IIASA

International Institute for Applied Systems Analysis • A-2361 Laxenburg • Austria
Tel: +43 2236 807 • Fax: +43 2236 71313 • E-mail: info@iiasa.ac.at • Web: www.iiasa.ac.at

INTERIM REPORT IR-98-102 / December

The Population Dynamics of Conflict and Cooperation

Karl Sigmund (ksigmund@esi.ac.at)

Approved by
Ulf Dieckmann (dieckman@iiasa.ac.at)
Project Coordinator, *Adaptive Dynamics Network*

Contents

About the Author

Karl Sigmund
Institut für Mathematik
Universität Wien
Strudlhofgasse 4
A-1090 Vienna, Austria
and
Adaptive Dynamics Network
International Institute for Applied Systems Analysis
A-2361 Laxenburg, Austria

The Population Dynamics of Conflict and Cooperation

Karl Sigmund

1 Introduction

The last decades have seen an explosive growth in biosciences, and astonishing progress in the mathematical modelling of fields as diverse as neurobiology, membrane formation, biomechanics, embryology, etc. (see e.g. J. Murray, 1990). The sequencing of biomolecules produces such a vast wealth of data on proteins and polynucleotides that the mere handling of the stored information becomes a computational challenge, let alone the analysis of phylogenetic trees and functional networks which is the main task of bioinformatics.

The recent advances in our understanding of the chemical mechanisms describing the interactions of specific molecules – how virus, for example, use binding proteins to attack and penetrate hosts cells – are spectacular, but do not suffice to tackle basic problems like disease progression or the co-evolution of hosts and parasites. It is *populations* of virus particles, or immune cells, or hosts, that regulate each other's frequencies. The feedback loops of these ecosystems are too complex to be understood by verbal arguments alone. The biological community has come to accept that basic aspects of immunology and evolutionary ecology can only be analysed by mathematical means.

This has not always been the case. The pioneering work in genetics due to Fisher, Haldane, Wright, and Kimura, as well as the epidemiological models of Kermack and McKendrick occupied a marginal position in biology for the most part of this century, while at the same time motivating important mathematical advances in statistics, stochastic processes and dynamical systems (Fisher (1918) on correlation, Kolmogoroff (1937) on travelling waves in a gene pool, May (1976) on chaos). The models of evolutionary biology cannot compete in mathematical depth and sophistication with those of theoretical physics, but they offer a wide range of questions of great intuitive appeal.

This lecture surveys mathematical models in ecology and evolution, emphasising the major feedback mechanisms regulating the population densities of the interacting self-replicating units – be they genes, virus particles, immune cells or host organisms. The great variety of biological examples made it necessary to economise on mathematical diversity, by keeping to the framework of ordinary differential equations. This is certainly not meant to imply that time delays, spatial heterogeneities and stochastic fluctuations are secondary effects. In fact, they have a major impact in many applications (see, e.g., the survey by Levin et al., 1997)

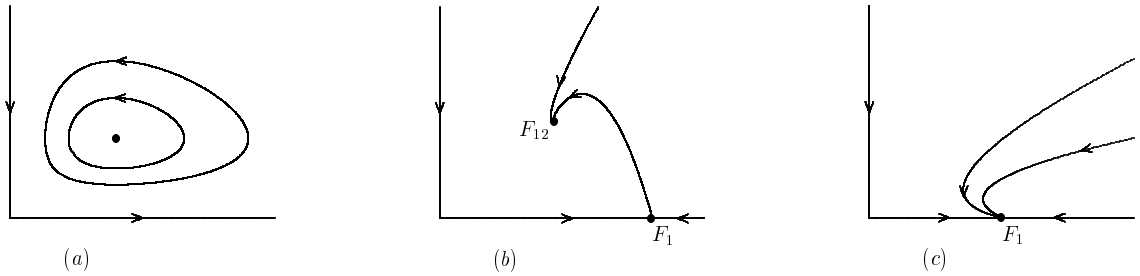

Figure 1: Predator-prey equations

2 Population Ecology

If we assume that n species live in an ecosystem, that x_i is the density of species i and that its per capita growth rate $\dot{x}_i/x_i$ depends on the densities of the interacting populations, then we obtain the ecological equation

$$\dot{x}_i = x_i f_i(\mathbf{x}). \tag{1}$$

The state space R^n_+ is invariant; so are its boundary faces, where one or several of the densities are 0; and the restriction of (1) to a face is again an ecological equation. If the f_i are affine linear, we obtain – as simplest example – the Lotka-Volterra equation

$$\dot{x}_i = x_i(r_i + \sum a_{ij}x_j) \tag{2}$$

$(i = 1, ..., n)$. It should be stated right at the outset that many ecological interactions display more complex interaction terms; but often, (2) offers a first approximation which is flexible enough to embody the main aspects of the community structure (Hofbauer and Sigmund, 1998). For instance, if 1 is a prey species and 2 its predator, we obtain

$$\dot{x}_1 = x_1(a - bx_2) \tag{3}$$

$$\dot{x}_2 = x_2(-c + dx_1). \tag{4}$$

where $a, b, c, d > 0$. In $intR^2_+$ there exists a unique fixed point $(c/d, a/b)$ which is surrounded by periodic orbits. If we add a self-limitation of the prey, i.e. set $f_1 = a - ex_1 - bx_2$ in (3), we obtain damped oscillations around the fixed point, or (if $e > 0$ is large) extinction of the predator (see Fig. 1):

On the other hand, if 1 and 2 are species competing for the same resources, we have to assume that the intrinsic growth rates satisfy $r_i > 0$ and the interaction terms $a_{ij} < 0$ $(i, j \in \{1, 2\})$. On each positive half-axis, there is one fixed point $\mathbf{F}_i$ corresponding to equilibrium of species i in the absence of the other species. Generically, there are three possible outcomes (see Fig. 2): (a) *dominance*: all orbits in $intR^2_+$ converge to $\mathbf{F}_i$; species i is said to dominate the other species;

(b) *coexistence*: there exists a fixed point $\mathbf{F}_{12} \in intR^2_+$ which is globally stable (i.e. attracts all orbits in $intR^2_+$);

(c) *bistability*: $\mathbf{F}_{12}$ is a saddle; almost all orbits in $intR^2_+$ converge to $\mathbf{F}_1$ or $\mathbf{F}_2$, depending on the initial condition.

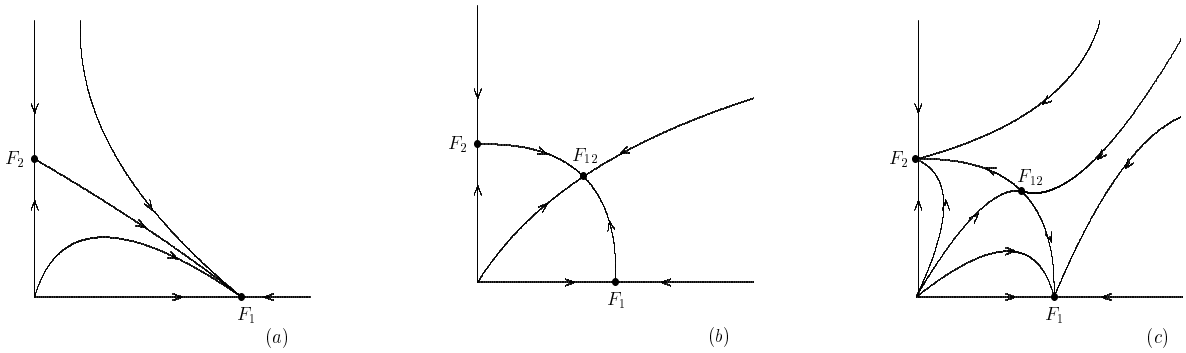

Figure 2: Competition equations

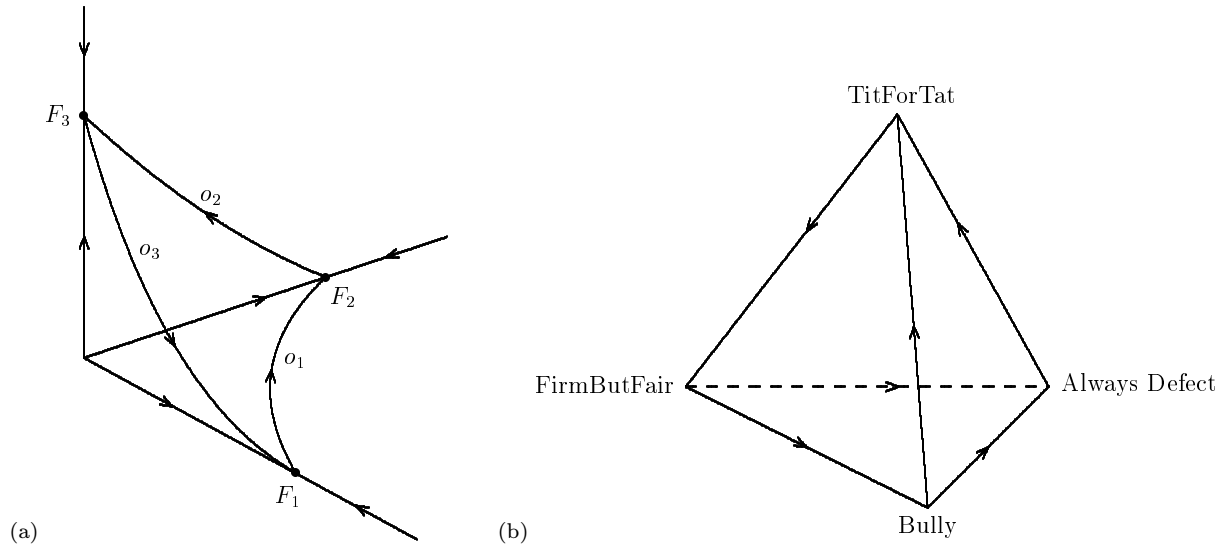

Figure 3: Heteroclinic orbits and networks

Because two-dimensional Lotka-Volterra equations admit no limit cycles, their dynamics can be easily classified; for three or more species, this is no longer the case. Systems with two competing species and one prey exhibit chaos (Gilpin, 1979) and systems with three competing species (which are monotonic and hence admit no chaos, see Hirsch, 1988) have not been classified yet, in spite of impressive progress (van den Driessche and Zeeman, 1997). One of the reasons is the existence of *heteroclinic cycles* see Fig. 3a (May and Leonard, 1975). If, in the absence of the third species, species 1 dominates 2, 2 dominates 3 and 3, in turn, dominates 1, then the boundary of R^3_+ contains a heteroclinic cycle consisting of three saddle points $\mathbf{F}_i$ (with only species i present) and three connecting orbits (orbit o_1 has $\mathbf{F}_2$ as α- and $\mathbf{F}_1$ as ω-limit etc). Depending on the products of the eigenvalues in the stable and unstable directions, this heteroclinic cycle can attract or repel the neighbouring orbits in $intR^3_+$. Three competing species with heteroclinic cycles have been found in laboratory populations. In higher dimensional ecological models, heteroclinic cycles become common. Such cycles are non-generic features for general dynamical systems, since saddle-connections can be destroyed by arbitrarily small

perturbations. Within the class of ecological equations, however, which leave the boundary faces of R^n_+ invariant, heteroclinic cycles and *networks* (where several cycles issue from one saddle) are usually robust. Such attractors offer a new brand of nonlinear dynamics: orbits approach saddle points ever more closely, and remain there for increasingly long times; furthermore, the sequence of saddles visited by an orbit can switch in arbitrary order from one cycle to another (Chawanya, 1995).

3 Permanence

If the orbit of an ecosystem reaches the neighborhood of a heteroclinic attractor on the boundary, some species are doomed. The ecosystem, in that case, is unstable: this kind of stability has nothing to do, however, with the usual asymptotic stability of a fixed point, which is a local notion. A more suitable stability notion in this context is that of *permanence*: (1) is said to be permanent if the boundary (including infinity) is a repellor, i.e. if there exists a compact set $K \subset intR^n_+$ such that whenever initially $\mathbf{x} \in intR^n_+$, then $\mathbf{x}(t) \in K$ for t sufficiently large. (After a transient phase, all densities are uniformly bounded away from 0). This notion has been extensively explored (see the survey by Hutson and Schmitt, 1992). Permanence implies the existence of a fixed point in $intR^n_+$, but this point need not be locally stable; and indeed ecologists view an ecosystem as stable even if it exhibits violent oscillations, as long as its species remain safe from extinction.

For a dissipative system (all orbits uniformly bounded from above), the most useful sufficient condition for permanence is the existence of an average Lyapunov function. This is a function P vanishing on the boundary and positive on the interior such that the continuous extension Ψ of the logarithmic derivative of P has the property that for every ω-limit point $\mathbf{x}$ on bdR^n_+ there is a $T > 0$ with

$$\int_0^T \Psi(\mathbf{x}(t))dt > 0. \tag{5}$$

Then P grows (in the long run) along every interior orbit sufficiently close to the boundary. In particular (2) is permanent if all orbits are uniformly bounded and the set

$$D := \{\mathbf{x} \in R^n_+ : r_i + \sum a_{ij}x_j < 0, i = 1, ..., n\} \tag{6}$$

is disjoint from the convex hull of the fixed points on the boundary. The condition is not necessary for permanence if $n > 3$. But if (2) is permanent, then there is a *unique* equilibrium $\hat{\mathbf{x}}$ with all species present, and it is the limit of all time-averages of orbits in the interior of the state-space. If D is the Jacobian at $\hat{\mathbf{x}}$, then $(-1)^n \det D > 0$, and trace $D < 0$. Furthermore, $(-1)^n \det A > 0$, where A is the matrix of the interaction terms a_{ij} (Hofbauer and Sigmund, 1998).

4 Invasion

Many studies have considered the assembly of ecological communities by sequential invasion (i.e. adding one species at a time). Will species $n+1$ grow when introduced in small numbers? If the resident system is in equilibrium $\mathbf{z} = (z_1, ..., z_n)$, this simply

means to check whether the growth rate $f_{n+1}(\mathbf{z}, 0)$ is positive. If the competition between two species is bistable, for instance, none can invade the other. If there is coexistence, each can invade, etc. Invasion is a question of *transversal stability*, which, if the resident system admits a chaotic attractor, offers subtle ergodic twists involving riddled basins of attraction etc (Ferriere and Gatto, 1995, Ashwin et al, 1996).

If the resident species obey a permanent Lotka-Volterra equation with fixed point $\mathbf{z} \in intS_n$, the condition $f_{n+1}(\mathbf{z}, 0) > 0$ implies that the lim sup of the invading species' density is positive, but tells nothing about the lim inf. The new attractor need not be close to the former one; the invading species can drive others to extinction, and even ultimately itself. Hofbauer (1998) has found conditions in terms of spatial or temporal averages of the initial growth rate which guarantee that the invasion of a permanent Lotka-Volterra community succeeds. His bifurcation analysis allows to decide whether, if a parameter changes so that invasion becomes possible, the new attractor is contained in a neighborhood of the resident attractor or not. The invasion of a heteroclinic cycle is a particularly arduous problem.

Evidence from field studies and numerical simulations suggest that ecosystems become increasingly harder to invade as time goes on, and that there is an upper limit to how 'closely packed' species can be; but so far, this has only been demonstrated under restrictive assumptions (Law and Blackford, 1992). Interestingly, predators can stabilise ecosystems: if a 'keystone' predator is removed from a permanent system, the remaining system is no longer permanent. For instance, if species 1 dominates species 2, or if the competition between species 1, 2 and 3 results in a heteroclinic attractor, then a suitable predator can mediate co-existence; Schreiber (1998) has produced systems with n competing prey, each with its specialised predator, such that removal of one predator species results in only one prey species surviving. Such ecosystems cannot be obtained by simply adding one species at a time; sequential assembly has to proceed in a more roundabout way, using species that are later eliminated like a scaffolding. These results agree well with the current emphasis of biologists on the role of contingency and history dependence in real ecological succession chronicles, and highlight the fact that a successful invasion can initiate a complex sequence of changes in the ecosystem (see Mylius et al, 1998).

5 Replicator Dynamics

Competition between conspecifics drives natural selection. The basic mechanism is simple: an inheritable trait which allows for a higher reproductive success spreads in the population. This can lead to extraordinary feats of adaptation due to relentless *optimisation under constraint*. In fact, some computational approaches to optimisation problems are mimicking the massively parallel algorithm of Darwinian evolution. Within 'populations' of possible solutions to a given problem (for instance in aerodynamics), those which perform better are allowed to multiply at the expense of the others. Occasionally, some 'offspring' is randomly altered, corresponding to the mutation or recombination of existing solutions. Such *genetic algorithms* allow to explore the space of solutions and often to home in on some optima (Forrest,

1993).

But in biology, it is the population itself that is often the problem. The efficiency of a wing shape may be independent on what the other birds are doing, but the success of a sex ratio or of a given degree of aggressivity is not. In a population with a surplus of males, it pays to produce females; it pays to escalate a conflict if the others are unlikely to escalate, but otherwise it is better to avoid escalating, etc. *Game theory*, rather than optimisation, is appropriate to deal with problems where the success depends on what the others are doing.

Assume that x_i is the frequency of the individuals using strategy i ($i = 1, ..., n$). A strategy, in this context, is simply a trait (behavioural, physiological, morphological) whose payoff, i.e. average reproductive success, depends on the frequencies $\mathbf{x}$ of the competing types. If the traits are inherited, the frequencies will evolve in time, depending on their success. If individuals breed true, the per capita rate of increase $\dot{x}_i/x_i$ is given by the difference $f_i - \bar{f}$, where $f_i(\mathbf{x})$ is the average payoff for using i if the population is in state $\mathbf{x}$, and $\bar{f} = \sum x_j f_j$ is the average success in the population. This yields the replicator equation

$$\dot{x}_i = x_i(f_i(\mathbf{x}) - \bar{f}) \quad i = 1, ..., n \tag{7}$$

on the simplex $S_n = \{\mathbf{x} \in R^n_+ : \sum x_i = 1\}$. This simplex is invariant, and so are its faces. The replicator equation is closely related to the ecological equation (1), of course. It introduces an ecological viewpoint into game theory.

Let us consider a conflict between pairs of individuals, for instance some contest over a resource, and assume that the strategies i correspond to different types of fighting behaviour, and that a_{ij} is the average payoff for using i if the co-player uses j. Then the payoff matrix $A = (a_{ij})$ determines the average payoff $(A\mathbf{x})_i = a_{i1}x_1 + ... + a_{in}x_n$ for strategy i in the population (assuming that individuals meet randomly) and (7) turns into

$$\dot{x}_i = x_i((A\mathbf{x})_i - \mathbf{x}^T A\mathbf{x}). \tag{8}$$

This equation is not only similar, but actually equivalent to a Lotka-Volterra equation for $n - 1$ species: a diffeomorphism from S_n (minus one face) to $\mathbf{R}^{n-1}_+$ maps orbits of one dynamical system onto the other, and vice versa (Hofbauer 1981). For $n = 2$, we obtain the same generic behaviour as for two competitors: dominance, coexistence or bistability. For $n = 3$, heteroclinic cycles show up (not just as a theoretic possibility: the mate guarding strategies of male lizards form a rock-scissors-paper cycle). With $n > 3$, limit cycles and chaotic attractors occur. (8) is permanent if there exists a $\mathbf{p} = (p_1, ..., p_n)$ with $p_i > 0$ for all i such that for every equilibrium $\mathbf{z}$ on the boundary,

$$\mathbf{p}^T A\mathbf{z} > \mathbf{z}^T A\mathbf{z} \tag{9}$$

(a conditions that can easily be checked by linear programming), etc.

Frequency-dependent selection will not optimise, in general. Only for very special interaction do replicator equations become gradients: if the game is symmetric, for instance ($A = A^T$) or more generally if the partial derivatives $f_{i,j} = \partial f_i/\partial x_j$ obey

$$f_{i,j} + f_{j,k} + f_{k,l} = f_{l,k} + f_{k,j} + f_{j,i} \tag{10}$$

for all i, j, k (one has to use a suitable Riemannian metric on S_n, cf. Hofbauer and Sigmund, 1998).

6 Other Game Dynamics

Among higher animals, and in particular humans, strategies can also spread by learning and imitation. Depending on the details of transmission, this leads to a large number of game dynamics for the frequencies x_i, often based on underlying stochastic processes. Again, the replicator dynamics is a kind of benchmark. Another example is the best reply dynamics (a differential inclusion)

$$\dot{x}_i \in \beta(\mathbf{x}) - x_i \tag{11}$$

where $\beta(\mathbf{x})$ is the set of strategies whose payoff (in a population where strategy i occurs with frequency x_i) is maximal. The idea is that in every short time interval, a small fraction of the players updates their strategy: these players know how to optimise, but do not anticipate that others will also update. The orbits of (11) are piecewise linear. Intriguingly, their asymptotic behaviour is often that of the *time averages* of the solutions of the replicator equation (8).

This brings one closer to classical game theory. Let us consider a game with payoff matrix A and assume that points $\mathbf{p} \in S_n$ are *mixed strategies* (p_i being the probability for a player to use strategy i). Then $\mathbf{p}$ is a best reply to $\mathbf{q} \in S_n$ if $\mathbf{p}^T A\mathbf{q} \geq \mathbf{x}^T A\mathbf{q}$ for all $\mathbf{x} \in S_n$. A point $\mathbf{p}$ is a (symmetric) Nash equilibrium if it is a best reply against itself. A Nash equilibrium is a fixed point for (8) (and every other decent game dynamics), but the converse need not hold. In fact, the Nash equilibria are precisely the fixed points of (8) which are *saturated* – missing pure strategies have no selective advantage. Every game with finitely many strategies has a Nash equilibrium, but there are games such that almost no solution, under any reasonable adjustment dynamics, converges to a Nash equilibrium.

Evolutionary game theory has originated with the concept of *evolutionarily stable strategies* (ESS). Intuititively, a strategy $\mathbf{q}$ is said to be an ESS if, whenever all members of the population adopt it, an invading (and sufficiently small) minority using a different strategy has no selective advantage (Maynard Smith, 1982). This means that $\mathbf{q}$ is Nash, and that whenever $\mathbf{p}$ is an alternative best reply to $\mathbf{q}$, then $\mathbf{q}^T A\mathbf{p} > \mathbf{p}^T A\mathbf{p}$. Equivalently, $\mathbf{q}$ is an ESS if

$$\mathbf{q}^T A\mathbf{x} > \mathbf{x}^T A\mathbf{x} \tag{12}$$

for all $\mathbf{x} \neq \mathbf{q}$ in a neighborhood of q. Not every game has an ESS. The connexion with the replicator equation is given by the following characterisation: $\mathbf{q} \in S_n$ is an ESS if and only if, whenever $\mathbf{q}$ is a convex combination of the (possibly mixed) strategies $\mathbf{p}^1, ..., \mathbf{p}^m$, the *mean population strategy* $\sum x^i_{\mathbf{p}}$ converges (under the replicator dynamics) towards $\mathbf{q}$ if initially it was close to $\mathbf{q}$ (Cressman, 1992). The idea that evolution always results in an ESS is not justified, however. There exist considerably more complex outcomes, as captured in the notion of an evolutionarily stable *attractor*, for instance (Rand et al, 1994).

7 Long-term Evolution

So far we have assumed that offspring are clones of their parent: 'like begets like'. The machinery of Mendelian inheritance is much more complex, and we have to

follow the frequencies of genes in the *gene pool* of the population. As long as the instruction is contained in one genetic locus (an address in the genome, housing two genes – one from the father and the other from the mother), the corresponding dynamics for the gene frequencies in the population is still of replicator type (7). But in general, the trait depends on several genetic loci, which can be recombined during reproduction, and the dynamics becomes challenging.

The state $\mathbf{x}$ of the gene pool determines the frequencies of the different types of individuals, who use different (pure or mixed) strategies. This determines the frequencies $\mathbf{p}(\mathbf{x})$ of the strategies in the population, and hence the reproductive success of each type, and therefore the rate of change in the gene frequencies $\mathbf{x}$. If the trait is determined by one genetic locus only, and if there are at most two pure strategies, or three types of genes which can occur on that locus, then an ESS $\mathbf{q}$ which is feasible is *strategically stable* in the sense that if a state $\hat{\mathbf{x}}$ of the gene pool satisfies $\mathbf{p}(\hat{\mathbf{x}}) = \mathbf{q}$, then every near-by state $\mathbf{x}$ remains close to $\hat{\mathbf{x}}$ and $\mathbf{p}(\mathbf{x})$ converges back to $\mathbf{q}$ (Cressman and Hofbauer 1997). For more complex genetic mechanisms, the relation between evolutionary stability and long-term stability (i.e. strategic stability against every invasion attempt) remains unclear, and offers a wealth of problems on normal forms and center manifold theory. The replicator dynamics can be used as a first approximation in the absence of more specific information on the genetic background. That kind of information is likely to be provided soon, and will act as a motivational booster for the population genetics of frequency-dependent selection.

At the present state, the best prospects for studying long-term evolution are offered by *adaptive dynamics*. It is based on the assumption that replication is only almost exact, and that occasional mistakes – mutations – occur so rarely that the fate of one mutation (its extinction or fixation under selection) is settled before the next mutation occurs (Metz et al, 1996). The population is thereby assumed to consist of one type only, which can be substituted by another type etc. This describes a dynamics in trait space which seems utterly remote from the description of population frequencies given by replicator dynamics but which, in important cases, reduces to it. In particular, if the trait space is a simplex (for instance, probabilities for certain types of behaviour) with a suitable Riemannian metric, one obtains (7) again. But this should not obscure the fact that replicator dynamics and adaptive dynamics adress fundamentally different processes operating on distinct time-scales. One describes short-term evolution – the population dynamics of the frequencies of a given set of genes, or traits; the other describes long-term evolution, the repeated introduction of new mutations (Eshel, 1996).

If the invader's reproductive success is a linear function of its trait, then an ESS is locally stable for each adaptive dynamics; but for many examples, this assumption does not hold, and the evolution in trait space may well lead *away* from an ESS.

Often, two players engaged in a biological 'game' belong to different populations, with different sets of strategies. Most of the previous results carry over to such two-role games, but the general tendency is that there is still less stability: for instance, no mixed strategy can belong to an ESS; there exists an incompressible volume form; heteroclinic cycles become more frequent, etc.

The interacting populations can be different species – for instance, predators

and their prey – and in this case adaptive dynamics leads to models of *co-evolution*. A typical question in this context is whether co-evolution may lead to interaction parameters such that the population numbers oscillate chaotically – a question on which the jury is still out. The interacting populations can also belong to the same species: males and females have conflicting interests about their amount of parental investment, owners and intruders about territorial issues, etc. In that case, role-specific strategies are likely to evolve, for example 'if owner, be prepared to fight to the end; if intruder, avoid escalation' or 'if male, try to desert your partner and fertilise another female; if female, insist on a long engagement period that commits your partner to a monogamous mating season'.

Before turning to some applications, it should be emphasised again that essential aspects can change completely if supplementary effects are included, for instance spatial distribution (Levin and Durett, Hassell et al 1991, Takeuchi, 1996), genetic or physiological heterogeneity (Dushoff and Levin, 1995), stochastic fluctuations (Durrett, 1991) or time lags (Gopalsamy, 1992).

8 Population Dynamics of Infectious Diseases

Applications of mathematical modelling to epidemiology, immunology and virology are of increasing biomedical relevance. They help to understand the course of infectious diseases both within organisms and within populations, and suggest guidelines for treatment and vaccination.

Within a population, the interactions of infected, susceptible and immune organisms lead to endemic or epidemic spread of the disease. In a commonly used epidemiological model (Anderson and May, 1992), if frequencies of uninfected and infected hosts are denoted by x and y, this becomes

$$\dot{x} = k - dx + cy - \beta xy \tag{13}$$

$$\dot{y} = y(\beta x - d - v - c) \tag{14}$$

where k is a constant birth (or immigration) term, d the mortality of uninfected, v the extra mortality due to the infection, i.e. the *virulence*, and c the rate of recovery (which in this simple model does not confer immunity). The model assumes that new infections occur through random contacts between infected and susceptibles. An infection can only spread if the frequency x of uninfected exceeds $(d+c+v)/\beta$. This threshold principle, a cornerstone of epidemiology, holds for most of the variants of the model (including immunity, other transmission mechanisms, periodic oscillations in susceptiblity, other birth and death rates, etc). For many diseases, one has to consider several classes of hosts (different risk groups, for instance, in the case of AIDS). Some of these extensions lead to chaotic dynamics (Grenfell and Dobson, 1995, Olsen and Schaffer, 1990).

Infections are caused by pathogens (virus, bacteria, protozoa), which can all be subsumed as parasites. In (13-14), the pathogen can invade only if the disease-free equilibrium $x = k/d$ is not saturated, i.e. if the *basic reproductive rate*

$$R_0 = \frac{k\beta}{d(d+v+c)} \tag{15}$$

(the number of secondary infections produced by an infected in a population of susceptibles) exceeds 1.

The population dynamics of disease-carrying parasites, and their impact on the population dynamics of the host, is an area of rapid growth. Even the simplest models display oscillations (which for discrete time dynamics can be chaotic). The relation between parasites and their host resembles that between predators and prey, of course: parasites can mediate permanent co-existence between competing strains of hosts, etc. Heteroclinic cycles are likely to occur, for instance when two strains of a host engaged in a bistable competition are beset by two suitably specialised strains of parasites: a resident population of host 1 can be invaded by parasite 1, the resulting equilibrium can be invaded by host 2 (eliminating hosts and parasites of type 1), which in turn allows parasite 2 to invade, etc.

The dynamics described so far deal with the course of an infection *within a population*. Its development *within an individual host* is no less dramatic, and constitutes a new chapter in biomathematics, dealing with the population dynamics and evolution of the 'biosphere' beneath the skin of the host organism. These ecological systems are ideally suited for modelling, since they involve huge populations and short generations, and are subject of intensive clinical tests.

HIV offers the most studied example. As is well-known, the full-blown symptoms of AIDS develop only after a latency period of some ten years. But this quietness is misleading. Clinial tests based on simple dynamical models have revealed a fierce battle between the virus and the immune system of the HIV-infected patient. The average rate of HIV production exceeds 10^{10} particles a day. Free virus particles are cleared within a few hours. Virus infected cells live on average two days.

HIV needs human cells (the 'target cells') to reproduce. In doing so, it kills these cells. Hence virus and target cells interact in much the same way as predators and their prey. But HIV is not only a predator, it is also a prey. The immune system contains a vast repertoire of possible responses (different types of antibodies, killer cells, etc), whose production is stimulated by specific pathogens. The immune responses attack and destroy the pathogens. Thus killer cells and virus also interact like predators and prey. Much clinical research has recently gone into finding out which role – prey or predator – has more relevance for HIV dynamics? At present, it appears that target cell limitation and immune control are of the same magnitude. This leads to prey-predator-superpredator systems which, as known from ecology (e.g. Hastings and Powell, 1991), exhibit complex dynamics. In our case, the simplest model reduces to

$$\begin{aligned} \dot{x} &= k - dx - \beta xv \\ \dot{y} &= \beta xv - ay - pyz \\ \dot{v} &= ry - sv \\ \dot{z} &= cyz - bz \end{aligned} \tag{16}$$

Here x (resp. y) are the frequencies of uninfected (resp. infected) cells, v that of free virus particles and z the abundance of the killer cells produced by the immune response (Nowak and Bangham 1996, DeBoer and Perelson 1998). There is a minimum threshold of infected cells to activate an immune response ($y > b/c$). The frequencies oscillate around an equilibrium value which can be stable or unstable,

i.e. subject to a Hopf bifurcation. The model shows that increasing the responsiveness c of the immune system decreases the abundance y of infected cells, but not necessarily the density z of the killer cells; in other words, there is no simple correlation between virus load and the magnitude of the immune response.

9 The Evolution of Virulence

Most pathogens evolve very quickly, due to their short generation time, their high mutation rate and the intensive selection pressure acting on them. HIV, for instance, spends on average 1500 generations within the body of a patient. During this time, its genetic diversity increases relentlessly, due to copying errors, so that the immune system is faced with ever new challenges.

Mathematical models of the interaction between virus replication and immune response led to completely new interpretations of disease progression in HIV infection (Nowak, 1991). HIV evolution can shift the steady state within an infected individual, and lead to escape from immune responses. Such immune responses are triggered by specific parts (so-called *epitopes*) of the virus. In the simplest model, the virus has two epitopes with two variants each, yielding an eight-dimensional predator-prey equation:

$$\begin{aligned} \dot{v}_{ij} &= v_{ij}(r_{ij} - x_i - y_j) \\ \dot{x}_i &= x_i[c_i(v_{i1} + v_{i2}) - b] \\ \dot{y}_j &= y_j[k_j(v_{1j} + v_{2j}) - b] \end{aligned} \tag{17}$$

where v_{ij} is the concentration of the virus with sequence i at the first and j at the second epitope ($1 \leq i, j \leq 2$), and x_i and y_j are the concentrations of antibodies directed at sequence i of the first resp. j of the second epitope. Generically, one or two of the four viral species and the same number of antibody species have to vanish, and the remaining densities oscillate (Nowak et al, 1995). A homogenous virus population induces an 'immunodominant' response against a single epitope, but a new variant at this epitope can cause the immune response to shift to the other epitope. Heterogenous virus population stimulate complicated fluctuating responses.

This dynamic picture of HIV infection was confirmed by detailed analysis of virus decay slopes in drug treated patients. Again mathematical models were at the core of this newly developing demography of virus infection.

The extreme mutability of HIV explains also why drug-resistant forms emerge so rapidly. Resistance against combinations of drugs requires several mutations. Mathematical models help in devising optimum treatment schedules based on combination therapy (Bonhoeffer et al, 1997).

This is one chapter of a 'Darwinian medicine' grounded in evolutionary biology. In this domain, the evolution of *virulence* (i.e. the parasite-induced mortality of host organisms) is of particular importance (Levin and Pimentel, 1980, Frank, 1995). Pathogens use the bodies of their hosts both as resource and as vehicle. Textbook knowledge presumed that parasites would always evolve towards decreased virulence, since it is better to milk the host rather than butcher it. If parasites become too virulent, they face extinction by depleting their reservoir of susceptibles. It was concluded that successful parasites all become benign. The most impressive example

of such an evolution towards harmlessness is the myxoma virus, released in Australia to kill rabbits: within few years, the death rate of infected rabbits dropped from more than 99 percent to less than 25 percent. Similar trends have been observed in many human diseases. Adaptive dynamics shows that evolution can actually turn parasites into mutualists necessary for the survival of their hosts (Law and Dieckmann, 1998).

But not all parasites become harmless. Selection for a higher basic reproductive rate R_0 often leads to conflicting demands on infectivity and long-term exploitation. If in (15), for instance, the virulence v is an increasing function of the transmission rate β, then R_0 need not necessarily decrease in v. And in the case of super-infection, i.e. when several strains compete *within a host*, selection on parasites does not optimise R_0. Roughly speaking, more virulent strains will have a selective advantage in the intra-host competition, and less virulent in the inter-host competition. Parasites face a so-called *tragedy of the commons*: the need to outgrow their rivals forces them to over-exploit the host, thus possibly driving their common resource to extinction. Game theoretical arguments help in analysing such situations. In general, there will be no evolutionarily stable strain (Nowak and May, 1996).

Of particular interest is the adaptive dynamics of viral particles which can spread either by horizontal transmission, i.e. by infecting of new hosts, or by vertical transmission, in the form of *provirus* integrated into the host's genome. Even if we assume that two strains cannot co-exist within one host (no superinfection), they can coexist within the population if one is favoured by vertical and the other by horizontal transmission (Lipsitch et al, 1996).

10 From the Red Queen to the Major Transitions

The parasite's ecology is further complicated by countermeasures of the hosts which tend to reduce virulence. Due to their short generation time, parasites can quickly adapt to prevailing host defenses, but sexual reproduction allows host organisms to recombine their genes and thus to present shifting targets to the pathogens trying to enter the cells. Many evolutionary biologists view this as the main reason for the prevalence of sexual reproduction (Hamilton, 1988). Indeed, the host faces a peculiar problem of frequency dependent selection. Gene combinations for successful immune systems tend to spread, but if they become too widespread, they cannot remain successful, since parasites will adapt. Sexual host species keep reshuffling their gene combinations, thus providing them with the advantage of being rare.

This is the so-called *Red Queen* theory of sex, named after a figure from the sequel of *Alice in Wonderland* in whose realm 'you have to run with all your speed' just to stay in place – a familiar feature in co-evolution. A species can never stop adapting since the other species do not stop either. Mathematical models for the resulting arms races display a profusion of limit cycles, irregular oscillations and heteroclinic attractors.

The Red Queen metapher makes evolution look like a treadmill rather than a ladder to progress. Nevertheless, evolution has come up with increasingly complex structures, through a sequence of *major transitions* (Maynard Smith and Szathmary, 1996). Cell differentiation, immune systems, or neural networks are examples

of breakthrough inventions. So are the genetic code (translating polynucleotide instructions into polypeptide machinery), the diploid genome (working with two homologous sets of genes), or sexual reproduction (producing offspring by recombining genes from two parents).

Understanding these major transitions necessarily requires thought experiments and mathematical modelling. A major issue for evolutionary biology is sex – a cooperative activity causing an endless series of conflicts. In the wake of the primary question – why should an organism transmit only half of its genes to its offspring? – many other problems surface: Why do sexually reproducing species have two sexes, rather than three, or one? Why are their roles asymmetrical (males producing tiny sperm cells and females large egg cells)? Why is the sex ratio close to one? Why are males fighting for females, and why are females choosier than males? And, since this is biology: why are there exceptions to all these rules? All these questions have been adressed by evolutionary game theory (see e.g. Hutson and Law, 1993, Karlin and Lessard, 1986, or Ywasa and Sasaki, 1987).

Some of the major transitions in evolution led to new levels of organisation, for instance self-replicating molecules, chromosomes, cells, multi-cellular organisms, colonies and societies. In most cases, this emergence of nested hierarchies was due to the fusion of formerly independent units into entities of higher order. These remain threatened by exploitation through mutinies of 'selfish' elements improving their own propagation at a heavy cost to the entire organism. Cancer cells grow without restraint; within a genome, so-called 'outlaw genes' subvert the segregation of chromosomes in a cell division; etc. Each such instance of cooperation is riddled with internal conflicts.

Selfishness may have been an issue since the dawn of life, when several types of self-replicating RNA molecules must have 'ganged up' in order to code for chemical functions. How could they co-exist? As one possible solution, Eigen and Schuster (1975) suggested the 'hypercycle', a closed feedback loop of chemical kinetics, with RNA of type M_i catalysing the replication of RNA of type M_{i+1} (counting the indices i mod n). The equation for the relative densities x_i of M_i is given by the replicator equation (7), with $f_i = x_{i-1}F_i(\mathbf{x})$ and $F_i > 0$ for all i. (If the F_i are constants, there exists a globally stable fixed point in $intS_n$ for $n < 5$, and a stable periodic orbit for $n \geq 5$, see Hofbauer et al, 1991.) This dynamics is always permanent, so that hypercyclic coupling does indeed guarantee the coexistence of all RNA types. But if there occurs an RNA type M which profits from M_i more than M_{i+1} does, then M will displace M_{i+1}, even if it confers no catalytic benefits to the other RNA; such a molecular parasite destroys the whole cycle.

11 The Evolution of Cooperation

Evolutionary history began with molecular networks and led to tightly-knit societies acting as coherently as single organisms do. Bee hives and termite states furnish striking examples. Their extraordinary degree of cooperation is due to the close kinship between all members of a society: a gene for helping one's sister is helping copies of itself. The close relatedness within a bee hive is due to the fact that only very few of its members reproduce. This type of cooperation can be explained by

kinship theory. It is based on the rule that an altruistic act costing c to the donor (in terms of reproductive success) and benefitting b to the recipient has a selective advantage if the relatedness between donor and recipient exceeds the cost-to-benefit ratio c/b.

In human societies, kinship accounts only for a small part of the cooperation: the larger part is due to economic rather than genetic factors. The simplest mechanism is *direct reciprocation*: as long as $c < b$ it pays to help others if they will return the help. This creates new opportunities for parasitism, by not returning help. Game theory provides a ready-made model succintely capturing this aspect. The *Prisoner's Dilemma* (PD) is a symmetric game between two players who can opt between the moves **C** (to cooperate) and **D** (to defect). The payoff matrix is

$$\begin{array}{c} \\ \mathbf{C} \\ \mathbf{D} \end{array} \begin{array}{c} \begin{array}{cc} \mathbf{C} & \mathbf{D} \end{array} \\ \begin{pmatrix} R & S \\ T & P \end{pmatrix} \end{array} \tag{18}$$

with

$$T > R > P > S \quad \text{and} \quad 2R > T + S \tag{19}$$

(the first condition means that the *reward* R for mutual cooperation is larger than the *punishment* P for mutual defection, but that the *temptation* T for unilateral defection is still larger, and the *sucker's payoff* S for being exploited ranks lowest. In our case, $R = b - c$, $P = 0$, $T = b$ and $S = -c$). Obviously, it is best to play **D**, no matter what the other is doing.

This changes if we assume that there is always a probability w for a further round, which is larger than $(T - R)/(T - P) = c/b$. The *iterated* PD game has a random number of rounds with mean $(1 - w)^{-1}$ and admits a huge set of strategies. This model led to a vast amount of investigations, often based on computer tournaments simulating populations of players meeting randomly and engaging in an iterated PD game. In Axelrod's first tournaments, (see Axelrod and Hamilton, 1981) the Tit For Tat strategy TFT (play **C** in the first round and from then on repeat the co-player's previous move) performed extremely well, despite its simplicity. But TFT is not evolutionarily stable (Selten and Hammerstein, 1984): indeed, the strategy of always cooperating can spread by neutral drift in a population of TFT players, and defectors can subsequently invade. Moreover, errors between TFT players lead to costly runs of alternating defections.

To analyse the iterated PD under *noise* (i.e. with a small probability of genetic or strategic errors), let us first consider *memory-one strategies* only. Such strategies are given by the probability to play **C** in the first round, and a quadruple $\mathbf{p} = (p_R, p_S, p_T, p_P)$, where p_i denotes the player's propensity for move **C** after having experienced outcome $i \in \{R, S, T, P\}$ in the previous round. Due to ocasional mistakes, the initial move plays almost no role in long interactions (w close to 1). The dynamics becomes extremely complex: for instance, restriction to the following four strategies leads to a heteroclinic network as attractor (see Fig. 3b): (1) Tit For Tat $(1, 0, 1, 0)$, (2) the more tolerant Firm But Fair $(1, 0, 1, 1)$ which forgives an opponent's defection if it was matched by an own defection, (3) the parasitic Bully $(0, 0, 0, 1)$ which cooperates only after punishment and (4) the strategy $(0, 0, 0, 0)$ which always defects.

But if we introduce occasional mutants, then long-term evolution leads (for $2R > T + P$, i.e. $b > 2c$) to the so-called *Pavlov* strategy $(1, 0, 0, 1)$ which coperates only if the co-player, in the previous round, acted like oneself (Nowak and Sigmund, 1993). This strategy embodies the simplest learning rule, called 'win-stay, lose-shift' by experimental psychologists. It consists in repeating the previous move if the payoff was high (R or T) and in switching to the other option if it was low (P or S). Pavlov players cooperate with each other; an erroneous defection leads in the next round to both players defecting, and then to a resumption of mutual cooperation. Furthermore, Pavlov populations cannot be invaded by other strategies, and in particular not by indiscriminate cooperators who, as mentioned before, pave the way for defectors. On the other hand, Pavlov cannot invade a strategy of defectors: this needs a small cluster of strongly retaliatory strategies like TFT, who eliminate unconditional defectors and then yield to Pavlov.

What about strategies with longer memory, or yet more general *finite-state automata*? Such strategies are defined by a finite set Ω of inner states, some (possibly stochastic) rule specifying which move to play when in state ω, and a rule specifying the transition to the next state as a function of the current state and of the outcome of the previous round (R, S, T or P). Together with the initial state, this defines a strategy for the iterated PD (Leimar, 1997). An example is given by the following table.

	R	S	T	P
1	1	2	3	1
2	1	2	1	2
3	1	1	3	3

It is easy to check that this example satisfies a variant of evolutionary stability: against a co-player using that strategy, it is best, at every stage of the game, to follow the same strategy. This defines a *social norm*. There are many such norms (including Pavlov, if $2R > T + P$), and it is not easy to decide which will get selected. But this example seems particularly successful, and it has an intuitive appeal, if we interpret state 2 as 'provoked' and state 3 as 'contrite': indeed, an erroneous defection by one player makes that player feel contrite, and the co-player provoked: the retaliation redresses the balance. Such inner states correspond to emotions, which are increasingly seen as tools for handling the complexities of social life.

12 Indirect Reciprocity

Obviously, the iterated PD captures only a part of the cooperative interactions in human societies. There is another, *indirect* reciprocity, whereby an altruistic act is returned, not by the recipient (as with direct reciprocity), but by someone else. Indirect reciprocity involves reputation. A simple model assumes that a *score* is attached to each player, which increases (or decreases) whenever the player provides (or witholds) help. Players help whenever the score of the potential recipient exceeds some threshold. This threshold is subject to selection. Punishing a low-scorer is costly, as it decreases one's own score; but if defectors are not punished, they take

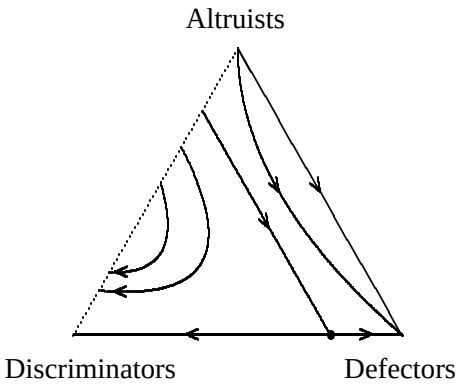

Figure 4: Dynamics of indirect reciprocity

over. Assuming that each player is engaged in a few rounds, both as potential donor and recipient (but never meeting the same co-player twice), one finds that mutation-selection chronicles lead toward cooperation, provided players know their co-players' score sufficiently well (Nowak and Sigmund, 1998).

Occasionally, waves of defection sweep through the population: they are provoked by an excessive frequency of indiscriminate altruists (who are too ready to help low-scorers). Cooperation is more robust if the society is challenged more frequently by invasion attempts of defectors (an intriguing parallel to immune systems). This can be nicely captured by an even simpler model involving only three types of players, with frequencies x_1, x_2, x_3, namely (a) indiscriminate altruists, (b) defectors, and (c) discriminate altruists who help except if the co-player witheld help. If we assume two rounds per player, for instance, both as a donor and as a recipient, the payoffs are

$$\begin{aligned} f_1 &= 2(b - c - bx_2) \\ f_2 &= 2bx_1 + bx_3 \\ f_3 &= 2(b - c) - cx_2. \end{aligned} \tag{20}$$

If discriminating altruists are too rare, i.e. if $x_3 < c/(b-c)$, defectors take over. But all orbits with $x_3 > c/(b-c)$ lead from the edge $x_2 = 0$ (no defectors) back to itself. A mixture of altruists gets established. We may expect that random drift makes the state *fluctuate* along this edge, which consists of fixed points only, and that occasionally, mutation introduces a small quantity x_2 of defectors. What happens? If $x_3 > 2c/b$, defectors cannot invade. If

$$\frac{2c}{b} > x_3 > \frac{c}{b-c}, \tag{21}$$

the invading defectors thrive at first, but are subsequently eliminated by discriminating altruists. After such an abortive invasion, the ratio of discriminators to indiscriminate altruists is so large that defectors can no longer invade. Only when random fluctuations cross the interval given by (21), will defectors take over. But this takes time. If defectors try too often to invade, they will not succeed (see Fig. 4). Imperfect as they are, such models show how cooperation emerges through

the selection of learning rules, moralistic emotions, social norms and reputation. Thus evolutionary models explain the ceaseless give and take prevailing in human societies, and lead game theory back towards its original economic motivation.

References

Anderson, A.M. and May, R.M. (1991) *Infectious Diseases of Humans: Dynamics and Control*, Oxford Univ. Press, Oxford.

Ashwin, P., Buescu, J. and Stewart, I.N. (1996) From attractor to chaotic saddle: a tale of transverse instability. *Nonlinearity* **9**, 703-37.

Axelrod, R. and Hamilton, W.D. (1981) The evolution of cooperation, *Science* **211**, 1390-6.

Bonhoeffer, S., May, R.M., Shaw, G.M. and Nowak, M.A. (1997) Virus dynamics and drug theraphy, *Proc. Nat. Acad. Sci. USA* **94**, 6971-6.

Chawanya, T. (1995) A new type of irregulat motion in a class of game dynamics systems. *Progress Theor. Phys.* **94**, 163-79.

Cressman, R. (1992) *The stability concept of evolutionary game theory*, Springer, Berlin.

DeBoer, R.J. and Perelson, A.S. (1998) Target cell limited and immune control models of HIV infection: a comparison, *Journ. Theor. Biol.* **190**, 201-14.

Durrett, R. (1991) Stochastic models of growth and competition, *Proc. Int. Cong. Math. Kyoto, Vol. II*, 1049–1056.

Dushoff, J. and Levin, S.A. (1995) The effects of population heterogeneity on disease invasion, *J. Math. Biosci.* **128**, 25-40.

Eigen, M. and Schuster, P. (1979) *The hypercycle: a principle of natural self-organisation.* Springer, Berlin-Heidelberg.

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.

Ferriere, R. and Gatto, M. (1995) Lyapunov exponents and the mathematics of invasion in oscillatory or chaotic populations, *Theor. Pop. Biol.* **48**, 126-71.

Fisher, R.A. (1918) The correlation between relatives on the supposition of Mendelian inheritance *Trans. Roy. Soc. Edinburgh* **52**, 399-433.

Forrest, S. (1993) Genetic algorithms: Principles of natural selection applied to computation, *Science* **261**, 872-9.

Frank, S.A. (1996) Models of parasite virulence, *Quart. Rev. Biol.* **71**, 37-78.

Gilpin, M.E. (1979) Spiral chaos in a predator-prey model, *Amer. Nat.* **113**, 306-8.

Gopalsamy, K. (1992) *Stability and Oscillations in Delay Differential Equations of Population Dynamics*, Dordrecht, Kluwer.

Grenfell, B.T. and Dobson, A.P. (1995) *Ecology of Infectious Diseases in Natural Populations*, Cambridge Univ. Press, Cambridge.

Hassell, M.P., Comins, H. and May, R.M. (1991) Spatial structures and chaos in insect population dynamics, *Nature* **353**, 252-258.

Hastings. A. and Powell, T. (1991) Chaos in a three-species food chain, *Ecology* **72**, 896-930.

Hirsch, M.W. (1988) Systems of differential equations which are competitive or cooperative III: Competing species. *Nonlinearity* **1**, 51-71

Hofbauer, J. (1998) Invasion, permanence and heteroclinic cycles. To appear.

Hofbauer, J., Mallet-Paret, J. and Smith, H.L. (1991), Stable periodic solutions for the hypercycle system, *J. Dynamics and Diff. Equs.* **3**, 423-36.

J. Hofbauer and Sigmund, K. (1998) *Evolutionary games and population dynamics*, Cambridge UP.

Hofbauer, J. (1981) On the occurence of limit cycles in the Lotka-Volterra equation, *Nonlinear Analysis* **5**, 1003-7.

Hutson, V. and Law, R. (1993) Four steps to two sexes, *Proc. Roy. Soc. London B* **253**, 43-51.

Hutson, V. and Schmitt, K. (1992) Permanence and the dynamics of biological systems, *Math Biosci.* **111**, 1-71.

Iwasa, Y. and Sasaki, A. (1987) Evolution of the number of sexes, *Evolution* **41**, 49-65.

Karlin, S. and Lessard, S. (1986) *Sex Ratio Evolution*, Monographs in Population Biology **22**, Princeton UP.

Kolmogorov, A.N. (1937) (with I.G. Petrovskij and N.S. Piskunov) Studies of the diffusion equation, combined with increase in the amount of matter and its application to a problem in biology, *Bul. Mosk. Gos. Univ. Mat. Mekh.* **1**, 1-26.

Law, R. and Blackford, J.C. (1992) Self-assembling food webs: a global viewpoint of co-existence of species in Lotka-Volterra communities *Ecology* **73**, 567-578.

Law, R. and Dieckmann, U. (1998) Symbiosis through exploitation and the merger of lineages in evolution, *Proc. Roy. Soc. London B* **265**, 1245-1253.

Leimar, O. (1997) Repeated games: a state space approach, *Journ. Theor. Biol.* **184**, 471-98.

Levin, S.A. and Pimentel, D. (1981) Selection of intermediate rates of increase in parasite-host systems *Amer. Nat.* **117** 308-1.

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.

Lipsitch, M.S., Siller, S. and Nowak, M.A. (1996) The evolution of virulence in pathogens with vertical and horizontal transmission, *Evolution* **50**, 1729-41.

May, R.M. (1976) Simple mathematical models with very complicated dynamics, *Nature* **261**, 459-467.

May, R.M. and Leonard, W.J. (1975) Nonlinear aspects of competition between three species, *SIAM J. Appl. Math.* **29** 243-253.

Metz, J.A.J., Geritz, S.A.H., Meszena, G., Jacobs, F.J.A. and van Heerwarden, J.S. (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.

Mylius, S.D., Doebeli, M. and Diekmann, O. (1998) Can initial invasion dynamics correctly predict phenotypic substitution? To appear.

Nowak, M.A., Anderson, R.M., McLean, A.R., Wolfs, T., Goudsmit, J. and May, R.M. (1991) Antigenic diversity thresholds and the development of AIDS, *Science* **254**, 963-9.

Nowak, M.A. and May, R.M. (1994) Superinfection and the evolution of parasitic virulence, *Proc. R. Soc. London B* **255**, 81-89.

Nowak, M.A. and Bangham, C.R. (1996) Population dynamics of immune responses to persistent viruses, *Science* **272**, 74-9.

Nowak, M.A. and Sigmund, K. (1993) Win-stay, lose-shift outperforms tit-for-tat, *Nature* **364**, 56-8.

Nowak, M.A. and Sigmund, K. (1998) Evolution of indirect reciprocity by image scoring, *Nature* **393**, 573-577.

Nowak, M.A., May, R.M. and Sigmund, K. (1995) Immune response against multiple epitopes, *Jour. Theor. Biol.* **175**, 325-53.

Maynard Smith, J. (1982) *Evolution and the Theory of Games*, Cambridge UP.

Maynard Smith, J. and Szathmary, E. (1995) *The Major Transitions in Evolution*, Freeman, Oxford.

Murray, J. (1990) *Mathematical Biology*, Springer, Heidelberg.

Olsen, L.F. and Schaffer, W.M. (1990) Chaos in infections? *Science* **249**, 499.

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.

Schreiber, S. (1998) On the stabilizing effect of predators on founder-controlled communities, to appear in *Canad. Appl. Math. Quarterly.*

Selten, R. and Hammerstein, P. (1984) Gaps in Harley's argument on evolutionarily stable learning rules and in the logic of tit for tat, *Behavioural and Brains Sciences* **7**, 115-6

Takeuchi, Y. (1996) *Global dynamical properties of Lotka-Volterra systems*, World Scientific, Singapore.

van den Driessche, P. and Zeeman, M.L. (1998) Three dimensional competitive Lotka-Volterra systems with no periodic orbits, *SIAM J. Appl. Math* **58**, 227-234.