

Why do we value fairness and cooperation over seemingly normal rational selfishness?
How can Darwinian generosity arise ?
Biologists and economists explain.

The Economics of Fair Play

By Karl Sigmund, Ernest Fehr and Martin A. Nowak (Scientist America, 2002)

Imagine that somebody offers you \$100. All you have to do is agree with some other anonymous person on how to share the sum. The rules are strict. The two of you are in separate rooms and cannot exchange information. A coin toss decides which of you will propose how to share the money. Suppose that you are the proposer. You can make a single offer of how to split the sum, and the other person – the responder – can say yes or no. The responder also knows the rules and the total amount of money at stake. If her answer is yes, the deal goes ahead. If her answer is no, neither of you gets anything. In both cases, the game is over and will not be repeated. What will you do?

Instinctively, many people feel they should offer 50 percent, because such a division is “fair” and therefore likely to be accepted. More daring people, however, think they might get away with offering somewhat less than half of the sum.

Before making a decision, you should ask yourself what you would do if you were the responder. The only thing you can do as the responder is say yes or no to a given amount of money. If the offer were 10 percent, would you take \$10 and let someone walk away with \$90, or would you rather have nothing at all? What if the offer were only 1 percent? Isn't \$1 better than no dollars? And remember, haggling is strictly forbidden. Just one offer by the proposer : the responder can take it or leave it.

So what will you offer?

You may not be surprised to learn that two thirds of offers are between 40 and 50 percent. Only four in 100 people offer less than 20 percent. Proposing such a small amount is risky, because it might be rejected. More than half of other responders reject offers that are less than 20 percent. But here is the puzzle: Why should anyone reject an offer as “too small”? The responder has just two choices: take what is offered or receive nothing. The only rational option for a selfish individual is to accept any offer. Even \$1 is better than nothing. A selfish proposer who is sure that the responder is also selfish will therefore make the smallest possible offer and keep the rest. This game-theory analysis, which assumes that people are selfish and rational, tells you that the proposer should offer the smallest possible share and the responder should accept it. But this is not how most people play the game.

The scenario just described, calls the Ultimatum Game, belongs to a small but rapidly expanding field called experimental economics. A major part of economic theory deals with large-scale phenomena such as stock market fluctuations or gross national products. Yet economists are also increasingly fascinated by the most down-to-earth interactions – the sharing and helping that goes on within office pools, households, families and groups of children. How does economic exchange work in the absence of explicit contracts and regulatory institutions?

For a long time, theoretical economists postulated a being called *Homo economicus* – a rational individual relentlessly bent on maximising a purely selfish reward. But the lesson from the Ultimatum Game and similar experiments is that real people are a cross-breed of *H.economicus* and *H. emoticus*, a complicated hybrid species that can be ruled as much by emotions as by cold logic and selfishness. An interesting challenge is to understand how Darwinian evolution would produce creatures instilled with emotions and behaviours that do not immediately seem geared toward reaping the greatest benefit for individuals or their genes.

Werner Güth of Humboldt University in Berlin devised the Ultimatum Game some 20 years ago. Experimenters subsequently studied it intensively in many places using diverse sums. The results proved remarkably robust. Behaviour in the game did not appreciably depend on

the players' sex, age, schooling or numeracy. Moreover, the amount of money involved had surprisingly little effect on results. In Indonesia, for instance, the sum to be shared was as much as three times the subjects' average monthly income – and still people indignantly refused offers that they deemed too small. Yet the range of players remained limited in some respects, because the studies primarily involved people in more developed countries, such as Western nations, China and Japan, and very often university students, at that.

Recently an ambitious cross-cultural study of 15 small-scale societies on four continents showed that there were, after all, sizable differences in the ways some people play the Ultimatum Game. Within the Machiguenga tribe in the Amazon, the mean offer was considerably lower than in typical Western-type civilisations – 26 instead of 45 percent. Conversely, many members of the Au tribe in Papua New Guinea offered more than half the pie. Cultural traditions in gift giving, and the strong obligations that result from accepting a gift, play a major role among some tribes, such as the Au. Indeed, the Au tended to reject excessively generous offers as well as miserly ones. Yet despite these cultural variations, the outcome was always far from what rational analysis would dictate for selfish players. In striking contrast to what selfish income maximisers ought to do, most people all over the world place a high value on fair outcomes.

Numerous situations in everyday life involve trade-offs between selfishness and fair play. A colleague, for example, invites you to collaborate on a project. You will be happy to do it, if you expect a fair return on your investment of time and energy or if he has helped you in the past. The pure Ultimatum Game, however, has artificial constraints that rarely apply in real-life interactions: haggling is impossible, people do not get to know each other, the prize vanishes if not split on the first attempt and the game is never repeated. But such constraints, rather than being a drawback, let us study human behaviour in well-defined situations, to uncover the fundamental principles governing our decision-making mechanisms. The process is somewhat like physicists colliding particles in a vacuum to study their properties.

Getting Emotional

Economists have explored a lot of variations of the Ultimatum Game to find what causes the emotional behaviour it elicits. If, for instance, the proposer is chosen not by a flip of a coin but by better performance on a quiz, then offers are routinely a bit lower and get accepted more easily – the inequality is felt to be justified. If the proposer's offer is chosen by a computer, responders are willing to accept considerably less money. And if several responders compete to become the one to accept a single proposer's offer, the proposer can get away with offering a small amount.

Overview/ Fair Play

- Universally across cultures, people value fairness highly and emotionally. Scientists study these behaviours in experiments such as the Ultimatum Game.
- In these experiments, players are often more generous than is predicted by game-theory analysis, which assumes that people selfishly seek to maximise their gains. Conversely, unfair play incites costly acts of revenge. The result is fairer play than is predicted.
- Humans could have evolved the emotions at work in these situations during the millions of years that we lived in small groups. Such emotions prompt us to behave in ways that would have benefited either us or our group in the long run.

These variations all point to one conclusion: in pairwise encounters, we do not adopt a purely self-centred viewpoint but take account of our co-player's outlook. We are not interested solely in our own payoff but compare ourselves with the other party and demand fair play.

Why do we place such a high value on fairness that we reject 20 percent of a

large sum solely because the co-player gets away with four times as much? Opinions are divided. Some game theorists believe that subjects fail to grasp that they will interact only once. Accordingly, the players see the offer, or its rejection, simply as the first stage of an incipient bargaining process. Haggling about one's share of a resource must surely have been a recurrent theme for our ancestors. But can it be so hard to realise that the Ultimatum Game is a one-shot interaction? Evidence from several other games indicates that experimental subjects are cognitively well aware of the difference between one-shot and repeated encounters.

Others have explained our insistence on a fair division by citing the need, for our ancestors, to be sheltered by a strong group. Groups of hunter-gatherers depended for survival on the skills and strengths of their members. It does not help to out compete your rival to the point where you can no longer depend on him or her in your contests with other groups. But this argument can at best explain why proposers offer large amounts, not why responders reject low offers.

Two of us (Nowak and Sigmund) and Karen M. Page of the Institute of Advanced Study in Princeton, N.J., have recently studied an evolutionary model that suggests an answer: our emotional apparatus has been shaped by millions of years of living in small groups, where it is hard to keep secrets. Our emotions are thus not finely tuned to interactions occurring under strict anonymity. We expect that our friends, colleagues and neighbours will notice our decisions.

If others know that I am content with a small share, they are likely to make me low offers; if I am known to become angry when facing a low offer and to reject the deal, others have an incentive to make me high offers. Consequently, evolution should have favoured emotional responses to low offers. Because one-shot interactions were rare during human evolutions, these emotions do not discriminate between one-shot and repeated interactions. This is probably an important reason why many of us respond emotionally to low offers in the Ultimate Game. We may feel that we must reject a dismal offer in order to keep our self-esteem. From an evolutionary viewpoint, this self-esteem is an internal device for acquiring a reputation, which is beneficial in future encounters.

The Ultimatum Game, in its stark simplicity, is a prime example of the type of games used by experimental economists: highly abstract, sometimes contrived interactions between independent decision makers. The founder of game theory, the Hungarian mathematician John von Neumann (one of the fathers of the computer) and the Austrian economist Oskar Morgenstern, collaborating in Princeton in the 1940s, used parlour games such as poker and chess for illustrating their ideas. Parlour games can certainly be viewed as abstractions of social or economic interactions, but most of these games are zero-sum: the gains of one player are the losses of another. In contrast, most real-life economic interactions are mixed-motive: they display elements of cooperation as well as competition. So-called Public Goods games model that situation.

Revenge Is Sweet

In one of the simplest Public Goods games, four players form a group. The experimenter gives each player \$20, and they have to decide, independently of one another, how much to invest in a common pool. The experimenter doubles the common pool and distributes it equally amount all four group members.

If every player contributes the full \$20, they all double their capital. Cooperation is highly rewarding. But the temptation to hold back on one's own contribution is strong. A selfish player ought to contribute nothing at all, because for every dollar he invests, only 50 cents return to his account. (The money is doubled by the experimenter but then divided by four among the players.) The experimenter makes sure that the players fully understand this, by asking them to figure out how much each would end up with if, say, Alice contributed \$10, Bob and Carol only \$5 each, and Dan nothing at all. After this preparation, the game is played for real. If everyone followed the selfish rational strategy predicted by economics, nothing would be invested and nobody would improve their \$20 stake. Real people don't play that way. Instead many invest at least half of their capital.

If the same group repeats the game for 10 rounds, subjects again invest roughly half of their capital during the first rounds. But towards the end, most group members invest nothing. This downhill slide from a high level of cooperation used to be interpreted as a learning process: players learn the selfish strategy the hard way – through a series of disappointing experiences. But this cannot be the right explanation, because other experiments have shown that most players who find themselves in new groups, with co-players they have not met before, start out again by contributing a lot. What explains these behaviours?

Experiments conducted by one of us (Fehr) and Simon Gächter from the University of St. Gallen in Switzerland show that the Public Goods game takes a dramatic turn if a new option

is introduced – that of punishing the co-players. In these experiments, players may impose fines on their co-players at the end of each round, but only at a cost. If Alice wants to impose a fine of \$1 on Dan, Alice has to pay 30 cents. Both the dollar and the 30 cents go back to the experimenter. The cost makes the act of punishment unjustifiable from the selfish point of view (Alice reduces her capital and gains nothing in return). Nevertheless, most players prove very willing, and even eager, to impose fines on co-players who lag behind in their contributions. Everyone seems to anticipate this, and even in a game of one round, less defection occurs than usual. Most significant, if the game is repeated for a known, preset number of periods, the willingness to contribute does not decline. Quite the contrary – the contributions to the common pool rise over time, and in the last few rounds more than 80 percent of all group members invest the whole capital: a striking difference to the outcome of the game without punishment.

In a repeated game, players can see punishment as a shrewd, selfish investment in co-player education: Tightwads are taught to contribute to the general benefit. Incurring costs to punish cheapskates can yield profits in the long run. But a recent variation of the Public Goods game shows that this economic aspect is only a side issue. In this version, numerous groups of four players are assembled, and after every round players are redistributed so that no two people ever meet twice. The punishment pattern (and also the high level of investments) does not change – free riders are punished as severely as when everyone stays in the same group, and again investments start high and may rise. This result is astonishing, because the “educational payoff” has been eliminated. As before, being fined usually increases a player’s future investment, but this increase never benefits the player who imposes the fine. Nevertheless, a lot of players show great eagerness to punish defectors. Participants seem to experience a primal pleasure in getting even with free riders. They seem to be more interested in obtaining personal revenge than in increasing their overall economic performance.

Why are so many players willing to pay the price to punish free riders without reaping any material benefit from it? Evolutionary economist Herbert Gintis of the University of Massachusetts has recently shown that this behaviour can provide fitness advantages. In his model, social groups with an above-average share of punishers are better able to survive events such as wars, pestilence and famines that threaten the whole group with extinction or dispersal. In these situations, cooperation among self-interested agents breaks down because future interactions among group members are highly improbable. Punishers discipline the self-interested agents so that the group is much more likely to survive. Subjects who punish are not, of course, aware of this evolutionary mechanism. They simply feel that revenge is sweet.

People expect fairness and solidarity within most groups, be it children in a summer camp or capos in the Mafia. Ultimately, moral guidelines determine an essential part of economic life. How could such forms of social behaviour evolve? This is a central question for Darwinian theory. The prevalence of altruistic acts – providing benefits to a recipient at a cost to the donor – can seem hard to reconcile with the idea of the selfish gene, the notion that evolution at its base acts solely to promote genes that are most adept at engineering their own proliferation. Benefits and costs are measured in terms of the ultimate biological currency – reproductive success. Genes that reduce this success are unlikely to spread in a population.

Darwinian Generosity

In social insects, the close relatedness among the individuals explains the huge degree of cooperation. But human cooperation also works among non-relatives, mediated by economic rather than genetic ties. Nevertheless, biologists have shown that a number of apparently altruistic types of behaviour can be explained in terms of biological success. (Others argue that a second form of evolution – an evolution of ideas, or “memes” – is at work. See “The Power of Memes,” by Susan Blackmore; *Scientific American*, October 2000).

It may seem callous to reduce altruism to considerations of costs and benefits, especially if these originate in biological needs. Many of us prefer to explain our generous actions simply by invoking our good character. We feel better if we help others and share with them. But where does this inner glow come from? It has a biological function. We eat and make love because we enjoy it, but behind the pleasure stands the evolutionary program commanding

us to survive and to procreate. In a similar way, social emotions such as friendship, shame, generosity and guilt prod us toward achieving biological success in complex social networks.

Centuries ago philosophers such as David Hume and Jean-Jacques Rousseau emphasized the crucial role of “human nature” in social interactions. Theoretical economists, in contrast, long preferred to study their selfish *Homo economicus*. They devoted great energy to theorising about how an isolated individual – a Robinson on some desert island – would choose among different bundles of commodities. But we are no Robinsons. Our ancestors’ line has been social for some 30 million years. And in social interactions, our preferences often turn out to be far from selfish.

Ethical standards and moral systems differ from culture to culture, but we may presume that they are based on universal, biologically rooted capabilities, in the same way that thousands of different languages are based on a universal language instinct. Hume and Rousseau would hardly be surprised. But today we have reached a stage at which we can formalise their ideas into game-theory models that can be analysed mathematically and tested experimentally.