

finally the association of groups, habitually scattered, but uniting in case of need, as we saw it with the bisons and other ruminants. It also takes higher forms, guaranteeing more independence to the individual without depriving it of the benefits of social life. With most rodents the individual has its own dwelling, which it can retire to when it prefers being left alone; but the dwellings are laid out in villages and cities, so as to guarantee to all inhabitants the benefits and joys of social life. And finally, in several species, such as rats, marmots, hares, etc., sociable life is maintained notwithstanding the quarrelsome or otherwise egotistic inclinations of the isolated individual. Thus it is not imposed, as is the case with ants and bees, by the very physiological structure of the individuals; it is cultivated for the benefits of mutual aid, or for the sake of its pleasures. * * *

MARTIN A. NOWAK, ROBERT M. MAY,
AND KARL SIGMUND

The Arithmetics of Mutual Help (1995)†

The principle of give and take pervades our society. It is older than commerce and trade. All members of a household, for example, are engaged in a ceaseless, mostly unconscious bartering of services and goods. Economists have become increasingly fascinated by these exchanges. So have biologists, who have documented many comparable instances in groups of chimpanzees and other primates. Charles Darwin himself was well aware of the role of cooperation in human evolution. In *Descent of Man* he wrote that "the small strength and speed of man, his want of natural weapons, & c., are more than counterbalanced by his . . . social qualities, which lead him to give and receive aid from his fellow-men."

Obviously, this is a far cry from the savage human existence that the philosopher Thomas Hobbes described as "solitary, poor, nasty, brutish, and short." Nevertheless, a number of Darwin's early followers emphasized the ferocious aspects of the "struggle for survival" to such an extent that the Russian prince Kropotkin felt compelled to write a book to refute them. In *Mutual Aid*, hailed by the *London Times* as "possibly the most important book of the year" (1902), he drew a vast fresco of cooperation acting among Siberian herds, Polynesian islanders and medieval guilds. Kropotkin was a famous ideologue of anarchism, but his dabbling in natural history was no mere hobby; for someone bent on getting rid of the State, it was essential to show that human cooperation was not imposed from an iron-fisted authority but had its origins rooted in natural conditions.

In a way, his arguments have succeeded far beyond what Kropotkin could ever have foreseen. A wealth of studies in anthropology and pri-

Handwritten note:
Mutual aid is a result of

† From "The Arithmetics of Mutual Help," *Scientific American* (June 1995): pp. 76-81. Reprinted with permission. Copyright © 1995 by Scientific American, Inc. All rights reserved. Martin A. Nowak (b. 1965) is professor of theoretical biology at the Institute for Advanced Study, Princeton, N.J. Robert M. May (b. 1936) is Royal Society research professor at Oxford University and at Imperial College, London. Karl Sigmund (b. 1945) is a professor at the Institute of Mathematics, University of Vienna.

matology point to the overwhelming role of reciprocal help in early hominid societies. Textbooks on animal behavior are filled with examples of mutual aid: grooming, feeding, teaching, warning, helping in fights and joint hunting. In ecology, symbiotic associations are increasingly seen as fundamental. Biologists find examples of cooperation at the level of cells, organelles and even prebiotic molecules.

But at the same time, the ubiquity of cooperation seems to have become ever more paradoxical. The Russian anarchist had failed to see how threatened it is by exploitation. What prevents mutualists from turning into parasites? Why should anyone share in a common effort rather than cheat the others? Natural selection puts a premium on individual reproductive success. How can this mechanism shape behavior that is altruistic in the sense that it benefits others at the expense of one's own progeny?

There are two main approaches to this question that go under the headings of kin selection and reciprocal aid. These concepts are not mutually exclusive, but they are sharply distinct. Kin selection is rooted in genetics. If a gene helps in promoting the reproductive success of close relatives of its bearer, it helps in promoting copies of itself. Within a family, a good turn is its own reward. But a good turn to an unrelated fellow being has to be returned in order to pay off. Reciprocal aid—the trading of altruistic acts in which benefit exceeds cost—is essentially an economic exchange. It works less directly than kin selection and is therefore more vulnerable to abuse.

Two parties can strike a mutually profitable bargain, but each could gain still more by withholding its contribution. In modern society an enormous apparatus of law and enforcement makes the temptation to cheat resistible. But how can reciprocal altruism work in the absence of those authoritarian institutions so despised by Kropotkin's anarchists? This difficult question is best answered by first considering simple, idealized systems.

The Prisoner's Dilemma

To demonstrate the conundrum, Robert L. Trivers, a sociobiologist (and, fittingly, a former lawyer), now at the University of California at Santa Cruz, borrowed a metaphor from game theory known as the Prisoner's Dilemma. As originally conceived in the early 1950s, each of two prisoners is asked whether the other committed a crime; their level of punishment depends on whether one, both or neither indicates the other's guilt. This situation can be viewed as a simple game. The two players engaged in it have only to decide whether they wish to cooperate with each other or not. In one illustration of the Prisoner's Dilemma, if both choose to cooperate, they get a reward of three points each. If both defect (by not cooperating), they get only one point each. But if one player defects and the other cooperates, the defector receives five points, whereas the player who chose to cooperate receives nothing.

Will they cooperate? If the first player defects, the second who cooperates will end up with nothing. Clearly, the second player ought to have defected. In fact, even if the first player cooperates, the second should

the ph.
of parasites?
X X

Kin selection
reciprocal aid

both cooperate 3pts each 3 years each
both defect 1pt each 1 year each

defect, because this combination gives five points instead of three. No matter what the first player does, the second's best option is to defect. But the first player is in exactly the same position. Hence, both players will choose to defect and receive only one point each. Why didn't they cooperate?

The prisoners' decisions highlight the difference between what is best from an individual's point of view and from that of a collective. This conflict endangers almost every form of cooperation, including trade and mutual aid. The reward for mutual cooperation is higher than the punishment for mutual defection, but a one-sided defection yields a temptation greater than the reward, leaving the exploited cooperator with a loser's payoff that is even worse than the punishment. This ranking—from temptation through reward and punishment down to the loser's payoff—implies that the best move is always to defect, irrespective of the opposing player's move. The logic leads inexorably to mutual defection.

Most people feel uneasy with this conclusion. They do often cooperate, in fact, motivated by feelings of solidarity or selflessness. In business dealings, defection is also relatively rare, perhaps from the pressure of society. Yet such concerns should not affect a game that encapsulates life in a strictly Darwinian sense, where every form of payoff (be it calories, mates or safety from predators) is ultimately converted into a single currency: offspring.

Virtual Tournaments

One can conceive a thought experiment in which an entire population consists of programmed players. Each of these automata is firmly wedded to a fixed strategy and will either always cooperate or always defect. They engage in a round-robin tournament of the Prisoner's Dilemma. For each contestant, the total payoff will depend on the other players encountered and therefore on the composition of the population. A defector will, however, always achieve more than a cooperator would earn in its stead. At the end of the imaginary tournament, the players reproduce, creating progeny of their own kind (defectors or cooperators). The next generation will, again, engage in a round-robin competition and get paid in offspring, and so on. In this caricature of biological evolution, where the payoff is number of offspring and strategies are inherited, the outcome is obvious: defectors will steadily increase from one generation to the next and will eventually swamp the population.

There are several ways to escape from this fate. In many societies the same two individuals interact not just once but frequently. Each participant will think twice about defecting if this move makes the other player defect on the next occasion. So the strategy for the repeated game can change in response to what happened in previous rounds.

In contrast to a single instance of the Prisoner's Dilemma, where it is always better to defect, countless strategies for the repeated version exist, and none serves as a best reply against all opponents. If the opposite player, for instance, decides always to cooperate, then you will do best by always defecting. But if your adversary decides to cooperate until you defect and

①
defectors
win
②
mutual cooperation

then never to cooperate again, you will be careful not to spoil your partnership: the temptation to cheat in one round and grab five points instead of three will be more than offset by the expected loss in the subsequent rounds where you cannot hope to earn more than one point.

The absence of a best choice is crucial. There is no hard-and-fast recipe for playing the repeated Prisoner's Dilemma. Success will depend on the other player's strategy, which one does not know beforehand. A strategy that does well in certain environments can fail miserably in others.

In the late 1970s the political scientist Robert Axelrod, at the University of Michigan, conducted round-robin tournaments of the repeated Prisoner's Dilemma on his computer. The contestants—programs submitted by colleagues—were quite sophisticated, but it turned out that the simplest entry ultimately won. This strategy is aptly called Tit-for-Tat. It starts with a cooperative response and then always repeats the opposing player's previous move.

Remarkably, a player applying Tit-for-Tat is never ahead at any stage of the repeated game, being always last to defect. The Tit-for-Tat player can nonetheless win the whole tournament, because the Prisoner's Dilemma is not a zero-sum game: it is possible to make points without taking them away from others. By its transparency Tit-for-Tat frequently persuades opponents that it pays to cooperate. In Axelrod's tournaments the Tit-for-Tat strategy (entered by the game theorist Anatol Rapoport) elicited many rewarding rounds of cooperation, whereas other players, among themselves, were apt to get bogged down in long runs of defection.

By winning the round-robin tournament, Tit-for-Tat obtained more representatives among the next generation than did other strategies. Moreover, those players who had cooperated tended also to receive more offspring than those who had not. With each generation Tit-for-Tat shaped a more congenial environment. The strategies that ruthlessly exploited cooperators succeeded only in depleting their own resources.

Unpredictable Adversaries

We recently performed computer simulations with an extended set of strategies that base their next move on the result of the previous round rather than just the opponent's previous move (as does Tit-for-Tat). A strategy based on prior outcome must determine the response for each of four eventualities: temptation, reward, punishment or loss. Two possible responses for each of four prior outcomes give 16 possible types of players.

We further allowed for "stochastic" strategies that respond to the four possible outcomes by changing only their statistical propensity to cooperate. Such strategies are not obliged to respond always in the same way to a given outcome. One form of stochastic player might, for example, cooperate 90 percent of the time after experiencing the reward. Such uncertainty simulates the inevitable mistakes that occur during real interactions.

The addition of stochastic responses resulted in a huge array of possibilities. Our computer searched for the most successful of these players by simulating the forces of natural selection, adding to every hundredth gen-

CC2:
Tit for Tat
wins

3 previous
round

eration some small amount of a new, randomly selected stochastic strategy. We followed many such mutation-selection rounds for millions of generations, not because the emergence of cooperation needed so many iterations but because this span allowed us to test a very large number of possible strategies. In spite of the rich diversity displayed in these chronicles, they led us invariably to some simple, clear results. The first is that the average payoff in the population can change suddenly. Indeed, the behavior we found is a showpiece for punctuated equilibria in biological evolution. Most of the time, either almost all members of the population cooperate, or almost all defect. The transitions between these two regimes are usually rare and abrupt, taking just a few generations. We found that later in the run, quiescent periods tended to last longer. And there was a definite trend toward cooperation. The longer the system was allowed to evolve, the greater the likelihood for a cooperative regime to blossom. But the threat of a sudden collapse always remained. * * *

CC3:
punctuated
equilibrium

Innate Cooperation

One can safely conclude that the emergence and persistence of cooperative behavior are not at all unlikely, provided the participants meet repeatedly, recognize one another and remember the outcomes of past encounters. These circumstances may seem familiar from daily life in the home or office, but among the larger world of living things, such requirements demand a high degree of sophistication. And yet we observe cooperation even among simple organisms that do not possess such abilities. Furthermore, the strategies discussed will work only if benefits from future encounters are not significantly discounted as compared with present gains. Again this expectation may be reasonable for many of the activities humans conduct, but for most simpler organisms delayed payoffs in the form of future reproductive success may count for little: if life is short and unpredictable, there is scant evolutionary pressure to make long-term investments.

conclusion
X
Po with
animals

But what of the creatures, such as many invertebrates, that seem to exhibit forms of reciprocal cooperation, even though they often cannot recognize individual players or remember their actions? Or what if future payoffs are heavily discounted? How can altruistic arrangements be established and maintained in these circumstances? One possible solution is that these players find a fixed set of fellow contestants and make sure the game is played largely with them. In general, this selectivity will be hard to attain. But there is one circumstance in which it is not only easy, it is automatic. If the players occupy fixed sites, and if they interact only with close neighbors, there will be no need to recognize and remember, because the other players are fixed by the geometry. Whereas in many of our simulations players always encounter a representative sample of the population, we have also looked specifically at scenarios in which every player interacts only with a few neighbors on a two-dimensional grid. Such spatial games are very recent. They give an altogether new twist to the Prisoner's Dilemma.

spatial
games
X

Fixed in Flatland

It should come as no surprise that cooperation is easier to maintain in a sedentary population: defectors can thrive in an anonymous crowd, but mutual aid is frequent among neighbors. That concept is clear enough. But in many cases, territorially structured interactions promote cooperation, even if no follow-up encounter is expected. This result favors cooperation even for the seemingly hopeless single round of the Prisoner's Dilemma. * * *