Published in final edited form as:

Nature. 2009 January 1; 457(7225): 79–82. doi:10.1038/nature07601.

Indirect reciprocity provides a narrow margin of efficiency for costly punishment

Hisashi Ohtsuki^{1,2}, Yoh Iwasa³, and Martin A. Nowak⁴

¹Department of Value and Decision Science, Tokyo Institute of Technology, Tokyo 152-8552, Japan

²PRESTO, Japan Science and Technology Agency, 4-1-8 Honcho Kawaguchi, Saitama 332-0012, Japan

³Department of Biology, Faculty of Sciences, Kyushu University, Fukuoka 812-8581, Japan

⁴Program for Evolutionary Dynamics, Department of Organismic and Evolutionary Biology, Department of Mathematics, Harvard University, Cambridge MA 02138, USA

Abstract

Indirect reciprocity¹⁻⁵ is a key mechanism for the evolution of human cooperation. Our behavior toward other people depends not only on what they have done to us, but also on what they have done to others. Indirect reciprocity works via reputation⁵⁻¹⁷. The standard model of indirect reciprocity offers a binary choice: people can either cooperate or defect. Cooperation implies a cost for the donor and a benefit for the recipient. Defection has no cost and yields no benefit. Currently there is considerable interest in studying the effect of costly (or altruistic) punishment on human behavior¹⁸⁻²⁵. Punishment implies a cost for the punished person. Costly punishment means that the punisher also pays a cost. It has been suggested that costly punishment between individuals can promote cooperation. Here we study the role of costly punishment in an explicit model of indirect reciprocity. We analyze all social norms, which depend on the action of the donor and the reputation of the recipient. We allow errors in assigning reputation and study gossip as a mechanism for establishing coherence. We characterize all strategies that allow the evolutionary stability of cooperation. Some of those strategies use costly punishment, while others do not. We find that punishment strategies typically reduce the average payoff of the population. Consequently, there is only a small parameter region where costly punishment leads to an efficient equilibrium. In most cases, the population does better by not using costly punishment. The efficient strategy for indirect reciprocity is to withhold help for defectors rather than punish them.

Human societies are organized around cooperative interactions. But why would natural selection equip selfish individuals with altruistic tendencies? This question has fascinated evolutionary biologists for decades. One answer is given in terms of direct reciprocity²⁶⁻²⁹.

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

Correspondence to: Hisashi Ohtsuki.

Author Information Reprints and permissions information is available at npg.nature.com/reprintsandpermissions. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to H.O. (ohtsuki.h.aa@m.titech.ac.jp)..

There are repeated encounters between the same two individuals: I help you, and you help me. More recently, indirect reciprocity has emerged as a more general model: I help you, and somebody helps me. Indirect reciprocity is based on reputation⁵. People monitor the social interactions within their group. Helping others establishes the reputation of being a helpful individual. Natural selection can favor strategies that help those who have helped others⁵⁻¹⁷. The consequences for widespread cooperation are enormous. Direct reciprocity is like an economy based on the exchange of goods, whereas indirect reciprocity resembles the invention of money. The money that feeds the engines of indirect reciprocity is reputation. For direct reciprocity, my strategy depends on what you have done to me; for indirect reciprocity, my strategy also depends on what you have done to others. Direct and indirect reciprocity are mechanisms for the evolution of cooperation³⁰.

Punishment refers to an action that implies a cost for the punished person. Costly punishment means that the punisher also pays a cost for exercising punishment. In certain experimental situations costly punishment has been called ‘altruistic punishment’, because the punishers cannot expect any material gain from their action^{20,21}. In reality, however, most punishment actions among humans are associated with the expectation of a delayed material gain, and therefore they are not altruistic.

The suggested idea for evolution of cooperation is that people might be more willing to cooperate under the threat of punishment. We note, however, that costly punishment is not a separate mechanism for the evolution of cooperation, but a form of direct or indirect reciprocity. If I punish you, because you have defected with me, then I use direct reciprocity. If I punish you, because you have defected with others, then indirect reciprocity is at work. In the setting of direct reciprocity, punishment is a form of retaliation²⁵. For indirect reciprocity, punishment works via reputation and also includes third party actions, which means that observers of an interaction are willing to punish defectors at a cost to themselves²¹. Therefore, any discussion of the evolution of costly punishment brings us immediately into the framework of direct or indirect reciprocity.

In general, the reputation score could be a continuous variable⁵, but here we consider a simple model with binary reputation. People have either a good reputation, G , or a bad reputation, B . At times, two random players are chosen from the population, one in the role of donor, the other in the role of recipient. The donor can either cooperate, C , defect, D or punish, P . Cooperation means the donor pays a cost, c , and the recipient gets a benefit, b . Punishment implies that the donor pays a cost, α , and the recipient incurs a cost, β . For defection there is no cost and no benefit.

The interaction between the donor and the recipient is observed by the other members of the population (Fig.1). The reputation of the donor is updated according to a social norm. First order assessment depends only on the action of the donor: for example, cooperation leads to a good reputation, while defection leads to a bad reputation. Second order assessment^{12,13} depends both on the action of the donor and the reputation of the recipient: for example, it could be deemed ‘good’ to cooperate with a good recipient, but ‘bad’ to cooperate with a bad recipient. In this paper, we study social norms that use second order assessment. The donor has three possible moves, C , D or P , and the recipient has one of two reputations, G or

B. Hence, there are 6 combinations and $2^6=64$ social norms with second order assessment. All detailed calculations are shown in the supplementary information (SI) online.

Any interaction either leads to a good or bad reputation of the donor. We assume that this process of reputation updating is subject to errors. There may be wrong observations or the spread of false rumors. With probability μ an incorrect reputation is assigned and adopted by all. In the simplest model, everyone has the same opinion of everyone else. There are no private lists of reputation. Triggering a wrong reputation affects everyone equally. The parameter $q=1-2\mu$ quantifies the ability of the population to distinguish between good and bad. We call q the ‘social resolution’. If $\mu=1/2$ then reputation is assigned at random, and there is no ability to distinguish between good and bad, $q=0$.

Games of indirect reciprocity contain social norms and action rules. The action rule specifies for the donor whether to cooperate, defect or punish a recipient who is either good or bad. For example, the action rule *CD* prescribes cooperation with a good recipient and defection with a bad recipient; this rule does not use costly punishment. In contrast, the action rule *CP* prescribes cooperation with good recipients and punishment of bad recipients. The action rules *CC*, *DD* and *PP* encode, respectively, unconditional cooperation, defection and punishment. In total, there are 9 possible action rules.

For each of the 64 social norms we study the competition of all 9 action rules. We assume that everyone in the population has the same social norm, and we evaluate if this norm allows the evolutionary stability of action rules that specify to cooperate with good recipients. There are only two candidates for such action rules, *CD* and *CP*, because *CC* is not stable against invasion by defectors, *DD*. Figure 2 shows all social norms that allow the evolution of cooperation. The action rule, *DD*, is evolutionarily stable for any social norm.

Social norms that stabilize the *CD* action rule have the following properties: (i) cooperation with a good recipient leads to a good reputation; (ii) defection against a good recipient leads to a bad reputation; and (iii) defection against a bad recipient leads to a good reputation. The three remaining positions in the norm can be either *G* or *B*. If the cost of cooperation is greater than the cost of punishment, $c>\alpha$, then punishing a good recipient must lead to a bad reputation; otherwise a donor can keep a good reputation by using the cheaper punishment option instead of the more expensive cooperation move.

Social norms that stabilize the *CP* action rule have the following properties: (i) cooperation with a good recipient leads to a good reputation; (ii) defection always leads to a bad reputation; and (iii) punishing a bad recipient leads to a good reputation.

CD action rules are evolutionarily stable, if the social resolution exceeds the cost-to-benefit ratio, $q>c/b$. In contrast, *CP* action rules are evolutionarily stable if $q>\max\{c,\alpha\}/(b+\beta)$. Note that costly punishment can stabilize cooperation even if $q<c/b$. Thus, costly punishment can in principle extend the stability range of cooperation. *DD* action rules are always evolutionarily stable.

We have performed computer simulations in heterogeneous populations of finite size in order to test the validity of our analytical calculations. We find that the *CD* and *CP* action

rules are stable against invasion by other action rules under the appropriate social norms and given the right parameter region. In the simplest simulations, everyone has the same information about the reputation of others. In the extended simulations, we drop this assumption. Now there are individual errors in assigning reputation. Consequently everyone has a private list of the reputation of others. These errors can destroy indirect reciprocity unless there is a mechanism for re-establishing coherence. Gossip is such a mechanism. We assume that individuals talk to each other and sample each other's opinions (as in a 'voter model'). If there are enough communication events, then we observe the evolutionary stability of our strategies as predicted. We have also studied errors in executing the wrong action ('trembling hand') or recalling an incorrect reputation ('fuzzy mind'). Our results are robust as long as these errors are not too frequent. All simulations are described in the SI.

For some parameter regions multiple action rules are evolutionarily stable. Therefore, we ask the following question: for all possible parameter regions which of the three action rules, *CD*, *CP* and *DD* are stable and which one is most efficient in the sense of leading to the highest average payoff at equilibrium? We obtain the following answer:

- i. If $q > c/b$, then *CD* is most efficient.
- ii. If $c/b > q > c/(b+\beta)$ then *CP* is stable and more efficient than *DD*, if the following two conditions hold:

$$q > \frac{\alpha}{b+\beta} \quad \text{and} \quad q > \frac{\alpha+\beta-b+c}{\alpha+\beta+b-c}. \quad (1)$$

Otherwise *DD* is more efficient than *CP*. If $b < c$ then *DD* is always more efficient than *CP*.

- iii. If $c/(b+\beta) > q$, then only *DD* is evolutionarily stable.

Thus, if the accuracy of assigning the correct reputation, q , is too low then only *DD* is efficient. If q is sufficiently large, then *CD* is efficient. For intermediate values of q there can be a small window where *CP* is efficient. But the existence of this parameter region depends on whether the key parameters, b , c , α , β , fulfill the constraints given by eq.(1). Let us consider a numerical example. If $b=2$, $c=1$, $\alpha=1/2$ and $\beta=2$, then *CD* is efficient for $q > 1/2$, while *CP* is efficient for $1/2 > q > 3/7$ and *DD* is efficient for $3/7 > q$. If we increase the effect of punishment to $\beta=5/2$ (or larger), then there is no region left where *CP* is efficient. Intuitively speaking, if *CD* is evolutionarily stable then it is always the most efficient equilibrium. If it is not stable, then the remaining parameter region where *CP* is stable and more efficient than *DD* is very small or non-existent. Figure 3 illustrates the narrow margin of efficiency of costly punishment.

These considerations of efficiency do not imply that all populations will evolve toward punishment free action rules. It is possible that a population is stuck at an inefficient equilibrium for a long time. A model with contingent movement allows us to study the competition of different social norms. We examine a simple scenario, where two groups have two different social norms. One norm stabilizes *CD*, while the other norm stabilizes *CP*. People only interact within their own group, but sometimes they compare their payoff

with individuals from the other group. If the other individual has a higher payoff, then they might move to the other group and adopt its social norm. We observe rapid selection of the efficient equilibrium (see SI).

In an experimental study, the observers of a Prisoner's Dilemma game between two other people sometimes punish defectors at a cost to themselves²¹. This behavior is a form of indirect reciprocity. In another experiment²³, a public goods game is followed by one round of punishment and then by one round of cooperation or defection. This setup is not directly comparable with our model, but the observation is that adding the third round reduces the amount of punishment that is being used in the second round. This particular finding is in agreement with our result: other possibilities of indirect reciprocity reduce the amount of costly punishment. In the context of our theory it would be important to extend both experiments to allow for reputation building over multiple rounds of interaction and a choice between cooperation, defection and costly punishment in every round. We predict that such an experiment will show that costly punishment is an inefficient behavior for most parameter regions.

We have studied the effect of costly punishment in an explicit model of indirect reciprocity. We have analyzed all social norms that use binary reputation and second order assessment. We find that both *CD* and *CP* action rules can stabilize cooperation. These rules reward good recipients with cooperation and ‘punish’ bad ones with either defection (*CD*) or costly punishment (*CP*). If both *CD* and *CP* action rules are evolutionarily stable, then the use of costly punishment leads to a lower equilibrium payoff and is therefore inefficient. It is even possible that costly punishment yields a lower payoff than all-out defection, *DD*. Costly punishment maximizes the group average payoff only for a very limited parameter region. This narrow margin of efficiency requires a fine-tuning of the key parameters. If the social resolution exceeds the cost to benefit ratio, $q > c/b$, then *CD* rules are always more efficient than *CP* rules. Therefore, the evolution of improved mechanisms of indirect reciprocity leads to societies where costly punishment between individuals is not an efficient behavior for promoting cooperation.

Methods Summary

An action rule, s , is formulated as a mapping from $\{G, B\}$ (the recipient's reputation) to $\{C, D, P\}$ (the prescribed action). A social norm, n , is a mapping from the product of $\{C, D, P\}$ (the donor's action) and $\{G, B\}$ (the recipient's reputation) to $\{G, B\}$ (the donor's new reputation). We search for the combination of an action rule, s , and a social norm, n , that satisfies the following two properties: (i) the monomorphic population where all players adopt s and n achieves full cooperation in the absence of errors, and (ii) the action rule, s , is evolutionarily stable under the social norm, n . We check these two criteria for each of all $9 \times 64 = 576$ possible combinations of action rule and social norm, (s, n) . From the first criterion, action rule s must use cooperation (C). Because of the symmetry in the binary labels, G and B , we can assume without loss of generality that the action rule prescribes cooperation to good recipients; i.e. $s(G) = C$. To study the evolutionary stability of the action rule, s , we take advantage of the method of dynamic optimization. We assume that social norm is n and that all players except the focal player adopt action rule s . Under this assumption we calculate

the *best-response* action rule, s^* , of the focal player. If s^* uniquely exists and matches s , then s is evolutionarily stable under n . Coexistence of action rules¹⁰ is not the scope of our analysis. See the supplementary information for further details.

Acknowledgements

Support from the John Templeton Foundation, the Japan Society for the Promotion of Science, Japan Science and Technology Agency (PRESTO), the NSF/NIH joint program in mathematical biology (NIH grant R01GM078986) and J. Epstein is gratefully acknowledged.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

1. Sugden, R. The Economics of Rights, Cooperation and Welfare. Blackwell; Oxford: 1986.
2. Alexander, RD. The Biology of Moral Systems. Aldine de Gruyter; New York: 1987.
3. Kandori M. Social norms and community enforcement. *Rev. Econ. Stud.* 1992; 59:63–80.
4. Okuno-Fujiwara M, Postlewaite A. Social norms and random matching games. *Games Econ. Behav.* 1995; 9:79–109.
5. Nowak MA, Sigmund K. Evolution of indirect reciprocity by image scoring. *Nature*. 1998; 393:573–577. [PubMed: 9634232]
6. Wedekind C, Milinski M. Cooperation through image scoring in humans. *Science*. 2000; 288:850–852. [PubMed: 10797005]
7. Dufwenberg M, Gneezy U, Güth W, van Damme E. Direct vs indirect reciprocity: an experiment. *Homo Oecon.* 2001; 18:19–30.
8. Fishman MA. Indirect reciprocity among imperfect individuals. *J. Theor. Biol.* 2003; 225:285–292. [PubMed: 14604582]
9. Ohtsuki H, Iwasa Y. How should we define goodness? - reputation dynamics in indirect reciprocity. *J. Theor. Biol.* 2004; 231:107–120. [PubMed: 15363933]
10. Brandt H, Sigmund K. The logic of reprobation: assessment and action rules for indirect reciprocation. *J. Theor. Biol.* 2004; 213:475–486. [PubMed: 15488525]
11. Bolton GE, Katok E, Ockenfels A. Cooperation among strangers with limited information about reputation. *J. Pub. Econ.* 2005; 89:1457–1468.
12. Brandt H, Sigmund K. Indirect reciprocity, image-scoring, and moral hazard. *Proc. Natl. Acad. Sci. USA*. 2005; 102:2666–2670. [PubMed: 15695589]
13. Nowak MA, Sigmund K. Evolution of indirect reciprocity. *Nature*. 2005; 437:1291–1298. [PubMed: 16251955]
14. Suzuki S, Akiyama E. Reputation and the evolution of cooperation in sizable groups. *Proc. R. Soc. B*. 2005; 272:1373–1377.
15. Chalub FACC, Santos FC, Pacheco JM. The evolution of norms. *J. Theor. Biol.* 2006; 241:233–240. [PubMed: 16388824]
16. Takahashi N, Mashima R. The importance of subjectivity in perceptual errors on the emergence of indirect reciprocity. *J. Theor. Biol.* 2006; 243:418–436. [PubMed: 16904697]
17. Pacheco JM, Santos FC, Chalub FACC. Stern-judging: a simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS Comp. Biol.* 2006; 2:1634–1638.
18. Yamagishi T. Seriousness of social dilemmas and the provision of a sanctioning system. *Social. Psychol. Q.* 1988; 51:32–42.
19. Clutton-Brock TH, Parker GA. Punishment in animal societies. *Nature*. 1995; 373:209–216. [PubMed: 7816134]
20. Fehr E, Gächter S. Altruistic punishment in humans. *Nature*. 2002; 415:137–140. [PubMed: 11805825]

21. Fehr E, Fischbacher U. Third-party punishment and social norms. *Evol. Hum. Behav.* 2004; 25:63–87.
22. Fowler JH. Altruistic punishment and the origin of cooperation. *Proc. Natl. Acad. Sci. USA.* 2005; 102:7047–7049. [PubMed: 15857950]
23. Rockenbach B, Milinski M. The efficient interaction of indirect reciprocity and costly punishment. *Nature.* 2006; 444:718–723. [PubMed: 17151660]
24. Sigmund K. Punish or perish? Retaliation and collaboration among humans. *Trends Ecol. Evol.* 2007; 22:593–600. [PubMed: 17963994]
25. Dreber A, Rand DG, Fudenberg D, Nowak MA. Winners don't punish. *Nature.* 2008; 452:348–351. [PubMed: 18354481]
26. Trivers RL. The evolution of reciprocal altruism. *Q. Rev. Biol.* 1971; 46:35–57.
27. Axelrod R, Hamilton WD. The evolution of cooperation. *Science.* 1981; 211:1390–1396. [PubMed: 7466396]
28. Colman, AM. *Game Theory and Its Applications in the Social and Biological Sciences.* Routledge; New York: 1995.
29. Rutte C, Taborsky M. The influence of social experience on cooperative behaviour of rats (*Rattus norvegicus*): direct vs generalised reciprocity. *Behav. Ecol. Sociobiol.* 2008; 62:499–505.
30. Nowak MA. Five rules for the evolution of cooperation. *Science.* 2006; 314:1560–1563. [PubMed: 17158317]

Action: cooperate, defect or punish

Figure 1. Indirect reciprocity with costly punishment

The donor chooses one of three actions, cooperate, defect or punish. The recipient has a binary reputation, which is either ‘good’ or ‘bad’. The donor's choice depends on the recipient's reputation. The donor's action is observed by other members of the population, who update the donor's reputation according to a social norm, which is shared by all. The donor is assigned an incorrect reputation with probability μ . The ‘social resolution’, $q=1-2\mu$, is a key parameter of indirect reciprocity: it defines the ability to distinguish between good and bad.

Cost of cooperation > cost of punishment:

	G	B			G	B	
C	G	*	4 social norms	C	G	*	2 social norms
D	B	G		D	B	B	
P	B	*		P	B	G	
	C D		action rule		C P		action rule

Cost of cooperation < cost of punishment:

	G	B			G	B	
C	G	*	8 social norms	C	G	B	2 social norms
D	B	G		D	B	B	
P	*	*		P	*	G	
	C D		action rule		C P		action rule

Figure 2. Social norms of cooperation

We have determined all social norms which allow the evolutionary stability of action rules that prescribe cooperation with good recipients. There are two such action rules, *CD* and *CP*. The former ‘punishes’ bad recipients by defection, *D*, the latter uses costly punishment, *P*. There is an intuitive summary of all successful social norms: following the action rule maintains a good reputation; deviating from it can lead to good or bad reputation; a deviation that is less costly than the prescribed action must lead to a bad reputation. The * symbol denotes *G* or *B*. The detailed parameter requirements for evolutionary stability are given in the main text and in the supplementary information.

Figure 3. The marginal efficiency of costly punishment

Projections of the five-dimensional (b, c, α, β and q) parameter space onto various two-dimensional planes. The parameters b and c denote the benefit and cost of cooperation. The parameters α and β denote the cost and the effect of punishment. The social resolution of the system is given by q : it is the probability to distinguish between good and bad in a world where errors in assignment of reputation are possible. The symbols CD , CP , and DD represent the region where the corresponding action rule is the most efficient equilibrium, which means the evolutionarily stable strategy with the highest average payoff. CD means

cooperate with good recipients and defect with bad ones. *CP* means cooperate with good recipients and punish bad ones. *DD* is unconditional defection. Costly punishment is an efficient equilibrium only for a very constrained parameter region. **(a)** Projections on the β - q plane for $b=3$ and $\alpha=0.8$ or $\alpha=1.2$. **(b)** Projections on the b - q plane for $\beta=3$ and $\alpha=0.8$ or $\alpha=1.2$. **(c)** Projections on the α - β plane for $b=1.5$, and $q=0.3$, $q=0.6$ or $q=0.9$. We always use $c=1$.