

16

Population Dynamics In Evolutionary Ecology

MARTIN A. NOWAK and KARL SIGMUND

16.1

Introduction

Ernst Mayr [1] argues that the 'greatest conceptual revolution that has taken place in biology' is the replacement of typological thinking by population thinking. In some way, the tremendous progress in molecular biology during the last decades has threatened to obscure this fact. It is certainly essential to understand the chemical mechanisms going on between specific molecules – how viruses, for example, use binding proteins to attack and penetrate hosts cells. But this does not suffice to tackle basic problems like, for instance, disease progression or the co-evolution of hosts and parasites. It is *populations* of virus particles, or immune cells, or hosts, or genes, that regulate each other's frequencies.

Although he had no mathematical training, Darwin was aware of the intricate dynamics of interacting populations. 'Throw up a handful of feathers', he wrote in *The Origins of Species*, 'and all must fall to the ground according to definite laws. But how simple is this problem compared to the action and reaction of the innumerable plants and animals.' And more explicitly, he describes how, 'if certain insectivorous birds were to increase in Paraguay, a species of flies would decrease', how this decrease would cause cattle to become abundant, because the flies parasitise newborn calves, and how this in turn 'would certainly greatly alter the vegetation,..., and this again would largely affect the insects; and this again the insectivorous birds ... and so onwards in ever-increasing circles of complexity.' As another example of 'how plants and animals, most remote in the scale of nature, are bound together by a web of complex relations', Darwin wrote: 'It is quite credible that the presence of a feline animal in large numbers might determine, through the intervention first of mice and then of bees, the frequency of certain flowers!' Such feedback loops are too complex to be understood by verbal arguments alone. The biological community has come to accept that basic aspects of evolutionary ecology, and related feedback mechanisms in immunology and epidemiology, can only be analysed by mathematical means.

The corresponding non-linear control mechanisms are gold mines for mathematical modellers. These mines have been exploited throughout the past century, and given rise to many important mathematical techniques and concepts (for instance

chaotic dynamics, branching processes or travelling waves, to name but a few). The mathematical work of Fisher, Haldane, Wright and Kimura has shaped genetics, that of Volterra and May ecology, and that of Kermack and McKendrick epidemiology. During the last decades, mathematical models have blossomed in these fields, and are increasingly used to shape experimental research in virology, epidemiology, conservation biology or animal behaviour. The degree of mathematisation is unlikely to reach that in physics, but mathematical concepts like evolutionary stability, replicator dynamics, or basic reproductive ratios have become essential tools in biological discussions.

We propose to sketch here a few recent developments. Our aim is to emphasise the basic similarity of frequency-dependent selection across all scales of nature, from RNA viruses to human societies. In order to highlight the interplay of demographic variables, we purposely neglect other factors which often affect the dynamics, like time-delays, stochastic fluctuations or spatial distributions (see e.g. [2] or [3]). In this survey, we sketch

- (1) the mathematical framework based on ordinary differential equations;
- (2) some applications in mathematical ecology, including in particular virus dynamics and epidemiology;
- (3) some concepts of evolutionary game dynamics and the dynamics of adaptation, and
- (4) their impact on the investigation of the emergence of cooperative societies.

16.2

The Framework: Population Dynamics

Let $x_i(t)$ denote the density of population i at time t and let \dot{x}_i be its time derivative. The per capita increase to population growth, i.e. \dot{x}_i/x_i , will in general depend on the frequencies of some or all of the other species in the ecosystem. This yields the ecological equation

$$\dot{x}_i = x_i f_i(x_1, \dots, x_n) \quad (1)$$

on the positive orthant R_+^n . This state space is invariant, and so are its boundary faces, where some of the population densities vanish. In the scenario considered here, if a population is absent, it remains so: we neglect the possibility that migration, or mutation, can introduce it into the ecosystem. Depending on the biological interaction, the growth rate f_i can take many forms. Often, for instance, a population is self-limiting, i.e. f_i is decreasing in x_i (in other situations, we may have an Allee effect, where f_i is an increasing function of the density x_i as long as it is small (more individuals would mean more chances to find a mate)). Whether f_i grows or decreases as a function of density x_j (with $j \neq i$) depends on whether population j competes for the same resources, or predaes on i , etc. The very simplest case, when the f_i are linear function, leads to the class of Lotka-Volterra equations

$$\dot{x}_i = x_i \left(r_i + \sum_{j=1}^n a_{ij} x_j \right) \quad i = 1, \dots, n \quad (2)$$

which covers a wide spectrum of possible types of interaction and displays many possible dynamical behaviours (for a survey, see [4] or [5]). In the case of a prey-predator interaction, for instance, the trajectories either converge to the boundary (if the prey is unable to support the predator population) or else they spiral in damped oscillations towards an equilibrium with both populations coexisting. In the case of two competing species (the r_i are positive and the a_{ij} all negative) three outcomes are generically possible: *dominance* of one species which, for every initial condition, eliminates the other species; *coexistence* of both species in a stable equilibrium; or *bistability* (one species eliminates the other, but the outcome depends on the initial condition).

Two-dimensional Lotka-Volterra equations are easy to classify, but for three dimensions, the behaviour is much more complex. If two predator species depend on one common prey species, chaotic dynamics are possible: irregular oscillations which do not damp down, and display a very sensitive dependence on the initial condition. With three competing species, chaotic behaviour can be ruled out, but nevertheless, a complete classification has not been achieved yet. One of the reasons for this is the existence of so-called *heteroclinic attractors*. It can happen that in the absence of the third species, species 1 is dominated by 2; similarly, 2 is dominated by 3, and 3 by 1 (see [6]). Such *rock-scissors-paper cycles* are not a mathematical artefact: they have been observed in lab experiments on microorganisms (see [7]), and are likely to play a role in real ecosystems.

Indeed, the heteroclinic cycles and networks become more and more likely if the number of interacting species increases. In their simplest form, such networks consist of finitely many equilibria which are all saddle-points, and of orbits connecting these saddle points in the sense that they lead, for $t \rightarrow -\infty$, to one equilibrium and for $t \rightarrow +\infty$ to another. The corresponding network is supposed to be transitive (which means that one can journey along such saddle connections from any equilibrium to any other one). It can be shown that such networks can occur as limit sets of orbits: there are trajectories which, for arbitrary large time, visit arbitrarily small neighborhoods of every point of such a network (and of no other point). The behaviour of such trajectories is quite striking: they hover for some time close to one equilibrium point, then switch relatively fast to another equilibrium, where they linger for a much longer time, then switch brusquely to the next equilibrium etc. This is not a periodic oscillation: the times spent near saddle points increase exponentially.

In the case of three competing species dominating each other in a rock-scissors-paper cycle, a trajectory attracted by such a heteroclinic cycle will spend more and more time closer and closer to the boundary. A small stochastic fluctuation will eventually wipe out whichever species happens to be here in a marginal quantity only. Of the two remaining species, one will be eliminated by the other, and hence we end up with one species only. But such an one-species equilibrium is unstable: if the 'right' missing species is introduced in a tiny minority, it will be able to invade and even-

tually take over; but the resulting one-species equilibrium is again prone to invasion by the next species, etc. It is quite likely that many of the sudden transitions and recurrent instabilities which we are witnessing in real ecosystems are based on such heteroclinic networks.

As we have seen, a heteroclinic cycle on the boundary of the state space R_+^n (where some x_i are 0) which attracts orbits from the interior of the state space (where all species are present) spells the doom for some species and hence implies a type of ecological instability. In contrast, we shall say that an ecosystem modelled by (1) is ecologically stable, or more precisely *permanent*, if the interior of R_+^n contains a compact set K which attracts all orbits in the interior. Whenever the system starts from a state with all populations present, then after some transition period, the trajectory will be close to K and hence cushioned away from the boundary: if fluctuations are sufficiently small and rare, they will not be able to eliminate some population. A system is permanent, therefore, if the boundary of the state space is a repeller. This new, ecologically motivated notion of stability has proved quite successful (see [4] and [8]). Especially in the case of Lotka-Volterra equations (2), there exist quite a few useful necessary conditions for permanence, and also some sufficient conditions (but so far, no useful necessary and sufficient condition).

Permanence concerns the extinction of some populations in the ecosystem. The other side of the coin is the invasion of a resident community by a small minority of a newly introduced population $n + 1$. If the resident community is at an equilibrium (z_1, \dots, z_n) , the success of such an invasion attempt will depend on the sign of the growth rate $f_{n+1}(z_1, \dots, z_n, 0)$, i.e. the eigenvalue at the equilibrium $(z_1, \dots, z_n, 0)$ which is transversal to the boundary face $x_{n+1} = 0$. If it is positive, the missing species can invade. This in itself, of course, says nothing about the further development. If population $n + 1$ invades, it can simply join the other populations, so that the number of resident populations increases, or it can eliminate some of the residents, or it can even, after an initial growth phase, eliminate itself! The problem becomes much more complex if the attractor of the resident system is not a point, or a periodic orbit, but a strange attractor. Obviously, in such a case the growth rate of population $n + 1$ has to be averaged, but with respect to which measure? This leads to interesting ergodic problems concerning intermittency, riddled basins etc. (see [9], [10], [11]).

16.3

Applications in Ecology, Epidemiology and Immunology

An intriguing aspect of classical population ecology concerns the top predator – a species at the top of the food chain, with no predator preying on it. Top predator populations are usually small, and vulnerable to extinction. What happens if a top predator is removed, either through natural causes or through human intervention? One would tend to assume that the elimination of such a supreme exploiter improves the conditions of the other species in the ecosystem. In fact, ecologists have repeatedly found that this is not the case. The diversity of the remaining ecosystem drops drastically, often to less than half the former number of species.

Mathematical models show that this should not come as a surprise. The fact that a permanent system admits only considerably smaller permanent subsystems is the rule rather than the exception. This has interesting applications to the field of community construction (see [12], [13], [14]). Let us consider a given (often fictitious) 'species pool' and construct an artificial ecosystem by drawing, at random, some members of that pool and introducing them, in small number, into some habitat. This experiment can be viewed as a sequence of invasion attempts. Some attempts will succeed, some will fail. The diversity of the resulting ecosystem can grow or decrease. It often happens, in such experiments, that at some stage no further invasion attempt can succeed. It can also happen that the composition of the ecosystem cycles indefinitely. Many of the resulting ecosystems cannot be constructed by simply adding one resident after another. It often happens, for instance, that the permanent n population system admits no permanent $(n - 1)$ -subsystem. The n -species system has to be constructed in a more roundabout way involving some extra species, in a way which resembles the construction of an arch: the extra species play the role of a scaffolding which is later removed. The current emphasis of population ecology on the contingency and history-dependence of evolving ecosystems fits well with this modelling approach.

A particularly important chapter of population dynamics concerns the spread of infectious diseases (see e.g. [15]). The simplest model in the monumental book by Anderson and May [16] describes the frequencies x and y of uninfected and infected individuals in one (well-mixed) population by the equation:

$$\begin{aligned}\dot{x} &= k - dx + cy - \beta xy \\ \dot{y} &= \beta xy - dy - vy - cy.\end{aligned}\tag{3}$$

In the absence of the infection ($y = 0$) the frequency x converges to the uninfected equilibrium k/d . But if infected individuals are present, then random contacts transform uninfected into infected (this corresponds to the term βxy). Conversely, infected can recover (cy) and become uninfected again (for the sake of simplicity we neglect to consider the effects of acquired immunity in this model). The term v describes the additional mortality due to the infectious disease – the so-called virulence. One sees immediately that the number of uninfected, x , can only grow if the number of infected y exceeds the threshold $\beta/(d + v + c)$. Hence the pathogen carrying the disease (usually a parasite, for instance a virus or a bacterium) can only invade if this threshold is smaller than k/d . This means that the so-called reproductive ratio

$$R_0 := \frac{k}{d} \frac{\beta}{d + v + c}\tag{4}$$

has to be larger than 1. R_0 is the average number of infections caused by one infected individual in an otherwise uninfected population. As soon as this model is extended to cover more realistic situations (by taking account of spatial or temporal effects, the multiplicity of risk groups, immune reactions etc) it becomes considerably more complex (see [17] and [18]).

Models of this type describe the population ecology of parasites and their host species. They are obviously closely related to predator-prey models: the more parasites, the worse for the host, the more host, the better for the parasite. Hence we may expect a similar dynamical behaviour. Thus parasites can, in their role as keystone species, mediate the permanence of an ecosystem. Many biologists see in microbes a basic cause for the biodiversity prevailing in nature. The models lead to damped or chaotic oscillations, and to heteroclinic cycles. In the simplest case, this happens when two strains of the host population live in a bistable equilibrium, each one attacked by its own specialised strain of parasite. In a population of host species 1, the host species 2 cannot take hold. But the parasite strain 1 can invade, and weaken the host population 1 to such an extent that host 2 can invade and, in fact, eliminate host 1 (together, of course, with its parasite 1). But now, parasites of type 2 can enter, etc.

So far, these models describe the evolution of the parasite population at large, and its interaction with the entire host population. But an important new chapter of population ecology describes the dynamics of parasite populations within *one* host organism. This has been first studied for HIV (for a survey of such models, see [19]). It is well known that a long latency period with a median value of ten years or more separates infection from the outbreak of the full-blown AIDS-symptoms. For a long time, this seemed to reflect a strange period of apparent quietness in the behaviour of the virus. But clinical research based on the applications of methods of mathematical demography has shown that during this so-called latency, the body of the infected host contains usually more than 10^{10} free virus particles, and that its immune system is engaged in a tremendous battle. On average, free virus can only survive for a few hours, and infected cells for one or two days. The fast dynamics and the huge population numbers offer excellent conditions for modelling the immune reaction by methods developed in mathematical ecology.

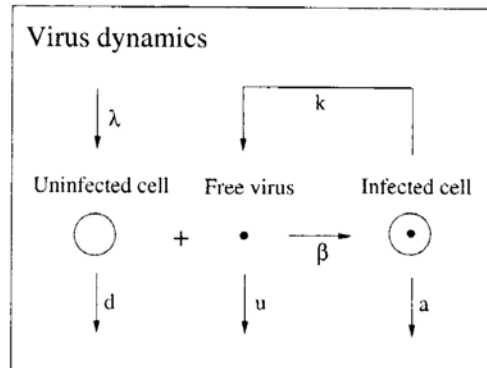
HIV plays a double role in this context (see [20]). On the one hand, the virus is obviously a predator exploiting certain cells of the infected body – the blood cells, the lymph cells, so-called target cells. On the other hand HIV is the prey of certain highly specialised cells of the immune system, whose mass production is stimulated by the infection – so-called killer cells (B-cells, CTH4 cells etc). Which role is more important for the viral population dynamics, that of predator or that of prey? Measurements have shown that they are of comparable importance [20]. Hence models have to include both effects. But as is well-known from mathematical ecology, the dynamics of prey-predator-superpredator models can be quite involved.

A basic model (see [19]) is the following:

$$\begin{aligned}\dot{x} &= k - dx - \beta xw \\ \dot{y} &= \beta xw - ay - \gamma z \\ \dot{w} &= hy - dw \\ \dot{z} &= \gamma z - bz.\end{aligned}\tag{5}$$

Here x is the frequency of uninfected target cells, and y that of infected target cells. w is the frequency of free virus and z that of killer cells. These killer cells destroy infected cells, and thereby increase their own reproduction (this corresponds to

Fig. 1 The basic model of virus dynamics [19]. Uninfected cells are produced at a constant rate. Upon contact with free virus, they turn into infected cells. Infected cells produce new free virus particles. Uninfected cells, infected cells and free virus die at rates d , a and u respectively. In HIV-1 infection, for example, productively infected cells have a half-life of 1–3 days.



the zy terms). When infected cells burst, they release free virus, and these enter uninfected cells (xw). This leads to damped or undamped oscillations and even to Hopf bifurcations. Mathematical demographers work in close contact with clinical research teams, since it is clear that verbal, qualitative arguments are not sufficient to deal with the complex feedback loops occurring in such systems.

As an example of the intricacies involved, let us stress that each HIV contains several loci (so-called epitopes) which stimulate specific immune responses which otherwise would remain dormant. There may be seven or ten such epitopes, but let us assume, for the sake of the argument, that there are only two such loci, A and B, each with two possible molecular configurations. A_1 and A_2 stimulate the production of immune responses with the frequencies x_1 resp. x_2 , whereas B_1 and B_2 stimulate responses of magnitude γ_1 resp. γ_2 . If v_{ij} describes the frequency of virus of the type $A_i B_j$ (with $i, j = 1, 2$) then the very simplest ansatz yields already an eight-dimensional Lotka-Volterra equation. This system is not permanent: the important part of the dynamics occurs close to the boundary of the state space. If, for instance, only the type $A_1 B_1$ is present in the viral population, then one immune response – for instance x_1 – prevails over all other responses (and in particular over γ_1). This is the so-called immunodominant response [21]. But if mutation on the epitope A introduces a small minority of the viral strain $A_2 B_1$, then the immune system can either not react at all, or react by producing x_2 , or by producing γ_1 . In the latter case, we observe a shift in immunodominance: the immune response is now centered on the other epitope, although it had not been affected by the mutation. This phenomenon must be taken into account for the design of drugs and vaccines. One has to consider whether it is better to attack an epitope with a high responsiveness, or an epitope which is conservative in the sense that it is unlikely to mutate and hence to escape from the immunodominant answer.

HIV spends on average 1500 generations in the body of an infected host. Since it replicates very sloppily, many mutations occur during this period, leading to a continuous increase in genetic diversity. It may well be this diversity which overcomes the immune system in the end. Other viral and bacterial parasites also display high mutation rates (although in general not of the same magnitude as HIV). This pro-

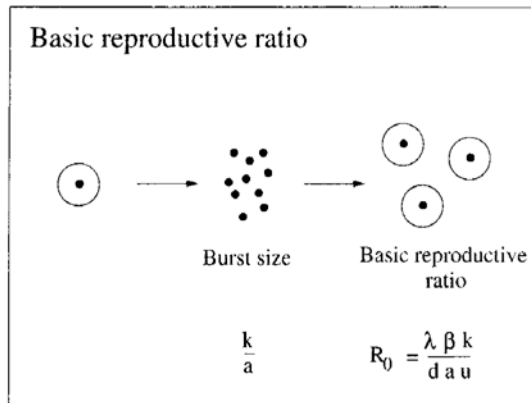


Fig. 2 The basic reproductive ratio of viral dynamics is defined a number of infected cells that are produced by any one infected cell (when almost all cells are still uninfected). In terms of the basic model of virus dynamics (Fig. 1) each infected cell produces on average k/a virus particles, a quantity that is usually denoted as 'burst size'. These virus particles produce R_0 newly infected cells. If $R_0 < 1$ the infection cannot proceed. If $R_0 > 1$ the infection can proceed. Successful anti-viral therapy has to reduce R_0 below 1. Resistance to antiviral therapy can be interpreted as R_0 values of certain mutants exceeding 1.

vides constantly new variants submitted to natural selection. On the one hand, these variants have to multiply quickly within the host, on the other hand, they must be good at transiting from one host organism to the next. These are in general quite different selective forces. Each genetic variant of the parasite corresponds to a different strategy.

For a long time biologists assumed that natural selection would always lead to a decrease in virulence: indeed, it is obviously in the basic, selfish 'interest' of the parasite to keep the host organism alive and mobile for as long as possible. Mathematical models have shown that this textbook wisdom was erroneous ([16], [22], [23]). An optimal parasite must have a basic reproductive ratio R_0 which is as large as possible. Equation (3) suggests that this entails a virulence v which is as small as possible. But this is wrong: in general, the infectiousness β is an increasing function of v , and therefore the optimal value for the virulence need not be the minimal value.

If co-infection occurs within the host organisms, i.e. if several strains of parasites compete within the same host, then optimisation of the virulence will not occur. A restraint in the exploitation of the host would only be to the benefit of the rival strains. The situation is similar to the so-called tragedy of the commons. Competition for a common resource leads to its over-exploitation. For such scenarios, optimisation arguments lead to the wrong conclusions. They have to be replaced by game theoretic arguments.

16.4 Evolutionary Games and the Dynamics of Adaptation

Natural selection essentially means that whenever an inheritable trait occurs in several variants within a population, then those variants which yield more offspring increase in relative frequency. Sub-optimal variants vanish if the reproductive success is independent of the frequencies of the variants. Such situations are doubtlessly common: for instance, the aerodynamical qualities of a birds wings are independent of what other birds are doing. On the other hand, the sex ratio is a trait whose success depends on the sex ratio of the other members of the population. If there is a surplus of males, it pays more to produce daughters, and vice versa.

Game theory can be used to analyse such instances of frequency dependent selection (see [24], [25] or [4]). The strategies are the different variants of some inheritable trait, and the payoff is the increase in fitness (i.e. reproductive success). The game dynamical replicator equation describes the time-evolution of the frequencies x_i of strategies i . They increase if i has more than average success, and decrease otherwise. But the changes in frequency can affect the success. The game dynamical replicator equation is given by

$$\dot{x}_i = x_i (f_i - \bar{f}) \quad i = 1, \dots, n \quad (6)$$

where $f_i = \sum_j a_{ij} x_j$ is the fitness of strategy i and $\bar{f} = \sum_i x_i f_i$ is the average fitness of the population. The payoff for i versus j is given by a_{ij} .

The state space is the unit simplex S_n (since $\sum x_i = 1$ must always hold). Again, the state space is invariant, and so are its boundary faces, where some strategies do not occur in the population. Transversal stability, permanence, invasion etc can be treated as before.

If there are only two strategies, then dominance, coexistence or bistability are again the three options. Let us assume, for instance, that there are only two possible behavioural types for conflicts within one species: individuals can either be ready to escalate the conflict until it is settled by the injury of one contestant, or else they can stick to some conventional display of strength and flee as soon as the rival starts to escalate the fighting. In this case, each of the two strategies will be able to invade a population where the resident type uses the other strategy: if all residents escalate, it is better to avoid a fight, whereas if all residents stick to conventional displays, it is better to escalate. As a result, natural selection will lead to a mixed population where the two behavioural morphs coexist, at a stable frequency. The proportion of individuals ready to escalate is given by G/C . Here, G is the gain in fitness from winning the object of the contest, and C is the cost of an injury (we assume $G < C$). It follows that the species which are heavily armed (high C) will be those most reliably sticking to conventional fighting – a result that was well known among students of animal behaviour long before Maynard Smith succeeded in explaining it by means of game theory.

If more than two strategies are present in the population, the dynamical outcome can be much more complex. In particular, heteroclinic cycles of rock-scissors-paper type can occur. Nothing may seem more abstruse than the idea of animals playing a

rock-scissors-paper game, but there exists actually a well documented case of three inheritable morphs A, B, C for male lizards of a certain species, where A is dominated by B , B by C and C by A again. These three morphs are characterised by different types of mating behaviour (which most conveniently correlate with throat colour). Type A males live monogamously, and guard their female jealously. They can be outcompeted by type B males, who have a larger harem, but cannot guard their females as closely. A type B population offers opportunities for type C males, who do not guard any females, but drift around and look out for opportunities to mate sneakily with an unguarded female. In a type C population, the more bourgeois type A can invade in turn, etc (see [26]).

By now, methods from evolutionary game theory are well established in animal behaviour. But these methods are also enriching classical game theory (see [27]), in spite of the fact that one of its basic assumptions, the rationality axiom, clearly has to go overboard. For instance, problems in bounded rationality or in equilibrium selection have become much more tractable. Arguments using population dynamical concepts (like migration, selection or invasion) are increasingly used even in social or economic applications of game theory. In particular, a plethora of approaches to learning mechanisms can be subjected to game dynamics. Strategies, here, are no longer inherited but imitated.

Although there always exist Nash equilibria, which are fixed points of the game dynamics, such equilibria need not be attractors. It may easily occur that under the effect of selection, or learning, almost no trajectory converges to such an equilibrium. Often, the behaviour is much more dynamical. This is the case, for instance, when we consider virulence as the strategy of a parasite. In the case of co-infection, or of superinfection (where it is assumed that the more virulent parasite instantaneously takes over in the host organism, see [28]), there is no evolutionarily stable degree of virulence. There is always the opportunity for new variants to enter the parasite population, and to eliminate some residents, in endless cycles which reflect the emergence and re-emergence of virulent strains.

Clearly, it is not only the parasite that evolves: the host is also submitted to a strong selective pressure. This leads to a co-evolutionary arms race, with alternating moves and countermoves by both populations. At first glance, the host seems hopelessly handicapped, since the parasite, usually, has a much shorter generation span. As soon as an immune system is fairly common, adaptation will bring forth parasites that are able to resist. Hence, an immune system can only be successful if it is rare. But if it is successful, it will spread and hence stop being rare.

However, there exists a way leading host species out of this impasse. By recombining sexually, they can keep all immune systems rare. Parasites will encounter a shifting target – or more precisely, a permanently reshuffled pool of genotypes controlling the immune answer. The corresponding models show endlessly repeated oscillations in the gene frequencies of both host and parasite species (see [29]).

This is an instance of the so-called Red Queen theory of coevolution. In the country of the Red Queen (a figure from *Beyond the Looking Glass*, the sequel of Lewis Carroll's *Alice in Wonderland*), all have to run as fast as possible just to stay at the same place. And indeed, the endless sequence of adaptations and counteradaptations may

appear like a treadmill leading nowhere. Nevertheless, there have been some undeniable instances of progress in biological evolution – the so-called major transitions in evolution, encompassing the discovery of sexual recombination, cell differentiation, the immune system of mammals or the nervous system (see [30]). In contrast to virulence, where the evolution takes place under our eyes, the major transitions are not subject to our direct observation. To understand them, one has to use through experiments and mathematical models.

The evolution of sex is arguably hardest to understand. Aside from the parasite-driven Red Queen theory for sex, there are several other contenders, most of them relying heavily on mathematical arguments (see [31]). And once we assume sexual recombination as given, other problems do come up (see [32]). Why are there usually two ‘sexes’ – two mating types A and B so that individuals of type A can mate with type B individuals, but not with type A individuals [33]? And why are there usually two and not three sexes? Why are the two sexes usually producing sex cells of different type, one of them producing few, but large gametes (the eggs) and the other one many tiny ones (the sperm), which often possess a high degree of mobility (see [34])? Once this basic difference between males and females has developed, other conflicts occur. Some of them concern parental investment. Males can have many more offspring than females. They are tempted to desert their mates, since females have much more to lose from failing to raise their brood. Females can counteract, for instance by insisting on a long engagement period. If females have the option of being coy (mate only after a long engagement) or fast, and males the options of being faithful (i.e. accept the engagement period) or philanderers, then adaptation and counteradaptation can lead to heteroclinic cycles in the strategies of the two sexes. Furthermore, the potentially higher reproductive success of males leads, on one hand, to combats between males, and on the other hand, to female choice based on male signals (for instance the peacock’s tail). All these are examples of frequency-dependent selection in action: the success of a trait depends on how abundant it is in the population. Evolutionary game theory has provided a theoretical framework for dealing with each of these questions (see [24]).

Sexual replication does not only bring drama into life. It also provides challenging tasks for mathematical modellers. They can no longer proceed on the assumption that like begets like if the genome of an individual stems from the combination of maternal and paternal genes. If we describe the frequencies x_i of the alleles (different types of genes at one chromosomal locus), we obtain again a replicator equation, but with more complicated interaction terms. In addition, most relevant traits depend not on one, but on several loci which can be recombined by genetic cross-over, and the investigation of the dynamics becomes particularly arduous (see e.g. [36]).

A promising alternative is to simplify the dynamics by ignoring the intricacies of some particular genetic system and analysing evolution, not in the space of gene frequencies, but in the trait space. In the simplest example, this consists in assuming that all individuals in the population are monomorphic (share the same trait-value), with the exception of occasional mutations introducing minorities having a slightly different trait value. Selection will then decide the fate of that minority. If one proceeds on the assumption that it will either get eliminated, or else eliminate the resi-

dent, before the next mutation introduces another minority, then one can model evolution by following the corresponding trait substitution sequence [36]. This sequence, which depends on the random order of arrival of mutants and therefore is a stochastic process, can often be approximated by a deterministic dynamics in trait space, the so-called adaptive dynamics. This adaptive dynamics usually provides a good tool for understanding long term evolution [37], even in the case of evolutionary chases of the 'Red Queen' type, which do not settle down to an evolutionary stable standstill – arms races between predators and their prey, for instance, or between hosts and parasites.

16.5

The Evolution of Cooperation

Returning to the major transitions which, in a way, transcend long term evolution, we note that occasionally, the fusion of independently replicating entities leads to entities of a higher order. For instance, the linkage of many genes results in a chromosome; the symbiosis of some protocells leads to the 'modern' eukaryotic cell, whose mitochondria and organelles retain only traces of their former independence; cells combine to form multicellular organisms; and organisms form societies. In each of these cases, such a cooperation can be threatened by the selfishness of parasitic exploitation. There exist so-called 'selfish genes', for instance, that rebel against the Mendelian rules of segregation and sabotage their opposite number on the partner chromosome, in order to be represented with more than their fair share in the next generation ([38] and [39]). This segregation distortion usually works to the detriment of the whole genome. Similarly, cancer cells multiply without regard to the multicellular organism, etc. Evolutionary game theory allows to model this tug-of-war between selfish interests and group benefits.

Not surprisingly, it is the last transition – the one leading to societies – which has attracted the most attention from evolutionary game theorists. After all, classical game theory was motivated by social questions. Societies built on mutual assistance are also one of the major themes of ethology. It is easy to see that they constitute a problem for Darwinism (and were in fact addressed by Darwin himself). Let us suppose that an act of help yields a benefit b for the recipient, and entails a cost c to the donor, with benefit and cost measured in terms of reproductive success. How can natural selection lead to the emergence of helping behaviour?

A first answer is based on kin selection. If the kinship between donor and recipient exceeds c/b , then a genetic disposition to help can spread under natural selection ([40] contains a historical introduction to this subject by its founder; see also [41]). Loosely speaking, a gene for helping relatives helps copies of itself. This explains cooperation in eusocial insects (in bee hives or termite hills, most individuals are very closely related, due to the fact that reproductive power is concentrated in a few queens). But it does not suffice to explain the cooperation in bands of hominids or stone age dwellers. An explanation, in this case, has to rely on economic rather than genetic arguments. An act of assistance will be of value to the donor whenever it is

returned with a high likelihood. But this, of course, provides an opportunity for cheating, by not returning help received.

In game theory, this free rider problem is often encapsulated in a game called the Prisoner's Dilemma. In this game, the two players have two options each, called **C** and **D** (to help or not to help, i.e. to defect). The payoff matrix is of the form

$$\begin{pmatrix} R & S \\ T & P \end{pmatrix} \quad (7)$$

with $T > R > P > S$ (if the role of a potential donor or recipient are equally likely, one has $T = b$, $R = b - c$, $P = 0$ and $S = -c$). Obviously **D** is the better move, regardless of whether the other player opts for **C** or **D**. Hence both players will play **D** and therefore obtain P instead of R – no cooperation.

Research in animal behaviour has uncovered many situations which are likely to reflect a Prisoner's Dilemma situation. Feeding or grooming each other, emitting warning calls, helping in territorial defense, inspecting or mobbing a predator are likely candidates (cf. [42]). In each of these cases, however, it remains possible to doubt whether the payoff matrix – whose terms are fitness increments – really satisfies the required inequalities $T > R > P > S$. It is extremely difficult to measure the fitness of fish darting in and out of shoals, of bats clustering in cave-roofs and of monkeys hiding in the bush. But recently, a well documented example of a Prisoner's Dilemma type of interaction has been uncovered for RNA-phages – virus reproducing in the inside of a bacterium ([43]). Here, the common type cooperates, whereas a certain mutant defects, by producing less than its share of the intracellular products needed for viral reproduction. Interestingly, the virus has found a way to overcome the dilemma, and therefore to do better than the fictitious rational player engaged in a Prisoner's Dilemma game. This makes one wonder if rationality is really the gift it is supposed to be.

It often happens that the probability to repeat another round of the Prisoner's Dilemma game *with the same partner* exceeds c/b . Computer simulations show that in this case, cooperation is quite likely to emerge. In the first series of computer tournaments addressing this issue, a very simple strategy did particularly well (see [44]): this was *Tit-For-Tat*, which tells a player to cooperate in the first round and then to opt for whatever move the co-player had chosen in the previous round. Such a strategy is nice, in the sense that it is never the first to defect, and it is quick to retaliate, answering a **D** with a **D** in the next round. Game theorists and evolutionary biologists have criticised the extreme error proneness of *TFT*, however. If, in an iterated game between two *TFT* players, one player defects by mistake, then the other player will retaliate, and this will cause a series of mutual defections which greatly reduces the average payoff and is only stopped by another mistake (such a mistake is equally likely to lead back to mutual cooperation as to make both players keep defecting simultaneously). Errors in the implementation of a move, or in the perception of the co-player's move, are inevitable in every real life type of interaction. What happens in the presence of such 'noise'?

In the simplest case this can be investigated by means of strategies specified by their first move and by a vector (p_R, p_S, p_T, p_P) , where p_k is the probability to play **C** if




Fig. 3 Evolutionary chronicles of the iterated Prisoner's Dilemma reflect cycles of cooperation and defection. Populations of selfish defectors (ALLD) can be invaded by small clusters of Tit-for-tat players. Once cooperation is established more forgiving strategies, such as Generous Tit-for-tat, are favoured by natural selection. Unconditional cooperators can invade TFT or GTFT populations by neutral drift, thereby undermining a population's ability to reciprocate, which in turn allows the invasion of defectors. These endless cycles between cooperation and defection are interrupted when cooperative populations are invaded by Win-stay, lose-shift (Pavlov) players. Unlike TFT and GTFT, Pavlov dominates ALLC.

the payoff in the previous round was $k \in \{R, S, T, P\}$. For instance, $(1, 0, 1, 0)$ is the strategy *TFT*. In the presence of a small error probability, this is dominated by the strategy $(1, 0, 1, 1)$ (called *FirmButFair*, *FBF*) which plays C after a mutual defection. This more tolerant strategy is dominated both by the strategy $(0, 0, 0, 0)$ (*Always Defect*) and by the strategy $(0, 0, 0, 1)$ (*Bully*), which both in turn are dominated by *TFT*. There is, of course, no reason to restrict oneself to these four strategies only (which constitute a heteroclinic attractor, incidentally). Mutation will keep introducing small amounts of new strategies, which then will be subjected to natural selection. Depending on the composition of the population, they will increase or decrease in frequency. It turns out (see [45]) that if $b > 2c$, then evolution leads in the long run to populations whose members all play the so-called *Pavlov* strategy $(1, 0, 0, 1)$.


Pavlov is the simplest conceivable learning rule. Whenever the payoff in the previous round was high (a *T* or an *R*), the previous move will be repeated by a *Pavlov*-player. After a low payoff *P* or *S*, the other move will be tried. Two *Pavlov* players will cooperate almost always. If one of them erroneously plays *D*, then he (or she) will repeat the *D* move in the next round (since it yielded the high payoff *T*). But the other player, who has been cheated, will switch from the former *C* move to *D*. As a result, both players defect (on purpose) after an erroneous defection. Both will be dissatisfied by the outcome *P* and therefore both will switch back to *C*.

It can be shown that *Pavlov* is a social norm in a precise sense. Within a population adopting such a norm, any individual who deviates at any stage of the game will be penalised, even if the deviation occurs only after a mistake and is not expected in the usual run of the interaction. However, it must be noted that *Pavlov* players can never, by themselves, invade a population of *Always Defect* players. For this, it needs a strongly retaliatory strategy like *TFT* or *Grim* (the strategy $(1, 0, 0, 0)$). It is only when such stern retaliators have eliminated the invertebrate defectors that a more tolerant strategy like *Pavlov* can come to the fore. Such thought experiments (mostly performed with the help of computer simulations) show how, under minimalistic assumptions, such interesting aspects as learning rules, social norms or historical stages come up quite naturally.

It should be noted that cooperation can also take hold if individuals play the game for several rounds, but *never* with the same co-player twice (see [46]). It is enough to assume that the information about the co-player exceeds the cost-to-benefit ratio c/b .

Cooperation is then provided by discriminating strategies which only provide help if the recipient, to their knowledge, has sufficiently often helped others. Such discriminators are sometimes refusing to provide help, which in itself can be risky because it reduces the probability that they will be helped at some later stage, by other discriminators. The game dynamics is quite complex, but a basic aspect can be understood easily by a further simplification of the model. Let us assume that each player, in his or her entire lifetime, plays only two rounds, never meeting the same partner twice, and consider the replicator equation describing the relative frequencies of three strategies: the discriminators, the indiscriminate altruists (who always help) and the invertebrate defectors, who never help. In the absence of the defectors, discriminating and indiscriminating altruists do equally well, since they always cooperate. This means that there is no selective pressure, and that the composition of a population consisting of these two types only will slowly change by neutral drift, with the frequency of discriminators sometimes increasing and sometimes decreasing according to a random walk. Random fluctuations will occasionally introduce, through mutation, a small minority of defectors. How will they fare? If the frequency of discriminators is sufficiently large, they will immediately be selected against and eliminated. If the frequency of discriminators is very low, on the other hand, then the defectors will be able to exploit the unconditional altruists to the hilt and will eventually take over the entire population. However, if the frequency of defectors is in a middle range, not too high and not too low, then the defectors will initially be able to grow, and to exploit the altruists. Eventually, however, the indiscriminate altruists will be decimated to such an extent that the defectors will have mostly to deal with discriminators; and then, they will be eliminated, and the population returns to a mixture of discriminate and indiscriminate altruists, but this time (in contrast to the initial situation) with such a high frequency of discriminators that any follow-up attempt by the defectors to invade again will be instantly repelled. The short-term increase of defectors, in this evolution, is a self-defeating move, therefore, a kind of Pyrrhic victory. The next chance for defectors to invade will only come if, by random drift, the frequency of discriminators has decreased again. The only way for cooperation to break

Fig. 4 Game dynamics of indirect reciprocity can reveal the surprising feature that invasion of a strategy is only successful when rare. The 3 corner points of the simplex refer to unconditional cooperators, e_1 , unconditional defectors, e_2 , and discriminators, e_3 , which cooperate with individuals who have a good reputation. There is neutral drift between cooperators and discriminators. If the frequency of discriminators is sufficiently high, defectors cannot invade. If the frequency of discriminators is very low, defectors can invade and take over. Between these threshold values exists a region where defectors can invade, but become extinct again after some time leaving the population in an uninvadable state with a high frequency of discriminators. Thus, if invasion attempts of defectors are frequent the population will often be in a state with many discriminators.



down permanently is when the time until the next invasion attempt by defectors is large enough for the random drift to decrease the frequency of discriminators by a considerable margin. It follows that the exploiters, if they are to take over, should not try too often to invade. In other terms, a cooperative population will only be proof against parasitic exploiters if it is challenged sufficiently often by their invasion attempts. This is an intriguing parallel to the immune system, which only remains efficient if it is challenged sufficiently often.

Such parallels make it seem plausible that essentially the same methods of evolutionary game dynamics are applicable in such a wide range, spanning all the way from the evolution of virulence or the mating behaviour of lizards to molecular evolution and the emergence of cooperative societies. Frequency dependent selection operates across all scales of biological communities, and is instrumental in shaping societies.

References

- 1 MAYR, E., *Populations, Species, and Evolution*, Harvard UP, (1983)
- 2 LEVIN, S.A. (1992) (Ed), *Mathematics and Biology: The Interface* Lawrence Berkeley Laboratory, Univ. of California, Berkeley, CA.
- 3 LEVIN, S.A., GRENFELL, B., HASTINGS, A. and PERELSON, A.S. (1997), Mathematical and computational challenges in population biology and ecosystems science, *Science* 275, 334–343
- 4 J. HOFBAUER and K. SIGMUND, *Evolutionary games and population dynamics*, Cambridge UP (1998)
- 5 TAKEUCHI, Y. (1996), *Global dynamical properties of Lotka-Volterra systems*, World Scientific, Singapore.
- 6 MAY, R.M. and W.J. LEONARD (1975) Nonlinear aspects of competition between three species, *SIAM J. Appl. Math.* 29 243–253.
- 7 WEISSING, F. (1991) Evolutionary stability and dynamics stability in a class of evolutionary normal form games. In R. SELTEN (ed) *Game Equilibrium Models I: Evolution and Game Dynamics*. Berlin, Springer, 29–97.
- 8 V. HUTSON and K. SCHMITT, (1992) Permanence and the dynamics of biological systems, *Math Biosci.* 111, 1–71
- 9 FERRIERE, R. and M. GATTO (1995) Lyapunov exponents and the mathematics of invasion in oscillatory or chaotic populations, *Theor. Pop. Biol.* 48, 126–71.
- 10 RAND, D., WILSON, H.B. and MCGLADE, J.M. (1994): Dynamics and evolution: evolutionarily stable attractors, invasion exponents and phenotype dynamics, *Phil. Trans. Roy. Soc. London B* 24, 261–283.
- 11 HOFBAUER, J. (1999) Invasion, permanence and heteroclinic cycles, to appear.
- 12 R. LAW and J.C. BLACKFORD (1992) Self-assembling food webs: a global viewpoint of co-existence of species in Lotka-Volterra communities *Ecology* 73, 567–578.
- 13 A. HASTINGS, 'Food web theory and stability', *Ecology* Vol. 69 (1988), pp. 1665–1668.
- 14 NEE, S. (1990), Community construction, *Trends in Ecology and Evolution* 5, 337–340.
- 15 GRENFELL, B.T. and A.P. DOBSON (1995), *Ecology of Infectious Diseases in Natural Populations*, Cambridge Univ. Press, Cambridge.
- 16 ANDERSON, A.M. and R.M. MAY (1991), *Infectious Diseases of Humans: Dynamics and Control*, Oxford Univ. Press, Oxford.
- 18 DIEKMANN O, HEESTERBEEK J A P, and METZ J A J (1990) On the definition and the computation of the basic reproductive ratio R_0 in models for infectious diseases in heterogeneous populations, *J Math Biol* 28, 365–382

- 19 NOWAK, M.A. and BANGHAM, C.R. (1996), Population dynamics of immune responses to persistent viruses, *Science* **272**, 74–9.
- 20 DEBOER, R.J. and A.S. PERELSON (1998) Target cell limited and immune control models of HIV infection: a comparison, *Journ. Theor. Biol.* **190** 201–14.
- 21 NOWAK, M.A., R.M. MAY and K. SIGMUND (1995) Immune responses against multiple epitopes, *Journ. Theor. Biol.* **175**, 325–53.
- 22 LEVIN, S.A. and D. PIMENTEL (1981) Selection of intermediate rates of increase in parasite-host systems *Amer. Nat.* **117** 308–1.
- 23 FRANK, S.A. (1996) Models of parasite virulence, *Quart. Rev. Biol.* **71**, 37–78.
- 24 MAYNARD SMITH, J. (1982) *Evolution and the theory of games*, Cambridge UP.
- 25 CRESSMAN, R. (1992) *The stability concept of evolutionary game theory*, Springer, Berlin.
- 26 SINERVO, B. and LIVELY, C.M. (1996), The rock-scissors-paper game and the evolution of alternative male strategies, *Nature* **380**, 240–43.
- 27 WEIBULL, J. *Evolutionary Game Theory*, Cambridge, MA, MIT Press.
- 28 NOWAK, M.A. and R.M. MAY (1994), Superinfection and the evolution of parasitic virulence, *Proc. R. Soc. London B* **255**, 81–89.
- 29 HAMILTON, W.D. (1980) Sex versus non-sex versus parasites, *Oikos* **35**, 282–90
- 30 MAYNARD SMITH, J. and SZATHMARY, E. (1995), *The major transitions in evolution*, Freeman, Oxford.
- 31 STEARNS, S.C. (ed) (1987) *The evolution of sex and its consequences*, Birkhäuser, Basel.
- 32 SIGMUND, K. (1995), *Games of Life*, Penguin, Harmondsworth.
- 33 HUTSON, V. and R. LAW (1993), Four steps to two sexes, *Proc. Roy. Soc. London B*, **253**, 43–51.
- 34 HOEKSTRA, R., The evolution of sexes, in [31], pp. 59–92.
- 35 IWASA, Y. and A. SASAKI (1987), Evolution of the number of sexes, *Evolution* **41**, 49–65.
- 36 METZ, J.A.J., S.A. GERITZ, G. MESZENA, F.J.A. JACOBS and J.S. VAN HEERWARDEN (1996), Adaptive dynamics: a geometrical study of the consequences of nearly faithful replication. In S.J. VAN STRIEN and S.M. VERDUYN LUNEL (eds), *Stochastic and spatial structures of dynamical systems*, 183–231. Amsterdam, North Holland.
- 37 ESHEL, I. (1996): On the changing concept of evolutionary population stability as a reflection of a changing point of view in the quantitative theory of evolution, *J. Math. Biol.* **34**, 485–510.
- 38 CROW, J.F. (1986) Basic concepts in population, quantitative and evolutionary genetics, Freeman, New York
- 39 HAIG, D. and A. GRAFEN, (1991) Genetic scrambling as a defense against meiotic drive, *J. Theor. Biol.* **153**, 531–58.
- 40 HAMILTON, W.D. (1996) *Narrow roads to gene land*, Oxford UP, Oxford.
- 41 FRANKS, S.A. (1998) *Foundations of Social Evolution*, Princeton UP, Princeton.
- 42 DUGATKIN, L.A. (1998) *Cooperation among Animals: an Evolutionary Perspective*, Oxford UP, Oxford.
- 43 TURNER, P.E. and CHAO, L. (1999) Prisoner's Dilemma in an RNA virus, *Nature* **398**, 441–3
- 44 AXELROD, R. (1989) *The Evolution of Cooperation*, Penguin, London.
- 45 NOWAK, M.A. and SIGMUND, K. (1993), Win-stay, lose-shift outperforms tit-for-tat, *Nature*, **364**, 56–8.
- 46 NOWAK, M.A. and SIGMUND, K. (1998), Evolution of indirect reciprocity by image scoring, *Nature* **393**, 573–7.