International Institute for
Applied Systems Analysis
IIASA www.iiasa.ac.at

The Economics of Fairness

Sigmund, K.

IIASA Interim Report
April 2002

International Institute for
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria

Tel: +43 2236 807 342
Fax: +43 2236 71313
E-mail: publications@iiasa.ac.at
Web: www.iiasa.ac.at

Interim Report **IR-02-020**

The Economics of Fairness

Karl Sigmund (karl.sigmund@univie.ac.at)

Approved by

Ulf Dieckmann, dieckman@iiasa.ac.at
Project Leader, Adaptive Dynamics Network

March 2002

Abstract

Experimental games are a small, but expanding branch of economics. The major part of economics deals with large-scale phenomena like stock market fluctuations, rates of exchange, and gross national products. The trend in economic life towards globalisation leads to ever more abstract and virtual forms of interaction, as can be seen by the rapid growth of global e-commerce, trading in licenses and options, and the like. Yet at the same time, paradoxically, economists become increasingly fascinated by interactions at the most down-to-earth level — the sharing and helping that goes on within office pools, households, families, or children groups. How does economic exchange work in the absence of explicit contracts and regulatory institutions?

About the Author

Karl Sigmund
Institute of Mathematics, University of Vienna
Strudlhofgasse 4, A-1090 Vienna, Austria
and
Adaptive Dynamics Network
International Institute for Applied Systems Analysis
A-2361 Laxenburg, Austria

The Economics of Fairness

Karl Sigmund

Experimental games are a small, but expanding branch of economics. The major part of economics deals with large-scale phenomena like stock market fluctuations, rates of exchange, and gross national products. The trend in economic life towards globalisation leads to ever more abstract and virtual forms of interaction, as can be seen by the rapid growth of global e-commerce, trading in licenses and options, and the like. Yet at the same time, paradoxically, economists become increasingly fascinated by interactions at the most down-to-earth level — the sharing and helping that goes on within office pools, households, families, or children groups. How does economic exchange work in the absence of explicit contracts and regulatory institutions?

Centuries ago, philosophers like David Hume or Jean Jacques Rousseau stressed the crucial role of 'human nature' in such social interactions, but many theoretical economists preferred to postulate a being which they called homo economicus — a rational agent relentlessly bent on pursing the maximisation of a purely selfish utility. A lot of theoretizing was spent on working out how an isolated individual — a Robinson on some deserted island — would choose between different bundles of commodities. But we are no Robinsons. Our ancestor's line has been social for some thirty million years. And in social interactions, our preferences often turn out to be far from selfish.

The Ultimatum Game is a particularly concise example. The experimenter offers 100 dollars to a pair of test persons who — this is essential — do not know each other and are unlikely to ever meet again. All they must do is split the sum while keeping to the following rules. First, 'Heads or Tails' decides who of them is going to propose how to share the money. This 'Proposer' has to offer part of the 100 dollars to the co-player. The co-player — the 'Responder' — can agree to the offer, and then receives the proposed share, while the Proposer keeps the rest. If this happens, the game is over. If the Responder rejects the offer, the game is also over: the experimenter pockets the money again and bides his leave. It has been made clear to the players that they are not going to play each other again. Also, all haggling is strictly forbidden. Just one offer by the Proposer, and the Responder can take it or leave it.

This does not seem a very exciting game. Few people are surprised when they learn that a majority of the offers — some sixty to eighty percent — are between 40 and 50 dollars. Only some three to five percent of the offers are below 20 dollars. It is dangerous to offer such small amounts. Indeed, as a rule, Responders reject offers of less than 20 percent of the available sum. Anyone offering less will have reason to fear that the offer will not be accepted. But here is the catch: why should anyone reject an offer as 'too small'? The only rational option for a selfish individual is to accept every offer — even one dollar is better than nothing. And a selfish Proposer who is sure that the Responder is selfish too will therefore make the smallest offer, and keep the rest. But this is not how most people do play.

The Ultimatum Game, first devised some twenty years ago by Werner Güth, has been studied intensively by experimental game theorists, in many different places and for very diverse sums. The results proved remarkably robust. They did not appreciably depend on the sex of the players, for instance, on their age, their schooling or their numeracy. Moreover, they depend only surprisingly little on the amount of money involved. In experiments performed in Russia or Indonesia, for instance, the sum could be as high as three times the average monthly income — and still people would refuse indignantly offers which they deemed too small.

Very recently, an ambitious cross-cultural study funded by the MacArthur foundation in four continents managed to show that there were, after all, sizeable differences in the way people play the Ultimatum Game. Within the Machiguenga tribe in Amazonia, the mean offer was considerably lower than in typical, western-type civilisations — only 26 (instead of 40) percent. On the other hand, many members of another tribe, the Au in Papua NewGuinea, offered even more than half of the pie. Cultural traditions in gift-giving, and the strong obligations that result from accepting a gift, seem to play a major role among some tribes. Nevertheless, the outcome was always far away from what homo economicus should do. Compared to other differences in life-style, the differences in offers were small. In striking contrast to what selfish income-maximisers ought to do, most people all over the world place a high value on fair outcomes.

It is doubtful that the Ultimatum Game is found in everyday economic life, be it ours or that of hominids. The constraints — no repetition, no haggling, and the prize vanishes if not split on the first attempt — are artificial. This is precisely why the game is so relevant to uncover hidden springs to our decision mechanisms — for much the same reason, physicists create vacuum to check that all objects fall with same speed. Economists have explored all kinds of variations of the Ultimatum Game to uncover the causes behind the widespread emotional behaviour it elicits. If, for instance, the better performance in a quiz decides on who is the Proposer, rather than the flip of a coin, then offers are routinely somewhat lower, and get accepted more easily — the inequality is felt to be justified. If the ‘co-player’ is a computer, Responders are willing to accept considerably less. If several Responders compete to become the one to accept a single Proposer’s offer, then the Proposer can get away with offering a small amount only. If, on the other hand, several Proposers have to compete for one Responder — who can accept only one offer, obviously the highest — then the offers become considerably higher.

These variations all point to one conclusion: in pairwise encounters, we do not adopt a purely self-centered viewpoint, but take account of our co-players’ outlook. We are not uniquely interested in our own payoff but compare ourselves with the other and demand equal rights.

Why do we place such a high value on fairness that we reject twenty dollars uniquely because the co-player gets away with four times as much? Opinions are divided. Some game theorists believe that test-persons fail to grasp that they 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 many of our ancestors. But can it really be so hard to realise that the Ultimatum Game is a one-shot interaction? There is evidence from several other experimental games indicating that experimental subjects are cognitively well aware of the difference between one-shot interactions and repeated encounters.

Others have explained our insistence on a fair division by the need, for our ancestors, to be sheltered by a strong group. The groups of hunter-gatherers depended for survival on the skills and strenghts of their members. It does not help to outcompete your rival to the point where you can no longer depend on him or her in your contests with other groups. Belonging to a strong group was often essential for survival. But this argument can at best explain why Proposers offer much, and not why Responders reject low offers.

It is possible that this last fact yields a cue for the ultimatum puzzle. Our emotional apparatus has been shaped by thousands of years of living in small groups where it is hard to keep secrets. Our emotions are thus not fine-tuned to one-shot interactions occurring under strict anonymity. We expect that our decisions will be noticed by our friends, colleagues and neighbours. If others know that I am content with a small share they are likely to make me low offers; and if I am known to become angry when facing a low offer and to reject the deal, the others have an incentive to make me high offers. This means that evolution should have favoured emotional responses to low offers. Due to the fact that one-shot interactions were rare during human evolution these emotions do not discriminate between one-shot and repeated interactions. This is why many of us respond emotionally to low offers in the ultimatum game. We may feel that we must reject a dismal offer in order to keep our self-esteem. But, 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. They are highly abstract, and sometimes even contrived instances of interactions between independent decision makers. The founders of game theory, the Hungarian mathematician John von Neumann (one of the fathers of the computer) and the Austrian economist Oskar Morgenstern, used parlour games like poker or chess for illustrating of their ideas, and these can indeed be viewed as abstractions of social or economic interactions. But most parlour games are zero-sum — the gains of one player are the losses of another — whereas most real-life economic interactions are ‘mixed-motive’: they display both elements of competition and cooperation.

Experimental economists have studied such situations in the context of so-called Public Goods games. In one such game, four players form a group. Each is given at the start 20 dollars by the experimenter. The players have to decide, independently of each other, how much of it to invest in a common pool. The experimenter, then, doubles this common pool, and distributes it equally among all four group members, irrespective of the amount (if any) that they have contributed.

If all players contribute maximally, they double thereby their capital. Thus cooperation is highly rewarding. But the temptation to hold back on one’s own contribution is strong. In fact, a selfish player ought to contribute nothing at all, for the good reason that for every dollar invested, only half a dollar returns to the personal account. (The money is first doubled by the experimenter, and then divided by four among the players.) The experimenter makes sure that the players understand this well, by asking them to figure out how much each would end up with if player A, say, contributed 10 dollars, players B and C only 5 dollars each, and player D nothing at all. After this preparation, the game is played for real. The selfishly rational strategy is, of course, for all to free-wheel. So nobody gains anything. In real experiments, this is not at all how people play. Many invest at least half of their capital.

If the same group repeats the game for ten rounds, subjects invest again roughly 50 percent of their capital during the first rounds. But towards the end, most group members invest nothing so that average investment drops to very low levels. 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.

Experiments conducted by Ernst Fehr and Simon Gächter from the University of Zürich show that the Public Goods game takes a dramatic turn if a new option is introduced — that of punishing the co-players. In their experiments players get the opportunity to impose fines on their co-players at the end of each round. It is not they who collect these fines: this is done by the experimenter. Punishing others is actually made costly by the rules: each dollar imposed as a fine costs the punisher thirty cents (which also go to the experimenter). These costs imply that the act of punishment cannot be justified from the selfish point of view. 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 one round game, there is less defection than usual. Most significantly, however, if the game is repeated for a known, finite, number of periods, no erosion of the willingness to contribute does set in. Quite the contrary, the contributions to the common pool even rise over time and in the last few periods more than 80 percent of all group members invest the whole capital: the difference to the Public Goods game without punishment is absolutely striking.

In a repeated game, punishing co-players can be seen as a measure of education: the co-players are taught to contribute to the general benefit. In this sense, incurring costs to punish others can yield profits in the long run, and hence may be seen as a shrewd, selfish investment. But a recent variation of the Public Goods game shows that this educational aspect is only a side-issue. If the groups of four players are re-arranged from round to round so that the same players never meet for a second time, the punishment pattern does not change — free-riders are punished with the same severity as in the game where the same people stay together in a group for many rounds. Moreover, although the same subjects never meet twice, investments are rather high (roughly 60 percent of the capital) and do not decline over time. This result is astonishing because the punishment reduces the material payoff of the punishing player. A fine increases, in general, the future investment of the fined player, but this increase never benefits the player who imposes the fine. Nevertheless, many players show great eagerness, indeed almost relish, to punish defectors. Participants in this game seem to experience a primal pleasure in ‘getting even’ with free-riders. They seem to feel intense satisfaction, more in obtaining their personal revenge rather than in increasing overall economic performance.

Fairness and solidarity are expected within every group — be it among children in a holiday camp or among 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 the recipient at a cost to the donor — seems hard to reconcile with the idea of the ‘selfish gene’. After all, 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. In social insects, the close relatedness between the individuals

explains the huge degree of cooperation. But human cooperation works also between non-relatives. It is mediated by economic rather than genetic ties. Biologists have shown that many apparently altruistic types of behaviour can be explained in terms of self-interest.

It may seem shocking to reduce altruism to considerations of costs and benefits, especially if these originate in biological needs. Many of us incline to explain our good actions simply by invoking our good character. We feel fine if we help others and share with them. But where does this 'inner glow' come from? It has a biological function. We eat and 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 like friendship, shame, generosity or guilt act to prod us towards achieving biological success in complex social networks.

Ethical standards and moral systems differ between cultures, but we may presume that they are based on universal, biologically rooted capabilities; in the same way as thousands of different languages are based on the same universal language instinct. David Hume or Jean Jacques Rousseau would hardly be surprised. But today, we have reached a stage where we can formalise their ideas into game theoretical models, and test them experimentally.

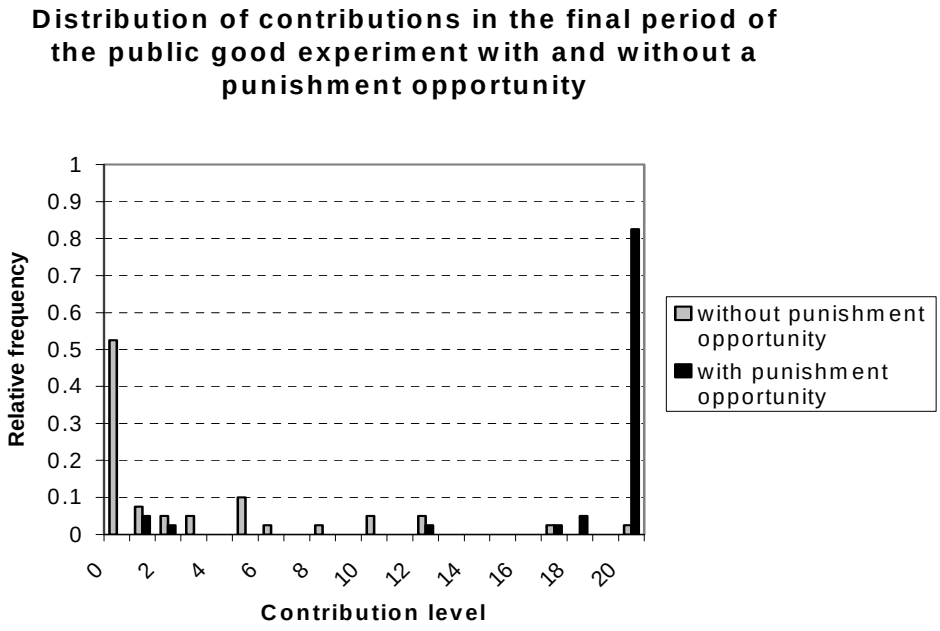

The figure above shows the distributions of contribution in the tenth and last round of a public goods game with and without the opportunity for punishment.There is a striking difference across conditions. Without a punishment opportunity 53 percent of all subjects contribute nothing in the final period and the rest contributes little. In contrast, when there is a punishment opportunity 83 percent of the subjects contribute the whole endowment, i.e., 20 dollars and zero subjects contribute nothing.