

Physicists suggest selfishness can pay

Findings that undermine thinking on the evolution of cooperation face a strong challenge.

Philip Ball

23 August 2012

It was long thought that game theory — a branch of mathematics that examines competitive decision-making — had the question of how altruism evolved all sewn up. But a recent finding shattered that comfortable belief, with a demonstration that selfishness could trump cooperation. Now the cooperators have fought back, with a study suggesting that selfishness still fails in the long term.

For several decades, one of the central results of game theory has seemed to be that self-interest can drive social cooperation, because in the long term, selfish behaviour hurts you as much as your competitors. Last May, two leading physicists, William Press of the University of Texas at Austin and Freeman Dyson of the Institute for Advanced Study in Princeton, New Jersey, argued otherwise. They showed how, in the classic 'game' called the iterated prisoner's dilemma, in which cooperation between players was thought to be a winning strategy, it is possible to be successfully selfish¹.

This apparently revolutionary idea has now been challenged by two evolutionary biologists at Michigan State University in East Lansing. In a preprint on the arXiv server², Christoph Adami and Arend Hintze say that the strategy proposed by Press and Dyson is "evolutionarily unstable". In a population of agents all seeking the best way to play the prisoner's dilemma, those using the new selfish strategy will eventually be bested by more generous players.

The roots of cooperation

The prisoner's dilemma is a simple 'game' that captures the fundamental problem faced by a population of organisms competing for limited resources: the temptation to cheat or freeload. You might do better acting together than working alone, but the temptation is to take a share of the spoils while letting others put in the effort and face any risks.

The simplest version of the game pits a pair players against each other. The players obtain particular pay-offs if they elect to cooperate or 'defect' (act selfishly). In a single bout it always makes sense to defect: that way you're better off whatever your opponent does. But if the game is played again and again — if you have repeated opportunities to cheat on the other player — you both do better to cooperate.

This 'iterated' prisoner's dilemma has been used to show how cooperation could arise in selfish populations: those who are genetically disposed to cooperate will be more successful than those who are predisposed to defect.

In studies during the 1980s, it seemed that the most successful strategy in the iterated game was a simple one known as Tit-for-Tat (TfT), which merely copies the opponent's behaviour from the last round: defection is met with defection, cooperation with cooperation. The moral message seemed reassuring: it pays to be nice and not be the first to defect, but nastiness should be punished.

However, in further studies it became clear that TfT might not always dominate in evolutionary games in which the most successful strategies are propagated from generation to generation. Slightly more forgiving strategies, which don't get caught in cycles of mutual recrimination by a single mistaken defection, can do better in the long run. In fact, there is no single best way to play the game — it depends on your opponents. Nonetheless, the iterated prisoner's dilemma seemed to explain how cooperation between unrelated individuals might evolve: why some animals hunt in packs and why we have altruistic instincts.

Unfair division

Press and Dyson seemed to shatter this cosy picture. They showed that there exists a class of strategies, which for technical reasons they call zero-determinant (ZD) strategies, in which one player can force the other to accept a less-than-equal share of the pay-off. The victim must either grit his teeth and accept this unfair division, or punish the other player at a greater cost to himself.



Image Source/Rex Features

In game theory, the 'prisoner's dilemma' is used to test strategies of cooperation versus selfishness.

Like a TFT player, a ZD player bases his next choice of cooperate/defect on what happened in the last round. But instead of being rigidly deterministic — so that the previous outcome dictates this choice absolutely — it is probabilistic: the choice to cooperate or defect is made with a certain probability for each of the four possible outcomes of the last round. A judicious choice of these probabilities enables one player to control the pay-off that the other receives.

The ZD strategy seemed to overturn several decades of consensus about the prisoner's dilemma and the evolution of cooperation.

"The paper caused quite a stir, because the main result appeared to be completely new, despite intense research in this area for the past 30 years", says Adami.

It wasn't totally new, however. In 1997, game theorists Karl Sigmund of the University of Vienna, Martin Nowak of Harvard University in Cambridge, Massachusetts, and Maarten Boerlijst of the University of Amsterdam discovered strategies that similarly allow one player to fix the other's pay-off at a specified level[3]. But they admit that "we didn't know about the vast and fascinating realm of zero-determinant strategies."

Switching strategies

Now, Adami and Hintze say that the extortionate world of the ZD player might just be transient. They find that, in an evolutionary iterated prisoner's dilemma game in which the prevalence of particular strategies depends on their success, ZD players are soon out-competed by others using more common strategies, and so they will evolve to become non-ZD players themselves. That's because ZD players suffer from the same problem as habitual defectors: they do badly against their own kind.

There is one exception: ZD players can persist if they can work out whether they are playing another ZD player or not. Then they can exploit the advantages of ZD strategies against non-ZD players, but will switch to a more advantageous non-ZD strategy when faced with their own kind.

Sigmund and Nowak, with their colleague Christian Hilbe of the Max-Planck Institute for Evolutionary Biology in Plön, Germany, have also shown in work not yet published that the ZD strategy is evolutionarily unstable, but can pave the way for the emergence of cooperators from a more selfish community.

Game theorist and economist Samuel Bowles at the Santa Fe Institute in New Mexico feels that these results demote the interest of the ZD strategies. "The question of their evolutionary stability is critical, and the paper makes their limitations clear. Because they are not evolutionarily stable, I'd call them merely a curiosity of little interest to evolutionary biology or any of the other biological sciences."

But Adami is not so sure that the same ZD strategies won't be found in the wild. "We don't usually have nearly enough information about either animal or cellular decisions — microbes play these kinds of games too," he says. "But in my experience, anything that is imaginable has probably evolved somewhere, sometime. To gather conclusive evidence about it is another matter."

Nature | doi:10.1038/nature.2012.11254

References

1. Press, W. H. & Dyson, F. J. *Proc. Natl Acad. Sci. USA* **109**, 10409–10413 (2012).
2. Adami, C. & Hintze, A. Preprint at <http://www.arxiv.org/abs/1208.2666> (2012).
3. Boerlijst, M. C., Nowak, M. A. & Sigmund, K. *Am. Math. Monthly* **104**, 303–305 (1997).