

Complex Adaptive Systems and the Evolution of Reciprocation

Karl Sigmund

*Institute for Mathematics, University of Vienna, Strudlhofgasse 4, A-1090 Vienna, Austria
and IIASA, Laxenburg, Austria*

ABSTRACT

Complex adaptive systems play a major role in the theory of reciprocal altruism. Starting with Axelrod's celebrated computer tournaments, a wide variety of computer simulations show that cooperation can evolve in populations of

selfish agents, both with direct and indirect reciprocation.

Key words: reciprocal altruism; evolutionary games; Prisoner's Dilemma game; evolution of cooperation.

INTRODUCTION

Most of the major transitions in evolution consist in assembling units of some type to form a new, hierarchically higher entity (Maynard Smith and Szathmari 1995). As a rule, the reproductive potential is monopolized by a minority of the previous units. This is most notably the case with multicellular organisms, where only germline cells transmit the genome, or with colonies of eusocial animals having large proportions of sterile workers. This channeling of reproductive opportunities entails a close relatedness between elements of the higher-order unit and enables cooperation based on kin selection.

In human societies, however, where cooperation is ubiquitous, there are remarkably few individual inequalities in reproductive potential. A few potentates have managed to obtain an almost unlimited control of their community and to sire several hundred offspring, but these are exceptions that occurred at a late, and probably transient, stage of cultural evolution. Both in modern mass societies and in bands of hunter-gatherers, social rules tend to level reproductive opportunities and to prevent the establishment of a global controller. Although a large part of the services and tasks performed within households can be explained by kin selection, the bulk of human cooperation is based, not on related-

ness, but on reciprocation. Not genetic ties, but economic exchanges, explain the cohesion of human communities.

This implies that individuals collaborate only if it is to their own good. The mathematical framework for studying the economics of interacting egoists is game theory: more precisely, that branch of game theory which is, somewhat misleadingly, called noncooperative game theory [cf. Binmore (1994)]. The term *noncooperative* means in this context that players cannot negotiate binding and enforceable agreements. Such players can nevertheless achieve cooperation, not by the command of a controller or by deference to the benefit of the group, but by following a myopic set of rules evolved to optimize their selfish interests. The challenge lies in showing how they do it in spite of the ever-present lure of defection.

Robert Trivers (1971) was the first to suggest *reciprocation* as a basis for mutual assistance in animal behavior (including human behavior) and to discuss it in terms of game theory. More precisely, he introduced the Prisoner's Dilemma game to bring the problem into focus [see also May (1987) and Boyd (1988)]. This approach was greatly expanded by Robert Axelrod and William D. Hamilton (1981), who applied evolutionary game theory [see Maynard Smith (1982)] to populations of players engaged in many rounds of the repeated Prisoner's Dilemma game against randomly chosen coplayers. Already at this early stage, computer simulations were essential to follow the complex nonlinear

dynamics of the frequencies of strategies submitted to natural selection (Axelrod 1984).

A further decisive step was taken when Axelrod adapted the genetic algorithms of John Holland (1975) to simulate the effects of evolutionary trial and error. This was the first application of genetic algorithms to a genuine evolutionary problem (rather than a technical optimization problem). It proved particularly successful because the binary coding of strategies, usually a moot point with genetic algorithms, was straightforward (Axelrod 1987) [see also Axelrod (1997) and Epstein and Axtell (1996)].

This success led to many further investigations, including the effects of noise, extending the memory, allowing for alternating moves, introducing more complex state-based strategies, increasing the number of players, changing the payoff structure, restricting the interactions to neighbors only, addressing continuous versions of the game, and so on. We present a short survey of this field, with pointers to the biological motivations behind these variants. The basic message is that details matter, but that cooperation robustly emerges from a bottom-up approach based on reciprocation. We finally deal with indirect reciprocity, which is, according to Richard Alexander (1987), the biological basis of human moral systems. In this setup, an act of assistance may be returned, not to the donor, but to a third party. A model based on the status of the players shows that cooperation can emerge even if any two individuals never interact more than once.

This brief survey concentrates on papers dealing with pairwise interactions and using evolutionary simulations. This is certainly not intended to deny the importance of other approaches to the evolution of cooperation [see Dugatkin and others 1992].

DIRECT RECIPROCATION

The Prisoner's Dilemma (PD) is a two-player game where both players have the same two strategies and the same payoffs. The two strategies are **C** (to cooperate) and **D** (to defect). If both players use **C**, both get the *reward* R for mutual cooperation; if both play **D**, both get the *penalty* P for mutual defection; a **D** player obtains the *temptation* payoff T for unilaterally defecting against a coplayer who plays **C**, whereas the coplayer receives the *sucker's* payoff S for being exploited. One assumes $T > R > P > S$ and $2R > T + S$. The first condition implies that **D** dominates **C**, in the sense that it is better no matter what the other player chooses, and the second condition entails that joint cooperation is better than sharing the payoffs after a unilateral defection.

Players have to choose **D** if they want to maximize their payoff, and this yields the payoff P , which is less than the reward R for mutual cooperation.

In the repeated PD game, players have to choose simultaneously, in every round, whether to play **C** or **D**. There is a constant probability w for another round, so that the average length of the game is random. The total payoff is given by $\sum A_n w^n$, where A_n is the payoff in round n . The strategies for the repeated PD can be arbitrarily complex, but in a biological context, it only makes sense to consider strategies given by simple knee-jerk rules, like Tit For Tat (TFT, which plays **C** in the first round and then simply repeats the coplayer's previous move). The first lesson from Axelrod's computer tournaments was that such simple rules, and in particular TFT, more than hold their own against more sophisticated strategies (based on stochastic optimization, for instance).

Complex adaptive systems are used to simulate artificial societies of players engaged in repeated PD games. In Axelrod (1987), strategies were considered that played **C** or **D** depending on the outcome of the previous three rounds. Binary strings of length 70 coded them. For every new generation, these strings were submitted to point mutation and recombination. The players then engaged in a round-robin tournament, accumulating payoff points that translated into number of offspring forming the next generation. An initial increase in defectors was often followed by a reemergence of cooperation.

Bendor and colleagues (1991), Nowak and Sigmund (1992), and Kollock (1993) have considered the influence of noise, which is particularly disruptive in a society dominated by TFT. When one allows only for reactive strategies given by different propensities to play **C**, depending on the last move of the coplayer, one finds that cooperation emerges based on a tolerant rule called GTFT (Generous TFT), which retaliates only with a certain probability after a **D**, but always cooperates after a **C** (up to mistakes in implementation). This rule cannot spread in a society of defectors, however; it is necessary first that TFT invades and paves the way, like a pioneering species in a plant community, for GTFT to take over (Nowak and Sigmund 1992).

If one considers strategies that depend on the moves of both players in the previous round, one finds a rich collection of far-from-equilibrium dynamics even if only three or four strategies interact (Nowak and Sigmund 1993a). Sufficiently long mutation-selection chronicles usually lead to metastable cooperative regimes, however, based (in case $P + T < 2R$) on Pavlov, a rule which cooperates if both players used the same move in the previous

round. Pavlov is a win-stay, lose-shift rule: players repeat their previous move if it led to a high payoff (R or T) and try the other option if the payoff was low (S or P). In a population of Pavlov players, unilateral defections due to mistakes cause one round of mutual defection, after which both players resume cooperation (Milinski 1993; Nowak and Sigmund 1993b). Pavlov exploits unconditional cooperators, preventing them from invading and thus offering targets for exploiters. Again, cooperation can emerge only after the invasion of stern retaliatory strategies like TFT, which then are superseded by the error-proof Pavlov. TFT acts in this sense as a catalyst for cooperation.

It should be stressed that these simulations of artificial populations based on extended mutation-selection chronicles show a high degree of history dependence and often display punctuated equilibrium. This is particularly clear in Lindgren (1991), where extensions of the memory were possible (caused by mutations reminiscent of gene duplications that introduce strategies based not just on the last round, but on a larger window of the past). Usually such simulations led to the emergence of cooperative strategies similar to Pavlov, or of variants defecting twice in a row, after a mistake.

There exist very simple finite-state automata that cannot be described by rules depending only on the outcome of a given number of previous rounds. An example is Contribute TFT (CTFT), a strategy that monitors its own standing and that of the coplayer. A player's standing is good except after defecting against a player with good standing. CTFT players cooperate except if they are in good standing and their coplayer is in bad standing. Boerlijst and colleagues (1997) showed that such strategies are good at invading populations of defectors and establishing a stable cooperative regime. These strategies, however, are immune only against errors in implementation, whereas the Pavlov rule also avoids errors in perception. Leimar (1997) showed that there are huge numbers of finite-state automata leading to limit-ESS (a version of evolutionarily stable strategies). This implies pronounced path dependence of evolutionary chronicles.

In the usual PD setup, both players are supposed to move at once. This is the case, for instance, in predator-inspection games by sticklebacks or guppies, where the fish take their mirror image for a coplayer (Milinski 1987). In many situations of mutual aid, however, players move alternately. This is the case, for instance, when a well-fed vampire bat feeds a hungry conspecific (Wilkinson 1984) or when a young male baboon diverts the attention of the dominant male so that his pal can mount an

estrous female (Packer 1977). Simulations of the alternating PD (Nowak and Sigmund 1994; Frean 1994; Hauert and Schuster 1998) lead again to cooperation based on error-proof strategies (not Pavlov, however, but Firm But Fair, a strategy that defects only after an unwarranted defection by the coplayer). Experiments by Wedekind and Milinski (1996) support this distinction between strategies for the simultaneous and the alternating PD.

If players do not interact at random, but only within a certain neighborhood structure, cooperation becomes much more readily established, even for the one-shot PD [see Nowak and May (1992, 1993), Sigmund (1992), Huberman and Glance (1993), Nowak and others (1994a, 1994b), Lindgren and Nordahl (1994), and Killingback and Doebeli (1996) and, for a general introduction to cellular automata in this context, Durrett and Levin (1994)]. Again, if players have the possibility of choice and refusal of partners, cooperation becomes much easier to achieve (Stanley and others 1994). In both cases, aggregation of cooperators occurs through local rules. In contrast, cooperation becomes much more unlikely if more than two players engage in the interaction. For extensive simulations of this setup, we refer to Hauert and Schuster (1997).

It should be stressed that the PD is certainly not the only game modeling aspects of cooperation. In particular, Sugden (1986) has stressed that the payoff rank ordering $T > R > S > P$ also makes sense in this context. In such an interaction, a player faced with a defector would nevertheless prefer to play C. This yields what game theorists call the Chicken game and biologists Hawk-Dove. Repeated Chicken need not lead to mutual cooperation based on reciprocity; it is more likely that an asymmetry becomes effective with one player always playing C and the other D. This could shed some light on the intruder inspections by female lions, where one can often distinguish leaders and laggards [see Heinsohn and Packer (1996)].

INDIRECT RECIPROCATION

In addition to reciprocation based on repeated interactions within a pair, there exists another, indirect reciprocity, where the donor does not obtain a return from the recipient, but from a third party. Donors provide help if the recipient has helped others in the past. This works if the cost of an altruistic act is offset by a raised "score," or status, which increases the chance to subsequently become the recipient of an altruistic act. Cooperation is channeled toward the "valuable" members of the community. For Richard Alexander, "indirect reci-

procuity involves reputation and status, and results in everyone in the group continually being assessed and re-assessed."

Nowak and Sigmund (1998a,b) modeled indirect reciprocity in a population of individuals who had the options of helping or not helping another individual. In each generation, a number of potential donor–recipient pairs are chosen randomly: if help is provided, there is a cost c to the donor, a benefit b to the recipient, and the donor's score is increased by 1. The score of a player refusing to help is decreased by 1. Initially all scores are 0. We consider strategies given by integers k : a player with such a strategy helps if and only if the score of the potential recipient is at least k . We can follow the frequencies of the strategies from generation to generation, allowing for occasional mutations.

A remarkably small number of interactions (for $b = 10$ and $c = 1$, an average of two interactions per lifetime suffices, for instance) can lead to the emergence of cooperative populations where most members use $k = 0$ or $k = -1$. If the simulation is continued, strategies that are less discriminating spread: players with $k = -3$, for instance, will rarely ever refuse to help, their score will therefore increase faster than average, and hence they will in turn be helped more often. But if the frequency of less discriminating players reaches a certain threshold, then defectors (players with $k = 3$, for instance, who hardly ever provide help) can take over, with the result that cooperation disappears in the population. Once this happens, the average k values will decrease again, leading eventually back to a cooperative regime of players with maximal discrimination (that is, $k = 0$). To summarize, random drift can subvert populations of discriminate altruists by indiscriminate altruists; once their frequency is large, defectors can invade; but as soon as the defectors have reduced the proportion of indiscriminate altruists, the discriminate altruists can fight back and eliminate the defectors. This leads again to a cooperative population that is proof against defectors, but not against indiscriminate altruists, and so on.

Such models show that indirect reciprocity based on image scoring works in principle. It should be stressed that the chance of two players ever meeting again is vanishingly small. All that needs to be known is the score of the coplayer [see also Pollock and Dugatkin (1992)]. Even this image scoring need not be public knowledge. One can modify the model to include, for every player, a private assessment of the other group members. In particular, one can assume that an interaction between two individuals is only observed by a small subset of the population. Only these onlookers will update their score of the

donor. For larger groups, it becomes more difficult to establish cooperation.

Models that are even more simplified help to explain analytically the cycling behavior, with its long bouts of cooperation interspersed by short periods of defection, which is reminiscent of the lack of stability near a critical state. Somewhat surprisingly, cooperation is more robust if the society is challenged more frequently by invasion attempts of defectors. One can compute the minimal amount of discriminators, the minimal number of rounds per generation, and the maximal size of the society, for indirect reciprocity to work. This yields as necessary condition for cooperation that the *degree of acquaintanceship* (the probability that a player knows the score of the coplayer) is larger than the cost–benefit ratio c/b . This result is analogous to Hamilton's rule which states that the *degree of relatedness* (the probability that an allele in the player's genome is also present in the coplayer) must exceed c/b .

DISCUSSION

Success in analyzing the iterated PD should not hide the fact that variants with more than two players lead only rarely to a cooperative outcome [see Hauert and Schuster (1997)]. Such N-person games have attracted much attention, for instance, as the Free Rider Problem or as the Tragedy of the Commons (Hardin 1968). The latter name suggests already, as a major application, the management of ecosystems.

More generally, how can individual restraint in the exploitation of a common resource emerge and establish itself in the absence of a global controller? This problem underlies the evolution of virulence or the concept of the prudent predator. In the human context, or more generally among individuals recognizing each other, the obvious solution—the punishment of selfish exploiters—leads to another dilemma. To punish defectors is a costly, possibly dangerous activity, and it is tempting to leave it to others. But this constitutes a second-order defection, which ought also to be punished, and so on. Several authors (Boyd and Richerson 1988; Gance and Huberman 1994) have addressed this problem. It seems plausible that the idea of indirect reciprocity can be applied in this context.

REFERENCES

- Alexander RD. 1987. The biology of moral systems. New York: Aldine de Gruyter.
- Axelrod R. 1984. The evolution of cooperation. Basic Books. New York [reprinted 1989, Harmondsworth: Penguin].

- Axelrod R. 1987. The evolution of strategies in the iterated Prisoner's Dilemma. In: Lawrence D, editor. Genetic algorithms and simulated annealing. London: Pitman, p 32-41.
- Axelrod R. 1997. The complexity of cooperation. Princeton: Princeton University Press.
- Axelrod R, Hamilton WD. 1981. The evolution of cooperation. *Science (Washington)* 211:1390-6.
- Bendor J, Kramer RM, Stout S. 1991. When in doubt... cooperation in a noisy Prisoner's Dilemma. *J Conflict Resolution* 35:691-719.
- Binmore KG. 1994. Fun and games. Lexington (MA): Heath.
- Boerlijst M, Nowak MA, Sigmund K. 1997. The logic of contribution. *J Theor Biol* 185:281-93.
- Boyd R. 1988. Is the repeated Prisoner's Dilemma a good model of reciprocal altruism? *Ethol Sociobiol* 9:211-22.
- Boyd R, Richerson P. 1988. The revolution of reciprocity in sizable groups. *J Theor Biol* 132:337-56.
- Dugatkin LA, Mesterton-Gibbons M, Houston AI. 1992. Beyond the Prisoner's Dilemma: toward models to discriminate among mechanisms of cooperation in nature. *TREE* 7:202-5.
- Durrett R, Levin SA. 1994. The importance of being discrete (and spatial). *Theor Popul Biol* 46:363-94.
- Epstein JM, Axtell R. 1996. Growing artificial societies. Cambridge, MA: Brookings Institution Press.
- Frean MR. 1994. The Prisoner's Dilemma without synchrony. *Proc R Soc Lond [B]* 257:75-9.
- Glance NS, Huberman BA. 1994. The dynamics of social dilemmas. *Sci Am* 270:76-81.
- Hardin G. 1968. The tragedy of the commons. *Science (Washington)* 162:1243-8.
- Hauert Ch, Schuster HG. 1997. Effects of increasing the number of players and memory size in the iterated Prisoner's Dilemma: a numerical approach. *Proc R Soc Lond [B]* 264:513-9.
- Hauert Ch, Schuster HG. 1998. Extending the iterated Prisoner's Dilemma without synchrony. *Forthcoming*.
- Heinsohn R, Packer C. 1995. Complex cooperative strategies in group-territorial African lions. *Science (Washington)* 269:1260-2.
- Holland J. 1975. Adaptation in natural and artificial systems. Ann Arbor: University of Michigan Press.
- Huberman BA, Glance NS. 1993. Evolutionary games and computer simulations. *Proc Natl Acad Sci USA* 90:7712-5.
- Killingback T, Doebeli M. 1996. Spatial evolutionary game theory: Hawks and Doves revisited. *Proc R Soc Lond [B]* 263:1135-44.
- Kollock P. 1993. An eye for an eye leaves everyone blind. *Am Sociol Rev* 58:768-86.
- Leimar O. 1997. Repeated games: a state space approach. *J Theor Biol* 184:471-98.
- Lindgren K. 1991. Evolutionary phenomena in simple dynamics. In: Langton CG, Taylor C, Farmer JD, Rasmussen S, editors. *Artificial life II*. Redwood City (CA): Addison-Wesley, p 295-312.
- Lindgren K, Nordahl MG. 1994. Evolutionary dynamics of spatial games. *Physica D* 75:292-309.
- May RM. 1987. More evolution of cooperation. *Nature (Lond)* 327:15-7.
- Maynard Smith J. 1982. Evolution and the theory of games. Cambridge: Cambridge University Press.
- Maynard Smith J, Szathmary E. 1995. The major transitions in evolution. Oxford: WH Freeman.
- Milinski M. 1987. Tit For Tat in stickleback and the evolution of cooperation. *Nature (Lond)* 325:434-5.
- Milinski M. 1993. Cooperation wins and stays. *Nature (Lond)* 364:12-3.
- Nowak MA, Bonhoeffer S, May RM. 1994a. Spatial games and the maintenance of cooperation. *Proc Natl Acad Sci USA* 91:4877-81.
- Nowak MA, Bonhoeffer S, May RM. 1994b. More spatial games. *Int J Bifurcation Chaos* 4:33-56.
- Nowak MA, May RM. 1992. Evolutionary games and spatial chaos. *Nature (Lond)* 359:826-9.
- Nowak MA, May RM. 1993. The spatial dilemmas of evolution. *Int J Bifurcation Chaos* 3:35-78.
- Nowak MA, Sigmund K. 1992. Tit for Tat in heterogeneous populations. *Nature (Lond)* 355:250-3.
- Nowak MA, Sigmund K. 1993a. Chaos and the evolution of cooperation. *Proc Natl Acad Sci USA* 90:5091-4.
- Nowak MA, Sigmund K. 1993b. Win-stay, lose-shift outperforms tit-for-tat. *Nature (Lond)* 364:56-8.
- Nowak MA, Sigmund K. 1994. The alternating Prisoner's Dilemma. *J Theor Biol* 168:219-26.
- Nowak MA, Sigmund K. 1998. Evolution of indirect reciprocity by image scoring. *Nature (Lond)* 393:573-77.
- Nowak MA, Sigmund K. The dynamics of indirect reciprocity. *Forthcoming*.
- Packer C. 1977. Reciprocal altruism in *Papio anubis*. *Nature (Lond)* 265:441-3.
- Pollock GB, Dugatkin LA. 1992. Reciprocity and the evolution of reputation. *J Theor Biol* 159:25-37.
- Sigmund K. 1992. On prisoners and cells. *Nature (Lond)* 359:774.
- Stanley EA, Ashlock D, Tesfatsion L. 1994. Iterated Prisoner's Dilemma with choice and refusal of partners. In: Langton CG, editor. *Artificial life III*. Redwood City (CA): Addison-Wesley, p 194-210.
- Sugden R. 1986. The evolution of rights: co-operation and welfare. Oxford: Blackwell.
- Trivers R. 1971. The evolution of reciprocal altruism. *Q Rev Biol* 46:35-57.
- Wedekind C, Milinski M. 1996. Human cooperation in the simultaneous and the alternating Prisoner's Dilemma: Pavlov versus generous tit-for-tat. *Proc Natl Acad Sci USA* 93:2686-9.
- Wilkinson GS. 1984. Reciprocal food-sharing in the vampire bat. *Nature (Lond)* 308:181-4.