Primer

Evolutionary game theory

Karl Sigmund* and
Martin A. Nowak†

Ever since Darwin read Malthus, the theory of evolution has benefited from the interaction of ecology with economics. Evolutionary game theory belongs to this tradition: it merges population ecology with game theory. Game theory originally addressed problems confronted by decision makers with diverging interests (for instance, firms competing for a market). The 'players' have to choose between strategies whose payoff depends on their rivals' strategies. This interdependence leads to a mutual 'outguessing', as with chess (she thinks that I think that she thinks…). There usually is no solution that is unconditionally optimal, no matter what the co-players are doing.

In the context of evolutionary biology, the two basic notions of game theory, namely strategy and payoff, have to be re-interpreted. A strategy is not a deliberate course of action, but an inheritable trait; payoff is Darwinian fitness (average reproductive success). The 'players' are members of a population, all competing for a larger share of descendants.

If several variants of a trait occur in a population, then natural selection leads to an increase in the frequency of those variants with higher fitness. If the success of a trait does not depend on its frequency, this will eventually lead to the fixation of the optimal variant. But if the success of a trait is frequency-dependent, then its increase may lead to a composition of the population where other variants do better; this can be analysed by means of game theory. This is similar to what happens in population ecology. If prey is abundant, predators will increase for a while. But this increase reduces the abundance of prey, and eventually leads to a decrease of the predators.

Hawks and doves

Intraspecific fights provide a first example of changes in a population that are dependent on the frequency of a trait. Assume that there are two behaviourally distinct morphs: 'hawks' escalate a fight until the injury of one contestant settles the outcome, 'doves' stick to some conventional display (a pushing match, for instance, where injuries are practically excluded) and take flight if the adversary escalates. If most contestants are doves, the hawk morph will spread; but if most contestants are hawks, escalating a conflict will lead with probability ½ to injury (see green box and Figure 1). Even this oversimplified example shows that heavily armed species (where the risk of injury is large) are particularly prone to conventional displays (see Figure 2).

A wide variety of behavioural traits — and purely morphological or physiological characters, like the height of trees or the length of antlers — are submitted to frequency-dependent selection and are amenable to game analysis. Such traits may influence conflicts of interest between two individuals, for instance, territorial disputes (between neighbours), the length of the weaning period (between parents and offspring), or the division of parental investment (between male and female). But frequency-dependent selection also occurs without antagonistic encounters. The sex ratio is an example of this (if it is biased towards males, it pays to produce daughters, and vice versa).

The prisoner's dilemma

The evolution of cooperation through reciprocation is a particularly extensive chapter of evolutionary game theory. Why should a selfish gene, or 'fitness maximiser', bear the cost of helping an unrelated individual? Suppose that the benefit (to the recipient) is b and the cost (to the donor) is c, with $c < b$. If an individual is equally likely to be potential recipient or donor in a given encounter, then a population of cooperators would earn, on average, $b - c$ per interaction, and be better off than a population of defectors earning 0. But no matter how the population is composed, an individual would always increase its

---

The hawk–dove game

In the hawk–dove game, G is the gain in fitness resulting from winning the contest, and C is the cost in fitness due to an injury (see Figure 1).

Figure 1

		Type of adversary	
		Hawk	Dove
Contestant who wins payoff	Hawk	$(G\text{-}C)/2$	G
	Dove	0	$G/2$

Current Biology

The payoff matrix for the hawk–dove game.

No morph is unconditionally better than the other. If the object of the fight is not worth the injury, then the dove morph will spread. Hawks can spread only if their frequency is below G/C. If G is less than C, a mixed population of hawks and doves will evolve. It is conceivable that some phenotype plays a mixed strategy, escalating only with a certain frequency. Such a mixed strategy escalating with probability G/C is evolutionarily stable; no mutant with a different propensity to escalate can invade. If there exists an asymmetry between the two contestants (larger versus smaller, for instance, or owner versus intruder), this will alter the game. In such cases a conditional strategy will emerge which uses the asymmetry as a cue, for instance the so-called 'bourgeois strategy': if owner, escalate; if intruder, display. Some intraspecific conflicts use a long assessment phase to detect an asymmetry in, for instance, size, by means of a pushing match or a parallel walk.

fitness by refusing help, and hence we would not see cooperation.

Game theorists have encapsulated this tug-of-war between common good and selfish interest in the so-called prisoner's dilemma game. In this game, a player has two possible strategies **C** (to cooperate) and **D** (to defect). Two **C** players will get a reward R which is higher than the punishment P obtained by two **D** players. But a **D** player exploiting a **C** player obtains a payoff T (temptation to defect) which is higher than R, and this leaves the **C** player with the sucker's payoff S, which is lower than P. So, because $T > R > P > S$, a rational player will always play **D**, which is the better move no matter what the co-player is doing.

Many species engage in interactions which seem to be of the prisoner's dilemma type. Vampire bats feed each other, monkeys engage in allogrooming, birds and vervet monkeys utter alarm calls, guppies and stickleback cooperate in predator inspection, hermaphroditic sea bass alternate as egg-spenders and lions engage in cooperative hunting or joint territorial defense.

But often, attempts to specify the payoff values of these behavioural interactions lead to doubts as to whether one is seeing a *bona fide* prisoner's dilemma. It is difficult to measure the fitness of fish darting in and out of shoals or of monkeys hiding in the bush. It may be that some of these are simply instances of by-product mutualism, in which both players are best served by cooperating and none is tempted to defect, or an instance of the so-called snowdrift game, in which the best reply to the co-player's **C** is a **D**, but the best reply to a **D** is a **C**. For instance, if two drivers are caught with their cars in a snowdrift and one defects by not shovelling, the other is still better off by digging, rather than defecting.

Recently, however, a striking example of a prisoner's dilemma type of interaction has been uncovered for an RNA phage — a virus that

Figure 2

A clash between male red deer. Only rarely does the fight escalate beyond a pushing match. (Photograph provided by Tim Clutton-Brock, Department of Zoology, University of Cambridge, UK.)

reproduces inside a bacterium. In this instance, one phage variant produces less of the intracellular products needed for replication than the other, and thus may be said to defect. The payoff values can be measured with some precision: they satisfy the rank ordering required for the prisoner's dilemma.

There are several ways in which the prisoner's dilemma can be overcome. With the phages, selection for a particular trait-group operates whenever the virus population is so small that most bacterial hosts are invaded by only one virus. A phage will then most likely interact only with members of its own clone, and this is when cooperators are better off than defectors. More generally, any form of associative interaction favours cooperation. Such association can be due to kinship, to partner choice, to ostracism of defectors or simply to spatial structure (or 'population viscosity').

Repeated interactions

Among higher organisms, interactions of the prisoner's dilemma type are probably repeated between the same two players — called the iterated prisoner's dilemma. Depending on their past experience, players can break up partnerships, or vary their amount of cooperation. But even without these options, the strategy of always defecting is not invariably the best option in the iterated prisoner's dilemma. If the probability of a

further round is sufficiently high, then the presence of even a small number of so-called 'retaliators' is enough to favour cooperation. The best known example of such a retaliatory strategy is 'tit-for-tat'. A tit-for-tat player cooperates in the first round and from then on always repeats whatever the co-player did in the previous round.

After promoting the emergence of cooperators, retaliatory strategies are often superseded by more tolerant strategies, for instance 'generous tit-for-tat' or 'win–stay, lose–shift'. A generous tit-for-tat player cooperates after an opponent's defection with a certain probability, whereas a win–stay, lose–shift player repeats the previous move after receiving a high payoff in a round, but otherwise switches to the other move. Both generous tit-for-tat and win–stay, lose–shift players return to cooperation after an erroneous defection, whereas tit-for-tat players do not. A great many theoretical results and computer simulations show that under very general conditions, populations of defectors can be invaded by small groups of stern retaliators, who pave the way for more tolerant populations, which, in turn, can eventually 'soften up' to such a degree that defectors may take over again.

The best examples of reciprocal strategies may be found in human societies. But here, reciprocation is often indirect. An act of assistance is returned, not by the recipient, but by a third party. Such indirect reciprocation can be based on score-keeping by 'discriminate altruists' (who help only those individuals that have not refused help too often). Withholding help can be costly, as it reduces the discriminators' chance of being helped. Simple models show that indiscriminate altruists can therefore spread at the expense of discriminate altruists; once their frequency exceeds a certain threshold, defectors can invade. Defectors will reduce the proportion of

indiscriminate altruists, however, and therefore sap the basis of their own success. Discriminate altruists will then be able to eliminate them again.

Other models based on evolutionary games have been used to analyse the emergence of moral systems, of division of labour or even of proto-languages in primitive societies. In particular, it has been shown that highly diverse social norms can be stably sustained by punishment; in many cases, deviating behaviour (for instance, the refusal to punish dissenters) cannot spread. In this sense, evolutionary game theory returns to some basic topics of classical game theory.

Population dynamics

The latest advances in evolutionary game theory are being made by means of population dynamics. This transfer from population ecology relies on the assumption that successful traits spread. If there are only two possible morphs A and B, for instance, then essentially only three scenarios are possible. In scenario 1, A can invade B and B can invade A. This leads to the coexistence of both types in stable proportions as, for instance, if A are hawks and B are doves. In scenario 2, A can invade B but B cannot invade A. In this case, the dominant strategy A will always outcompete B. This happens with the prisoner's dilemma, if A players are the defectors and B players cooperate. In scenario 3, no type can invade the other. This is a bistable situation; whoever exceeds a certain (possibly quite small) threshold will outcompete the other. This happens with the iterated prisoner's dilemma if A is a tit-for-tat player and B always defects.

As soon as three morphs A, B and C compete, the situation becomes considerably more complex, because cycles can occur: A dominates B, B dominates C, and C dominates A. This apparently bizarre situation holds for the lizard *Uta stantibus*, where three morphs (with different throat colors) correspond to

inheritable male mating strategies. Type A forms no lasting bonds but looks for sneaky matings; type B lives monogamously and closely guards the female; and type C guards a harem of several females, of course less closely. With still more strategies, the dynamics can become ever more complex, and the composition of the population need not converge to an equilibrium: the frequencies of the strategies keep oscillating in a regular or chaotic fashion.

Long-term evolution

We have so far discussed the evolution of frequencies of traits under selection alone. But, in addition, mutation will introduce new variants from time to time. This proceeds at another time scale. Game theory can be used to analyse both short-term and long-term evolution. In order to deal with the latter, it is often advantageous to assume that the transient effects following a random mutation have settled down before the next mutation occurs. As long as the population remains monomorphic, this leads to a trait substitution sequence: selection settles the fate of a mutant (fixation or elimination) before the next mutation comes along.

Along with the two time scales come two notions of stability: internal stability against perturbations in the frequencies within the resident population, and external stability against invasion attempts by any mutant morph. For a monomorphic population, the latter notion corresponds to an evolutionary stable strategy: no conceivable alternative strategy can invade. But adaptation need not converge to an evolutionarily stable strategy: the outcome can be a so-called 'Red Queen' type of dynamics, leading for instance to potentially endless 'arms races', to 'runaway selection' or to cyclic chases in trait space.

Evolutionary game theory deals with phenotypes. But with sexual replication, it cannot be assumed that

'like begets like', so that more successful types can also produce less successful variants and vice versa. The complications caused by Mendelian segregation, pleiotropy and sexual recombination can, in principle, be integrated to model frequency-dependent selection acting within the gene pool, but this often results in intractable dynamics.

But arguments from evolutionary game theory can fail if genetic constraints come into play. (If the optimal genotype can only be realized by heterozygotes, for instance, then it can never reach fixation in the population.) Theoretical studies tend to show, however, that the genetic system is flexible enough to overcome such constraints. In the absence of specific information on the genotype–phenotype map, evolutionary game theory often provides a reasonably efficient tool for understanding adaptation without getting bogged down by complications at the genotype level. Here again, evolutionary game dynamics follow the traces of Darwin — who, it will be noted, read Malthus, not Mendel.

Key references
Axelrod R: *The Complexity of Cooperation*. Princeton: Princeton University Press; 1998.
Dugatkin LA: *Cooperation Among Animals: an Evolutionary Perspective*. Oxford: Oxford University Press; 1997.
Hofbauer J, Sigmund K: *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press; 1998.
Maynard Smith J: *Evolution and the Theory of Games*. Cambridge: Cambridge University Press; 1992.
Milinski M: **Tit-for-tat in sticklebacks and the evolution of cooperation.** *Nature* 1987, **325:**434-435.
Nowak MA, May RM: **Evolutionary games and spatial chaos.** *Nature* 1992, **359:**826-829.
Nowak MA, Sigmund K: **Evolution of indirect reciprocity by image scoring.** *Nature* 1998, **393:**573-577.
Sigmund K: *Games of Life.* London: Penguin Books; 1995.
Trivers R: **The evolution of reciprocal altruism.** *Quarterly Rev Biol* 1971, **46:**35-57.
Turner PE, Chao L: **Prisoner's dilemma in an RNA virus.** *Nature* 1999, **398:**441-443.

Addresses: *Institute of Mathematics, University of Vienna, and IIASA, Laxenburg, Austria. †Institute for Advanced Study, Princeton, New Jersey, USA.