance of approximately 7% of the population with $A > 0$, which is sufficient for the concentrating effect of punishment to take place. A large number of generations is required for the selection–mutation balance to establish itself, accounting for the time required for altruism to evolve in Fig. 3. An implication is that altruism will not evolve from a starting point of $A = 0$ when the mutation rate (M) is sufficiently low or the cost of punishing (C) is sufficiently high, which we will demonstrate below.

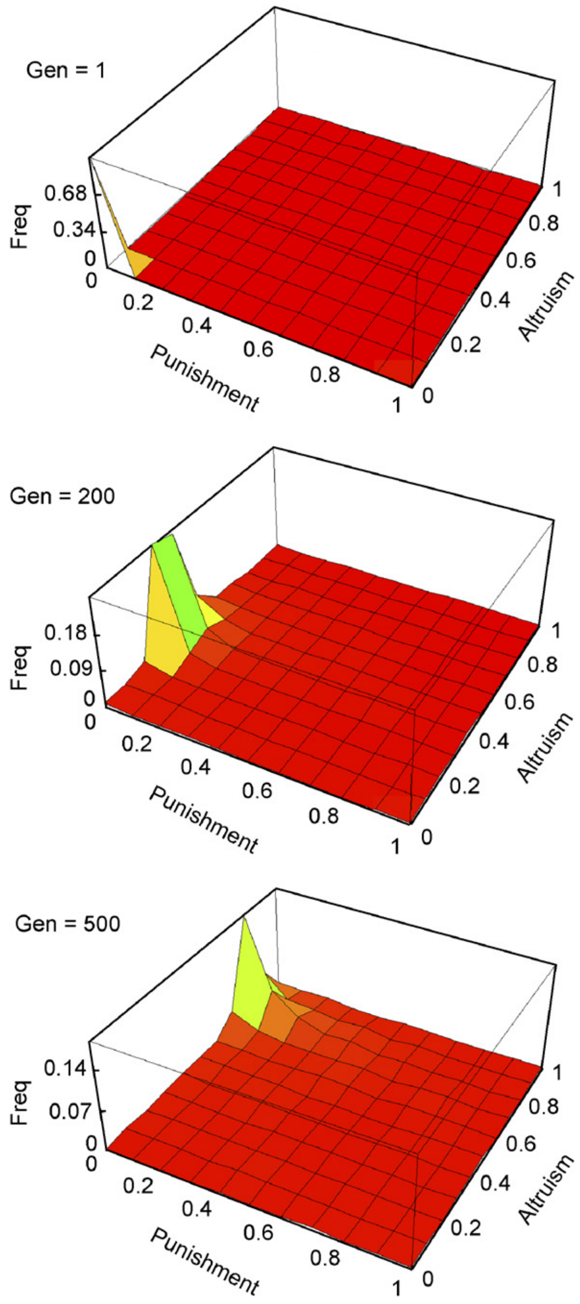

Fig. 3. Selected generations of a simulation run started from an initial population composed of only the selfish non-punishing strategy. Topography figures illustrate the phenotypic distribution of the population, starting from the initial selfish non-punisher distribution (generation 1), then the increase of altruism in the population (generation 200), and the further increase of altruism (generation 300) until reaching a stable equilibrium depicted in Fig. 1. This run consisted of 1000 groups of $N = 4$ individuals created at random from the total population every "generation". Within each group, the game was played for $R = 6$ rounds. The maximum cost of detecting and excluding a cheater was $C = 80\%$ of one's endowment and the efficiency of removal of a cheater upon detection was $D = 0.5$.

Now that we have described the dynamics of the model during a single simulation run, we will vary single parameters of the model while keeping the others at their default values (shown in Table 1 and Fig. 1).

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

3.2. *Cost of punishment*

The cost of punishing others is modelled by the parameter *C*, which represents the proportion of the endowment that is required to detect a cheater with certainty. Fig. 4 shows that the equilibrium levels and altruism and punishment decline as *C* is increased, although they remain at moderate levels even at the highest value of *C*. The lower part of Fig. 4 shows that the correlation between the altruism and punishment is close to zero when punishing others is nearly cost-free, but becomes increasingly negative when punishing others becomes costly. Thus, our model suggests that *the concept of selfish punishment is especially relevant when punishing others is costly*.

3.3. *Group size*

The equilibrium levels of altruism and punishment also decline with group size (Fig. 5), although remaining at moderate levels even at the highest value tested. The negative correlation between altruism and punishment reaches its lowest value at a value of $N = 7$ and then rises slightly. These trends reflect a number of factors. When $N = 2$, the concept of a public good is not applicable because an individual must punish if there is to be any

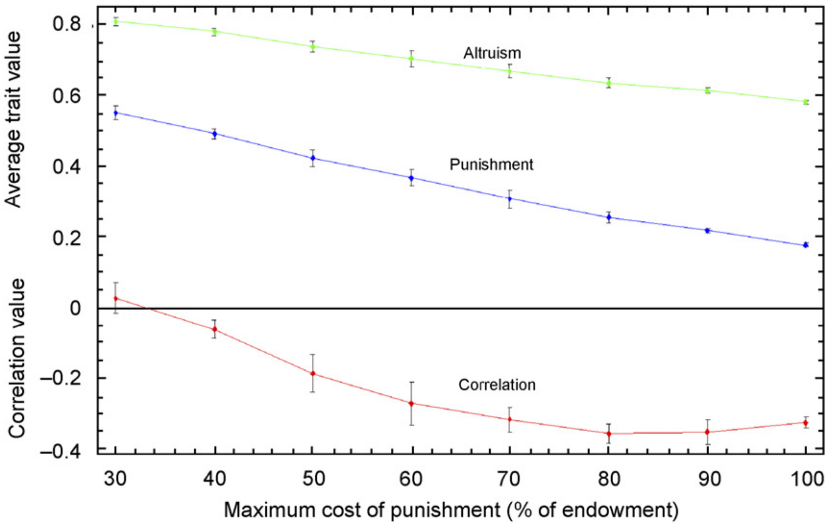

Fig. 4. The average correlation between altruism and punishment as the cost of detecting and excluding cheaters (*C*) increases. Error bars indicate the range of five replicate simulations for each set of parameter values. A low cost indicates that the maximum probability of detecting and removing a cheater can be achieved with a low proportion of one's endowment. Within this range, the amount that an individual invests is based on the value of its punishment trait and the amount of cheating that took place during phase 1 of the game. Other parameter values are the same as Figs. 1 and 2.

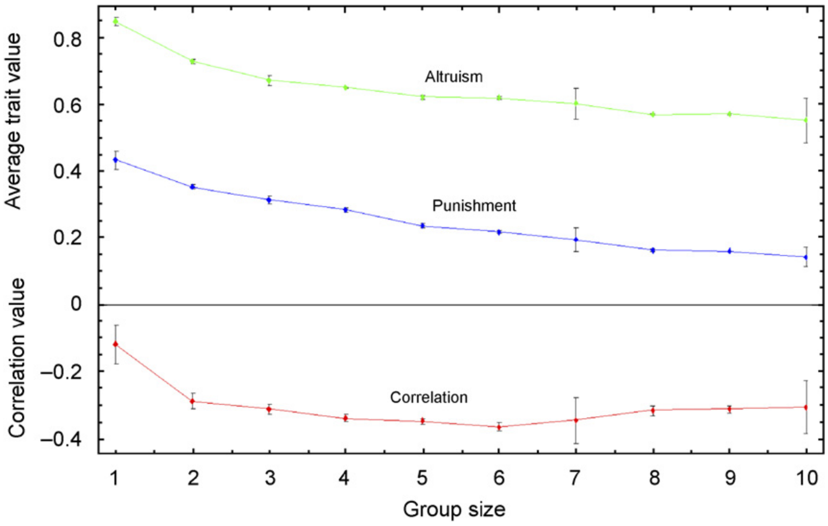

Fig. 5. Levels of altruism and punishment decline however remain at moderate levels when increasing group size (*N*). The correlation between altruism and punishment declines with group size until its maximum value at group size $n = 7$ then begins to raises slightly. Error bars indicate the range of five replicate simulations for each set of parameter (*N*) values. Other parameters values are the same as Figs. 1 and 2.

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

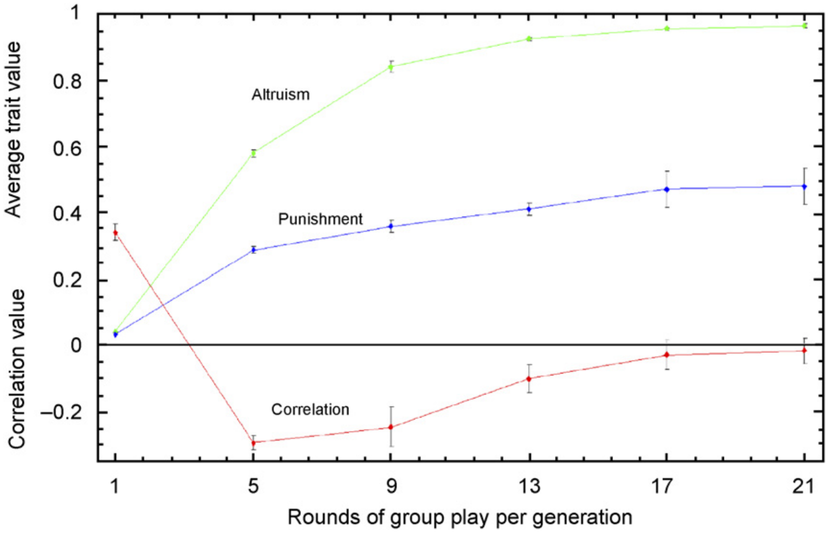

Fig. 6. Levels of Altruism and punishment increase with increasing the number of rounds (R) of group play per generation. Altruism and punishment levels at parameter value $R = 1$ indicate that both cooperation and punishment are unsuccessful in one-shot games. These values increase in simulations of successively longer iterated games. The correlation between altruism and punishment approaches zero with increasing rounds of group play indicating less overall punishment (see text for explanation indirect effect of round length on punishment costs). Error bars indicate the range of five replicate simulations for each set of parameter (R) values. Other parameters values are the same as Figs. 1 and 2.

punishment of its single partner. At high values of N, the capacity to punish is limited by round length. Only a maximum of two cheaters can be removed during each round, which means that when group size exceeds round length, intermediate cheaters are immune to punishment.

3.4. Round length

When groups exist for only a single round of play, punishment cannot maintain altruism and both traits go to zero, except for mutations, as shown in Fig. 6. The positive correlation between the two traits when $R = 1$ is based on the fact that most individuals have values of zero for both traits while a few (the mutants) have positive values of both traits. Punishment becomes increasingly effective at maintaining cooperation as round length increases. At the highest value of R tested, an average value of approximately 0.4 for the punishment trait maintains altruism at close to its maximum value. Round length has an indirect effect on the cost of being punished and the benefits of excluding cheaters from one's group. When cheaters are excluded from their group, they keep their earnings but "sit out" the remaining rounds of the game. This cost becomes increasingly severe as R increases, especially relative to those who remain in the game. Thus, increasing R is similar to decreasing the cost of punishing cheaters, causing the correlation between altruism and punishment to decrease, as in Fig. 4.

3.5. Mutations

For simulations that began with an equal distribution of all types, model results were not affected by the two assumptions about mutations within the range of $M = 10^{-2}$–10^{-7}. For models that began with $A = 0$, $P = 0$, a sufficiently small mutation rate (given either assumption) can prevent the frequency of altruists at mutation–selection balance from achieving the threshold required for the concentrating effect of punishment to take place. Similarly, when the cost of punishment is increased (C), a corresponding increase in mutation rate is required for altruism to evolve from a mutation frequency.

4. Discussion

Punishing others requires time, energy, and risk, just like any other trait. In dyadic interactions, the costs and benefits of punishing a cheater can be calculated in a straightforward fashion because the punisher is the sole beneficiary of the punishment (Clutton-Brock and Parker, 1995). In larger groups, the benefits of punishment are shared by other members of the group who do not share the costs, creating a public goods problem that increases with the cost of punishing others.

The proverbs "It takes a thief to catch a thief" and "there is no honor amongst thieves" imply that no one is better at finding a cheater than another cheater and that cheaters themselves interact competitively. Cheaters might have a number of special advantages for detecting and excluding other cheaters, such as familiarity with cheating strategies or experience at fighting. These special advantages are not included in our model. Instead, we made the conservative assumption that altruism and punishment are separate and (initially) uncorrelated traits. An extreme altruist is just as capable of detecting and excluding cheaters as an extreme cheater. Nevertheless, a negative correlation between the altruism and punishment traits robustly develops based on the model dynamics.

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

One way to interpret selfish punishment is as an entirely selfish strategy whereby cheaters maintain and protect "flocks" of cooperators for their own advantage, similar to the way that the mafia offers protection for a price. Alternatively, selfish punishment can be regarded as a division of altruistic labor, with some individuals providing the first-order public good of cooperation and others providing the second-order public good of punishment, similar to the way that human communities support a police force. Division of labor evolves because altruistic punishers suffer a double cost whereas selfish punishers in the same group are "compensated" for the cost of punishment by being "exempted" from cooperation during the first round. Regardless of how it is interpreted, selfish punishment can cause altruism to evolve and be maintained at a high frequency without the problems usually associated with altruistic punishment. There is a threshold frequency of altruism that must be crossed before altruism can be positively selected, but it is sufficiently low that mutation–selection balance is sufficient, at least for certain combinations of M and C. See Wilson and Dugatkin (1997) for a discussion of "the problem of origination" for the evolution of altruism in models that assume quantitative variation vs. discrete traits.

The correlation between selfishness and punishment becomes increasingly strong as the cost of punishing others (C) increases. Examples of low-cost punishment in human social interactions include gossip and collective decisions that cannot be opposed because the group is so much stronger than any particular individual (see Sober and Wilson, 1998, Chapter 5 for ethnographic examples from a random sample of cultures). In these cases, punishment and altruism should remain uncorrelated. The effects of group size and round length can also be interpreted in terms of public benefits and private costs. Increasing group size makes punishment ineffective because at most only one cheater can be removed during each round of the game. Increasing round length enables more cheaters to be excluded and increases the differential between those who are excluded and those who remain.

Nakamaru and Iwasa (2006) model of selfish punishment considers four discrete strategies; altruistic punisher (AP), altruistic non-punisher (AN), selfish punisher (SP) and selfish non-punisher (SN). The interactions are dyadic and punishment causes one's selfish partner to pay a "fine" at an expense to the punisher. Individuals exist on a two-dimensional lattice and interact either with their four nearest neighbors (lattice model) or with four individuals chosen at random from the total population (completely mixing population). In both cases, individuals compete with their four nearest neighbors based on their payoffs, either in terms of survival (score-dependent viability model) or reproduction (score-dependent fertility model). The SP strategy can invade and persist in some of these conditions but not others. As in our model, it can facilitate the evolution of altruistic strategies by virtue of its negative effect on other punishers.

Our model allows gradations of altruism and punishment, assumes interactions in randomly formed groups of size N rather than dyadic interactions on a lattice, and is intended to emulate the public goods games that experimental economists use to study altruism and punishment in human social interactions. Given these assumptions, we observe a robust negative correlation between altruism and punishment, although the magnitude of the correlation varies with the parameter values. Obviously, these two models only begin to explore the different kinds of social settings and population structures in which selfish punishment might exist as a successful behavioral strategy (Price et al., 2002; Shinada et al., 2004; Sigmund et al., 2001).

Considerable evidence for altruism maintained by competition among selfish individuals exists for non-human species, from insects to vertebrates. Wenseleers et al. describe a "corrupt policing" strategy in tree wasps *Dolichovespula sylvestris*, where workers that police other workers lay their own eggs (Wenseleers et al., 2005). Scrub jays that tend to steal caches from other scrub jays are also more defensive of their own caches (Emery and Clayton, 2001). In addition to our empirical study on humans that inspired our simulation model (Eldakar et al., 2006), the history of medieval knights provides a potential historical example of selfish punishment. Much as the knights of old are revered in mythology and popular culture, the first Castellans are better described as selfish thugs who fought among themselves to exploit the defenseless, and therefore altruistic, peasants (Bisson, 1994). As Pope Gregory VII put it during the 11th century (quoted in Bisson, 1994, p. 42), "Who does not know that kings and princes derive their origin from men ignorant of God who raised themselves above their fellows by pride, plunder, treachery, and murder?"

In this human example and many non-human examples, the dynamics of altruism and punishment are complicated by power asymmetries such as social dominance (e.g. Clutton-Brock and Parker, 1995; de Waal, 1986; Giraldeau and Caraco, 2000; Kim, 2006; Monnin and Ratnieks, 2001; Ruttan and Borgerhoff Mulder, 1999). Our model shows that selfishness and punishment can become correlated even in the absence of power asymmetries and other factors that give cheaters an intrinsic advantage in punishing other cheaters.

Acknowledgments

We thank A.B. Clark, C. Laramee, D. Smail, T. Wenseleers, N. Prince and members of EvoS, Binghamton University's Evolutionay Studies Program, for helpful discussion.

References

Bisson, T.N., 1994. The feudal revolution. Past Present 142, 6–42.
Bowles, S., Gintis, H., 2002. Homo reciprocans. Nature 415, 125–128.

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024

Bowles, S., Gintis, H., 2004. The evolution of strong reciprocity: cooperation in heterogeneous populations. Theor. Popul. Biol. 65, 17–28.

Boyd, R., Gintis, H., Bowles, S., Richerson, J., 2003. The evolution of altruistic punishment. Proc. Natl. Acad. Sci. 100, 3531–3535.

Clutton-Brock, T.H., Parker, G.A., 1995. Punishment in animal societies. Nature 373, 209–216.

de Waal, F.B.M., 1986. Integration of dominace and social bonding in primates. Q. Rev. Biol. 61, 479–549.

Eldakar, O.T., Wilson, D.S., O'Gorman, R., 2006. Emotions and actions associated with altruistic helping and punishment. Evol. Psychol., in press.

Emery, N.J., Clayton, N.S., 2001. Effects of experience and social context on prospective caching strategies by scrub jays. Nature 414, 443–446.

Fehr, E., 2004. Don't lose your reputation. Nature 432, 449–450.

Fehr, E., Fischbacher, U., 2003. The nature of human altruism. Nature 425, 785–791.

Fehr, E., Gachter, S., 2000. Cooperation and punishment in public goods experiments. Am. Econ. Rev. 90, 980–994.

Fehr, E., Gachter, S., 2002. Altruistic punishment in humans. Nature 415, 137–140.

Fehr, E., Fischbacher, U., Gachter, S., 2002. Strong reciprocity, human cooperation, and the enforcement of social norms. Hum. Nat.—Int. Bios. 13, 1–25.

Gintis, H., 2000a. Strong reciprocity and human sociality. J. Theor. Biol. 206, 169–179.

Gintis, H., 2000b. Game Theory Evolving. Princeton University Press, Princeton, NJ.

Gintis, H., Bowles, S., Boyd, R., Fehr, E., 2003. Explaining altruistic behavior in humans. Evol. Hum. Behav. 24, 153–172.

Giraldeau, L.A., Caraco, T., 2000. Social Foraging Theory. Princeton University Press, Princeton, NJ.

Kim, S., 2006. For whom do we reciprocate? The effects of dominance relationships on the use of incentives in collective action. Unpublished manuscript.

Monnin, T., Ratnieks, F.L.W., 2001. Policing in queenless ponerine ants. Behav. Ecol. 50, 97–108.

Nakamaru, M., Iwasa, Y., 2006. The coevolution of altruism and punishment: role of the selfish punisher. J. Theor. Biol. 240, 475–488.

Price, M.E., Cosmides, L., Tooby, J., 2002. Punitive sentiment as an anti-free rider psychological device. Evol. Hum. Behav. 23, 203–231.

Ruttan, L.M., Borgerhoff Mulder, M., 1999. Are east African pastoralists truly conservationsits? Curr. Anthropol. 40, 621–652.

Shinada, M., Yamagishi, T., Ohmura, Y., 2004. False friends are worse than bitter enemies: "altruistic" punishment of in-group members. Evol. Hum. Behav. 25, 379–393.

Sigmund, K., Hauert, C., Nowak, M., 2001. Reward and punishment. Proc. Natl. Acad. Sci. 98, 10757–10762.

Sober, E., Wilson, D.S., 1998. Unto Others, The Evolution and Psychology of Unselfish Behavior. Harvard University Press, Cambridge, MA.

Wenseleers, T., Tofilski, A., Ratnieks, F.L.W., 2005. Queen and worker policing in the tree wasp *Dolichovespula sylvestri*. Behav. Ecol. Sociobiol. 58, 80–86.

Wilson, D.S., Dugatkin, L.A., 1997. Group selection and assortative interactions. Am. Nat. 149, 336–351.

Please cite this article as: Eldakar, O.T., et al., Selfish punishment: Altruism can be maintained by competition among cheaters. J. Theor. Biol. (2007), doi:10.1016/j.jtbi.2007.07.024