



© 1997–2009, Millennium Mathematics Project, University of Cambridge.

Permission is granted to print and copy this page on paper for non-commercial use. For other uses, including electronic redistribution, please contact us.

12/11/2008

News

Trust me, I've evolved



Why are some people generous and others selfish? There's no doubt that both strategies pay off under certain circumstances, but research (as well as everyday experience) shows that we are not mere opportunists some people simply *are* nicer than others. This raises a question which intrigues evolutionary psychologists: is there a selective force that works in favour of a wide range of personalities, preventing us from all evolving the same optimal character trait? A possible answer has recently been published by mathematicians from the universities of Bristol and Exeter. Their game theoretical model of human interaction shows that diversity, once it has taken root in a population, can be self-perpetuating and ensure its own existence for generations to come.



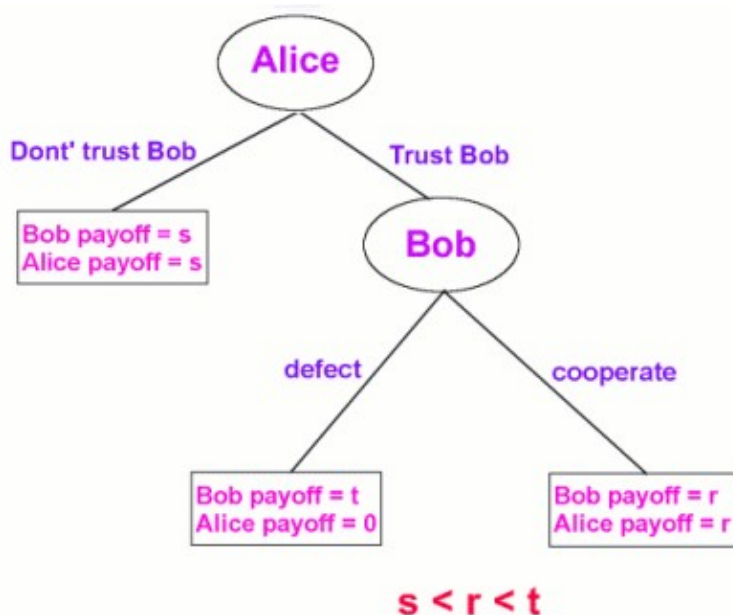
Who do you trust?

"What we found is that watching each other's behaviour produced individuals who were more socially aware," explains John McNamara, who led the research, "which in turn exaggerated the personality traits of both players. Some became more cooperative because they became aware of the impact their decisions were making on their reputations while others became less cooperative and exploited trusting individuals for personal gain."

Trust me, I've evolved

McNamara's team, which also includes Philip Stephens and Sasha Dall, work in *evolutionary game theory*, an area that over the last few decades has gained popularity with people who study evolving systems, including economists, biologists, sociologists and psychologists. As the name suggests, the theory likens inherited traits to strategies in the great game of evolution. An advantageous trait has a positive payoff in terms of fitness to survive and have offspring, and this is reflected in a larger proportion of the next generation having that particular trait. Using simple mathematical rules to describe the payoff and replication of traits in the next generation, researchers can sit back and let a computer work out whether a particular trait survives, becomes dominant, or eventually dies out.

McNamara, Stephens and Dall based their model on a simple game of trust and cooperation. Imagine that two individuals, call them Alice and Bob, have been thrown together and have to decide whether or not to cooperate on a certain task. Alice is the one who has to make the first move. She has to signal whether or not she is willing to trust Bob. If she does, and Bob reciprocates, then both receive a certain payoff, which is higher than the payoff they would both get if Alice had decided not to trust Bob. If, however, Bob decides to defect once Alice has decided to trust him, he will receive a higher payoff still, while Alice will get nothing.



As it stands, only a very naive Alice would decide to trust Bob, since he doesn't have much incentive to cooperate. In terms of evolution, you would expect naive Alices to eventually die out, which would in turn make altruistic, cooperating Bobs redundant. But now imagine that social awareness enters the game. If Bob is aware that he is being watched by future Alices, he may well decide to cooperate if trusted, for if he doesn't, he may never be trusted again and therefore always lose out on the higher reward. If Alice has some way of finding out how Bob has behaved on previous occasions, she may be more willing to trust. Being a trusting Alice and a cooperating Bob may become viable evolutionary strategies. The question is whether Alices with varying degrees of trust and Bobs with varying degrees of trustworthiness can exist stably alongside each other.

To find out, the researchers imagined a population of individuals who are randomly paired off in Alice–Bob situations. Each individual plays the role of Alice in some interactions and the role of Bob in others. When cast in the role of Alice, an individual always behaves in the same way, according to an inherited trait: some always trust Bob, some never trust Bob, some trust Bob if he has been seen to cooperate in at least one of n previous occasions, some trust him if he has been seen to cooperate in at least two of n previous occasions, and so on. The latter classes represent Alices with social awareness and varying degrees of trust. When cast in the role of Bob, each individual may or may not cooperate if trusted, and the individual's chance of

Trust me, I've evolved

cooperating is measured by a number p between zero and one. The continuously varying values of p represent Bobs of varying degrees of trustworthiness. The researchers assumed that an individual's degree of trust, when being Alice, and trustworthiness, when being Bob, are independent. This isn't entirely realistic, but it frees the researchers from having to keep track of the association between genes that control the traits.

Each interaction comes with a certain payoff for the two individuals involved. The exact value of this payoff depends on the individuals, but using statistics it's possible to calculate a *mean payoff*, a sort of average, for the traits in each role. For example, when acting as Alice an individual who uses some social awareness to assess Bob may have a higher mean payoff than one who trusts Bob unquestioningly. Similarly, an individual who in the role of Bob has a 50% chance of cooperating may have a higher mean payoff than one whose chance of cooperation is zero. These mean values are used to calculate the proportion of the next generation that display a certain trait: the researchers devised a formula which ensures that the higher the expected payoff for a trait (such as the degree of trust in an Alice, or the degree of trustworthiness in a Bob), the higher the proportion of the next generation with that same trait.

With these rules in place, it would be possible to calculate how the various degrees of trust and trustworthiness perform over many generations, but there is another important issue to take into account: in real-life evolution individuals don't just acquire traits through inheritance, but also through mutation. This is important because a society of nice altruists could be turned on its head within a few generations by the influence of a selfish mutant. The researchers built mutation into their model by introducing mutation rates which slightly adjust the replicating formulae for the next generation.

The next step was to set the computers chomping. Starting with some initial distribution of traits and values for the various parameters that define the game the payoffs that result from trusting and cooperating, and the rates of mutation the researchers waited to see how the traits fared generation after generation. For certain initial values, the proportions of individuals with the various traits eventually settled down they hardly changed at all from generation to generation. Such an unchanging configuration of traits was deemed to be *evolutionary stable*.



Social awareness makes for nice people.

Unsurprisingly, when the researchers initialised their model without social awareness with only universally trusting or untrusting Alices the number of trusting Alices became negligible, while the most popular chance p of cooperation was zero. Without social awareness, it is impossible to maintain trustworthiness and trust. However, in stable situations that involved social awareness, two types of Bob could be seen to exist

Trust me, I've evolved

alongside each other: a non-cooperating type, and a type with quite a high degree of trustworthiness. Universally trusting Alices were also seen to exist alongside universally untrusting ones, and those who display social awareness and varying degrees of trust a whole range of character traits existing in an evolutionary stable situation.

The researchers put the phenomenon down to a positive feedback loop: "Variation encourages social awareness," explains Dall, "which favours greater personality divergence, maintaining the need for social awareness. In other words, because people operate in such different ways, we need information to decide whether or not trust to them. This encourages a really diverse range of responses which, in turn, makes social awareness all the more important."

Real-life people and their interactions are of course a whole lot more complicated than the virtual individuals in the model. But what the research shows is that diversity is no mystery: it can arise from the very rules of the game and maintain itself. The research also shows that models of human behaviour, used widely in economics, can no longer ignore complex and changing interpersonal relationships, as they have tended to do so far.

[Post a comment on this story](#)

Marianne Freiberger



Plus is part of the family of activities in the Millennium Mathematics Project, which also includes the [NRICH](#) and [MOTIVATE](#) sites.