

The background of the entire page is a repeating pattern of small, teal-colored, eight-pointed star or snowflake shapes. These shapes are arranged in a regular grid across the white background.

Evolutionary games and population dynamics

Josef Hofbauer, Karl Sigmund.

Evolutionary Games and Population Dynamics

Josef Hofbauer, Karl Sigmund

江苏工业学院图书馆
藏书章



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge CB2 1RP, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, United Kingdom
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Cambridge University Press 1998

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without the
written permission of Cambridge University Press.

First published 1998

Printed in the United Kingdom at the University Press, Cambridge

Typeset in 10pt Monotype Times [TAG]

A catalogue record for this book is available from the British Library

ISBN 0 521 62365 0 hardback
ISBN 0 521 62570 X paperback

Preface

This book replaces our *Theory of Evolution and Dynamical Systems*, which was published in 1988 and has been reprinted several times. It now deserves to be put to rest.

The present text, which is totally restructured and contains a lot of new material, is no longer an interdisciplinary exploration but a tightly organized mathematical textbook on replicator dynamics and Lotka–Volterra equations.

Two important developments during the last decade have made it imperative to write this new book. Within the social sciences, game theory has gained a lot of ground; and within game theory, evolutionary and dynamical aspects have exploded. In our former book, it took us 150 pages of biological motivation to tentatively introduce the notion of a replicator equation. This is no longer warranted today: replicator dynamics is a firmly established subject, and it has grown so tremendously that our old volume definitely looks dated today.

It was exciting for us to see how many mathematical results obtained within the last decade could now be added to the curriculum. Of course, we had to economize elsewhere. This meant that the chapters on ecology, genetics and sociobiology which introduced the biological ideas underlying the mathematical models had to go. No regrets! All these aspects have been covered by one of us, in more readable form and for a general audience, in the Penguin book *Games of Life*. But we have kept to the principle of introducing every new aspect of replicator dynamics by a simple example of basic importance.

The need for a dynamical approach to game theory was felt by John von Neumann and Oskar Morgenstern already. In fact, one can argue that the very term of *moves* of a game suggests *motion* already. The static ‘solutions’ of classical game theory obtained by analysing the behaviour of ‘rational

agents' are fairly unrealistic. ~~Clearly, every form of social life is shaped by what we call~~ Such stepwise adaptation can occur through ~~individual learning or through natural selection~~. It is no coincidence that the single most decisive impetus for evolutionary game theory came from a theoretical biologist, namely John Maynard Smith. This was not the only impetus, of course. With hindsight, it was understood that John Nash (a winner of the Nobel prize in 1994) had a population dynamical setting in mind when he conceived his *equilibrium* notion, and that Reinhard Selten (who shared that prize) had taken, with his principle of the *trembling hand*, an essential step away from the rationality doctrine.

Dynamical models, like the method of fictitious play by Brown and Robinson or the replicator equation introduced by Taylor and Jonker, were originally used as tools for studying equilibria. But during the last few years, it has become increasingly clear that the analysis of equilibria cannot be nearly enough.

This book approaches game theory as a branch of dynamical systems. The first part, accordingly, sets up the dynamical framework by an elementary study of Lotka–Volterra equations, which form the backbone of ecological modelling and are equivalent to replicator equations. The second part, which is the core of the book, offers a systematic introduction to non-cooperative games via replicator dynamics and many other game dynamics. Part three explores the global properties of replicator dynamics, and in particular the notion of permanence. Part four turns to genetics and investigates, in particular, the connection between the strategic approach of evolutionary game theory and the genetic mechanisms of selection, mutation and recombination.

The book is divided into chapters and sections. The theorems and exercises are referred to by their sections, and numbered consecutively. Many results are given in the form of exercises.

We are indebted to E. Akin, E. Amann, A. Arneodo, U. Berger, K. G. Binmore, I. Bomze, P. Brunovský, R. Bürger, C. Cannings, F. B. Christiansen, J. E. Cohen, R. Cressman, P. Couillet, T. Czárán, O. Diekmann, U. Dieckmann, M. Eigen, I. Eshel, H. I. Freedman, B. Garay, A. Gaunersdorfer, S. A. H. Geritz, K. P. Hadeler, W. H. Hamilton, P. Hammerstein, W. G. S. Hines, V. Hutson, V. Jansen, Y. M. Kaniovski, G. Kirlinger, M. Koth, M. Krupa, Y. A. Kuznetsov, R. Law, O. Leimar, S. Lessard, K. Lindgren, M. Lipsitch, D. O. Logofet, V. Losert, R. M. May, J. Maynard Smith, J. A. J. Metz, P. Molander, D. Monderer, T. Nagylaki, J. Nash, S. Nee, M. A. Nowak, G. Parker, M. Plank, M. Posch, D. Rand, K. Ritzberger, K. H. Schlag, P. Schuster, R. Selten, H. L. Smith, J. W. H. So, P. F. Stadler, Y. M. Svirezhev, J. Swinkels, E. Szathmáry, Y. Takeuchi, B. Thomas, J. van

Baalen, E. van Damme, P. van den Driessche, G. T. Vickers, P. Waltman, J. Weibull, F. Weissing, P. Young, E. C. Zeeman and M. L. Zeeman for helpful discussions.

M. Posch deserves our special thanks for producing the figures and helping with T_EXnical questions.

Introduction for game theorists

The decline and fall of the rational player

Evolutionary game theory has been a latecomer in the evolution of game theory. The initial aim of game theorists was to find principles of rational behaviour, by means of thought experiments involving fictitious players who were assumed to know such a theory, and to know that their equally fictitious co-players would use it. At the same time, it was expected that rational behaviour would prove to be optimal against irrational behaviour too. It turned out that this was asking for too much.

The fictitious species of rational players reached a slippery slope when the so-called 'trembling hand' doctrine became common practice among game theorists. According to this eminently sensible approach, a perfect strategy would take into account that the co-player, instead of being a faultless demigod, occasionally does the wrong thing. How often is 'occasionally', one may ask, and what does it matter whether the players lucidly conceived the right move but failed to implement it? From allowing for an infinitesimal margin of error to assuming that the faculties of the players are limited, it takes only a small step; and once the word of 'bounded rationality' went the round, the mystique of rationality collapsed. It was like observing that the Emperor was only boundedly covered by his new clothes.

For game theory, the rout of the rational players proved an unmixed blessing. It opened up a vast realm of applications in the social sciences ranging from ethics to economics, and from political affairs to animal behaviour. As soon as players were no longer constrained to be rational, they could learn, adapt, and evolve.

It became a major task of game theory to describe the dynamical outcome of model games defined by strategies, payoffs, *and* adaptive mechanisms, rather than to prescribe 'solutions' based on a priori reasoning. The simplest

examples showed that it could not be taken for granted that dynamic adaptations would always lead to stationary solutions.

John von Neumann

Let us start with a highly simplified version of Poker. Suppose that there are only two players in the game, say Johnny and Oscar, and two cards, for instance an Ace and a King. At the start, both players pay one dollar into the pot. Then Johnny draws a card, and looks at it. He now has the option of folding (in which case Oscar gets the pot) or of raising the stakes by adding one dollar to the pot. Now it is Oscar's turn. Oscar, who does not know Johnny's card, can either fold (thereby losing his original dollar) or add one further dollar to the pot. In this case, Johnny has to show his card. If it is the Ace, he wins the pot; if it is the King, Oscar does.

Each player has two strategies. Johnny may as well decide in advance what he is going to do, before even looking at his card. He can choose between two strategies, which we call 'bluff' (raise the stakes no matter which card is drawn) and 'no bluff' (raise only if the card is an Ace). Oscar, who has to act only if Johnny raised the stakes, can also decide in advance what to do: he can choose between 'call' (add a dollar if Johnny did) and 'no call'.

The outcome is uncertain, since it depends on the card Johnny draws. But the expected payoff for Johnny is easy to compute in every case. If he chooses 'no bluff' and Oscar chooses 'call', for instance, then Johnny's expected payoff is fifty cents. Indeed, with probability $1/2$, Johnny will draw the Ace, raise the stakes, and win 2 dollars. With probability $1/2$, Johnny will draw the King, fold, and lose 1 dollar. The payoff for Johnny is encapsulated in the following matrix. Note that what one player wins is what the other loses: this is a zero-sum game.

	if Oscar calls	if Oscar does not call
if Johnny bluffs	0	1
if Johnny does not bluff	$1/2$	0

Clearly, if Johnny never bluffs, Oscar should never call. But if Johnny can count on Oscar never calling, he will always bluff. Of course, Oscar should wise up and start always calling. But then, Johnny should stop bluffing altogether, etc. The players will quickly start to use random strategies, and bluff (or call) with certain probabilities only.

Let x be the probability that Johnny bluffs, and y the probability that Oscar calls. Johnny's expected payoff is $x + y/2 - (3/2)xy$. Suppose he chooses

an $x > 1/3$. If Oscar can guess this, he will certainly call, and thus minimize Johnny's payoff to $(1 - x)/2$, a number less than $1/3$. If Johnny bluffs with a probability $x < 1/3$, however, then Oscar, if he can guess this, won't call, thereby minimising Johnny's payoff to x , which again is less than $1/3$. But if Johnny bluffs with a probability exactly equal to $1/3$, then he will have an expected gain of one-third of a dollar, no matter what Oscar does. Thus Johnny has found a way of *maximising his minimal payoff*. Oscar can do just the same. If he chooses $y = 2/3$, he can guarantee that Johnny gains not more than one-third of a dollar, on average. Every other value of y would allow Johnny to get away with more — *if* he manages to guess it. The strategies given by $x = 1/3$ and $y = 2/3$ are maximin strategies — the best if one assumes the worst. But why should one always assume the worst?

Suppose that Johnny, for instance, is a timid person and does not dare to bluff with a probability as high as $x = 1/3$. He will not be penalized for deviating from the right probability, as long as Oscar keeps to his equilibrium value $y = 2/3$. And Oscar will not be penalized either. Of course he would be better off if he switched to calling less frequently. The right reply for Johnny, in that case, is not to meekly reassume his maximin strategy, but to overcome his timidity and switch all the way to always bluffing. Thus the maximin solution seems a rather spurious equilibrium. If Johnny deviates from his maximin strategy, he is not led back to it. Rather it is Oscar who is led to deviate too.

John Nash

The maximin solution is used, not just for this simplified version of Poker, but for all zero-sum games where the gain of one player is the loss of the other. But most games are not zero-sum. Consider the following game, usually called Chicken: Johnny and Oscar have the option to escalate a brawl or to give in. If both give in, they get nothing. If only one player gives in, he pays 1 dollar to the other. But if both escalate the fight, each has an expected loss of 10 dollars, say, for medical treatment. The payoff matrix for Johnny is given by the following matrix:

	if Oscar escalates	if Oscar yields
if Johnny escalates	−10	1
if Johnny yields	−1	0

Clearly, by giving in, Johnny can maximize his minimal payoff. Oscar is in the same position: he also could maximize his minimal payoff by giving

in. But will both players give in? Hardly so. If they guess that the other will give in, they will certainly escalate. But if both escalate, both are worse off.

If x and y are the probabilities that Johnny and Oscar escalate, then the expected payoff for Johnny is $-10xy + x - y$. If Oscar escalates with a probability larger than $1/10$, Johnny should quit. If Oscar escalates with a probability smaller than $1/10$, Johnny should escalate. If both Johnny and Oscar escalate with a probability of exactly $1/10$, they are in a *Nash equilibrium*: neither of them has anything to gain by deviating unilaterally from his equilibrium. But neither Johnny nor Oscar has any reason *not* to deviate from $1/10$, either, as long as the other sticks to $1/10$. Oscar has no reason to care one way or the other, as long as Johnny escalates with a probability of $1/10$. But if Johnny has any reason for believing that Oscar escalates with a higher probability, he should never escalate; and if Oscar suspects that Johnny has such a reason, then he should certainly escalate. What should Oscar do, for instance, if Johnny has escalated twice in the first five rounds? Should he conclude that this was a statistical fluke? Even if it were such a fluke, Johnny could reasonably suspect that Oscar would attribute it to a higher propensity to escalate. Again, the argument for $1/10$ looks rather spurious.

John Maynard Smith

Oddly enough, it was a biologist who offered a convincing explanation. John Maynard Smith, who was studying animal contests at the time, viewed the Chicken game in a *population-dynamical* setting. There were no longer just Johnny and Oscar engaged in the game, but a large number of players meeting randomly in contests where they had to decide whether to escalate or not. It makes a lot of sense, now, to assume that the players escalate with a probability of $1/10$. Indeed, if the overall probability were higher, it would obviously pay to escalate less often, and vice versa. In this sense, *self-regulation* leads to the value of $1/10$ — self-regulation, not between two players, but within a population.

The value $1/10$, then, is an example of an evolutionarily stable strategy. It is the result of a population-dynamical approach which considers questions like: When will the frequencies of certain strategies increase? When will they reach a stable equilibrium?

Enter ecology

Such self-regulation is reminiscent of the self-regulation encountered in population ecology. Interacting species can regulate their population densities.

A scarcity of prey, for instance, will cause a population of predators to dwindle; as a result, the number of prey will increase. This, in turn, will cause the frequency of predators to increase, eventually leading to a decline in the prey, etc. This looks a lot like the up and down of the frequencies of calling and bluffing in Poker.

On second thoughts, it therefore appears considerably less surprising that a biologist would approach game theory with population dynamics in mind. It comes quite natural to naturalists to think of *self-regulation via frequency dependence*. This is a long-established theme among ecologists. Charles Darwin had already been thrilled by this dynamical aspect. He relished working out how, if 'certain insectivorous birds were to increase in Paraguay', a species of flies would decrease; how — since these flies parasitize newborn calves — this decrease would cause cattle to become abundant; which 'would certainly greatly alter the vegetation'; and 'how this again would largely affect the insects; and this again the insectivorous birds ... and so onwards in ever-increasing circles of complexity.' The mathematics underlying this complexity is the theory of dynamical systems.

Evolutionary game theory proved very popular with economists who had not felt too comfortable with the classical approach of analysing the behaviour of unboundedly rational players. It even led to a modest revival of rationality; it turned out more than once that the prescriptions for rational play agreed with the outcomes of game dynamics. And in retrospect, it was even discovered that John Nash had had a population setting in mind when he introduced his equilibrium notion. In his unpublished thesis he wrote 'it is unnecessary to assume that the participants have ... the ability to go through any complex reasoning processes. But the participants are supposed to accumulate empirical information on the various pure strategies at their disposal We assume that there is a population ... of participants ... and that there is a stable average frequency with which a pure strategy is employed by the "average member" of the appropriate population'.

This *mass action interpretation* foreshadows the population dynamical point of view of evolutionary game theory by more than twenty years. Why was it not pursued for such a long time? One possible reason is that evolutionarily stable equilibria do not always exist (whereas Nash equilibria do). It may well be that the adaptation process between players does not converge to a standstill. But the self-regulation of population densities in an ecosystem needn't converge either. The possibility of chaotic population oscillations was understood at about the same time as the ideas of evolutionary game theory started spreading.

One may safely conclude that ecology is the godfather of evolutionary

game theory. We stress this theme throughout our book. The most common models for the dynamics of population numbers in ecosystems (the Lotka–Volterra equations) and the most common models for the dynamics of frequencies of strategies (the replicator equations) are mathematically equivalent. It pays to study them together.

Introduction for biologists

Striking a balance

If we repeated Noah's experiment — starting a new ecosystem with one couple of each species — we would certainly not expect a restoration of the *old régime*. Numbers matter. The fate of a population depends on the frequencies of other populations.

The interdependency of different species can be wonderfully intricate. Darwin relished working out 'how plants and animals, most remote in the scale of nature, are bound together by a web of complex relations', pointing out, as an instance, that bumble-bees are indispensable to the fertilization of heartsease, and that field-mice cause havoc among the nests and combs of bumble-bees. Since the number of mice is largely dependent on the number of cats, it is consequently '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!'

This self-regulation of population frequencies has been a dominant theme of mathematical ecology. It started in the 1920s with Alfred Lotka modelling the cycle of mosquitoes and humans in transmitting malaria, and Vito Volterra analysing the dynamics of predators and prey among fish in the Adriatic. They came up with differential equations describing the dynamics of such systems. But the first generations of mathematical ecologists concentrated mostly on investigating static aspects of ecological communities. Some of their main problems, like the validity of the exclusion principle (when can there be more species than niches?) and the relation between the complexity and the stability of an ecosystem (do species that are more interconnected produce assemblies that are more robust?), were phrased in terms of stability properties of equilibria. Only in the 1970s did the prevalence of irregular oscillations, which had always been known to field workers, filter

down to theoreticians. Ecological models became a major impetus to chaos theory.

Survival vs. equilibrium

For a long time, the efforts of mathematical ecologists to analyze the stability of bio-communities were marred by a misunderstanding. If one simply adopts the stability concepts of physicists or engineers, then one will call an ecosystem stable if population numbers converge to an equilibrium which is promptly re-assumed after every small perturbation. But field ecologists would not expect to find, in the wild, the static, well-controlled state of affairs implied by such a stability notion. For ecologists unspoilt by physics courses, the proverbial lynx-hare cycle, whose undamped oscillations have been recorded for two hundred years, epitomizes stability. They little care whether the population numbers converge, or oscillate in a regular or chaotic fashion. For such ecologists, stability means that population numbers do not vanish; that the species making up the ecosystem do persist. Survival, not equilibrium, is what counts. This second form of stability is one of the main themes of our book.

All biological communities are transient, of course; but some are more so than others. They collapse right away, without having to wait for the construction of an interstate highway, the mutation of a parasitic strain, or a series of harsh winters. They are doomed from the start: they are *unsustainable* — that is, impossible in the long run. They do occur in nature, but it is hard to make them out before they are replaced by less fleeting configurations. Ecologists who wish to understand what happened have no time for a leisurely *post mortem*. Yet they must know why communities fail if they want to learn about those which persist. An empirical approach to this question is always difficult and often painful; mathematical models are less risky and, if they are cleverly set up, more revealing.

The history of ecological communities is chronicled in terms of invasions and extinctions. The fate of a population depends mostly on what happens when it is rare. Can an invader spread? Can a population recover after being decimated? Every new adaptation is initially rare, and every lineage has been tested in countless bottlenecks. A population must be able to gain a foothold when in a minority. It must be able to grow in a world which is essentially determined by the others. Only then will its numbers become large enough to affect the environment, the frequencies of the other species, and its own growth. If we translate this into static notions, we find that we have to investigate two types of stability for an equilibrium: *inner stability* (will a

small perturbation of the prevailing distribution be offset by self-regulation?) and *outer stability* (will a new population entering as a minority be able to grow?). Obviously the same two questions apply to dynamic regimes, too: (a) what happens if one adds or removes a few heads of a population currently present in the ecosystem, and (b) what happens if one adds a few heads of a new, intruding species?

Breeding and games

Every growth will of course eventually be checked. Darwin termed this the 'struggle for life', stressing that it often had little to do with animals actually fighting each other. He compared the 'face of nature' to a yielding surface, 'with ten thousand sharp wedges packed close together and driven inwards by incessant blows.' And he emphasized that this struggle of life was almost invariably most severe between individuals and varieties of the same species.

This suggests applying ecological modelling to animal behaviour, since different types of behaviour correspond to different varieties. One of the first to think about behaviour in terms of invasion and self-regulation of frequencies was John Maynard Smith. He applied it originally to actual contests within one species, using it to explain the prevalence of conventional fighting. Stags, for instance, engage in roaring contests, a parallel walk and a pushing match with interlocked antlers. Only rarely do they escalate to an all-out fight likely to have a lethal conclusion. Their usual restraint is obviously good for the species, but this advantage cannot explain it.

John Maynard Smith and John Price couched the contests in game theoretic terms. Stags, in their thought experiment, came in two brands: the 'hawks' ruthlessly escalating every contest, and the 'doves' sticking to conventional displays and fleeing whenever their opponent gets rough. In a population of doves, hawks do well and will spread, since they win all their contests. But in a population of hawks, they have only a fifty per cent chance of winning; it is just as likely that they will end up seriously injured. Doves, who avoid this fearsome risk, will do much better and spread. Hence hawk populations can be invaded by doves and vice versa; the outcome should be a mixture where the frequency p of hawks is inversely proportional to the cost of an injury (a cost expressed in fitness, i.e. reproductive success). Among heavily armed species like stags, this cost is very high, so that escalated conflicts will be rare.

A fighting behaviour corresponds in game theoretical terms to a strategy. Hawks and doves are so-called pure strategies, and each is a best reply to the other. The structure of the struggle is exactly that of a once popular pastime

of American teenagers, which went by the name *Chicken* and attracted the interest of early game theorists. The thing to avoid, in such a game, is to act like your adversary — but you do not know beforehand what he will do. It turns out that it is best to let chance decide — if you know how to weight the dice properly.

Hedging the bets

What about mixed strategies in the hawk-dove game, i.e. behavioural programs telling the player to escalate with a such and such a probability? It is easy to see that there exists one such strategy — namely escalating with the probability p given by the frequency of hawks in a hawk-dove mixture — which is a best reply to itself, a so-called Nash equilibrium. If all stags use this mixed strategy, no stag can expect to do better by escalating with a different probability. Actually, stags that do so will not do worse either. Nevertheless, they cannot invade, and this for a rather subtle reason. The offspring of the mutant (who inherit its propensity for escalation) do as well as the resident population in all contests against the residents; but they do less well than the residents in the (admittedly rare) contests against their own. Hence their type cannot spread: it checks its own increase even when rare.

A similar argument explains the prevalence of the sex ratio 1/2. If the sex ratio (the proportion of males among the offspring) were different, it would pay to produce offspring of the rarer sex. The success of one strategy (produce more sons) or the other (produce more daughters) depends on the mean sex ratio in the population, and hence on the frequencies of the strategies. Success, as always in biology, means reproductive success; hence the successful strategies spread, change the composition in the population, and therefore affect their own success. Again, one can show that investing equally in the production of the two sexes is evolutionarily stable. If everyone does it, any deviation will be self-defeating, because it affects the frequencies of the sexes in the wrong way.

The definition of an evolutionarily stable strategy (if the residents adopt it, no mutant can invade) is based on an implicit dynamics. It is easy to make this dynamics explicit, by assuming that like begets like. This yields the replicator equation describing the evolution of the frequencies of different strategies in a population. This dynamics is not merely a prop to sustain arguments from equilibrium theory. For many games, equilibria alone do not suffice to describe what happens, and a static outcome cannot be expected. The simplest example is the rock-scissors-paper game, where strategy A is

beaten by *B*, which is beaten by *C*, which is beaten by *A*. There exists a mixed equilibrium with so much of *A*, *B* and *C* each. But it may happen that this equilibrium is never reached. Instead, the cyclic succession of populations which are almost entirely composed of *A*, *B* or *C* builds up to an increasingly jerky roundabout of upheavals.

It used to be thought that the rock-scissors-paper game was just a conundrum devised for the amusement of theoreticians, until it was found out that lizards do play it: one of their species, *Uta stansburiana*, has three types of male with different mating strategies (they are conveniently distinguished by their throat colour). Type *A* keeps one female and guards it closely; type *B* keeps several females, and necessarily guards them less closely; and *C* guards no female at all and looks out for sneaky matings with unguarded females. The three types can invade each other cyclically.

Similar ratchets can occur in parasitology. The immune system of the host acts as a combination lock which the parasites try to break. By trial and error, they will eventually succeed. Of course they are usually most efficient in attacking the most common immunotype in their host. For the host, it can be deadly to adopt a combination code which is currently widespread in the population. It pays to belong to a minority. But since it pays in offspring — the Darwinian currency — such a minority will yield a new majority, and come under concentrated attack. Every solution is self-defeating in the long run. This leads to arms races without a finish line.

Short vs. long term

In studying evolutionary chronicles, we are led to consider two time scales. Short-term evolution describes how the frequencies of adaptive traits regulate each other via natural selection, i.e. how the distribution of the types actually present in a population changes from generation to generation. Long-term evolution describes how new types can invade through mutation. Not surprisingly, this leads to rather different dynamics, and hence also to different stability notions. In particular, an evolutionarily stable strategy need not attract. It may well be that a population is invasion-proof in the sense that mutations cannot lead out of it, and at the same time inaccessible, so that mutations will lead any nearby populations further away from it. Such 'Garden of Eden' configurations can never occur as evolutionary outcomes. The widespread idea that a population will somehow evolve until it happily reaches the safe haven of evolutionary stability is not always valid. The dynamics of evolution can be a lot more exciting.

Replicator dynamics and adaptive dynamics describe the short-term and