



International Institute for
Applied Systems Analysis
Schlossplatz 1
A-2361 Laxenburg, Austria

Tel: +43 2236 807 342
Fax: +43 2236 71313
E-mail: publications@iiasa.ac.at
Web: www.iiasa.ac.at

Interim Report

IR-11-015

Thoughts on the geometry of meso-evolution: Collecting mathematical elements for a post-modern synthesis

Johan A.J. Metz (j.a.j.metz@biology.leidenuniv.nl)

Approved by

Ulf Dieckmann
Program Leader, EEP

June 2011

Interim Reports on work of the International Institute for Applied Systems Analysis receive only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work.

Thoughts on the geometry of meso-evolution: collecting mathematical elements for a post- modern synthesis

J.A.J. (= Hans) Metz

to be published in
The Mathematics of Darwins Legacy
edited by
F. A. C. C. Chalub and J. F. Rodrigues
Birkhauser, Basel 2011

Contents

1. Introduction
 - 1.1. On micro-, macro- and meso-evolution
 - 1.2. The so-called modern synthesis
 - 1.3. Reinstating the missing components
2. A short introduction to adaptive dynamics and its ramifications
 - 2.1. Codifying the ecological principles of meso-evolution
 - 2.2. Fitness
 - 2.3. Evolutionarily statics: ESS theory
 - 2.4. Adaptive dynamics I: on traits, PIPs, MIPs and TEPs
 - 2.5. Adaptive dynamics II: the canonical equation
 - 2.6. Links between adaptive dynamics and Evo-Devo
 - 2.7. Adaptive dynamics III: evolutionarily singular strategies
 - 2.8. Adaptive speciation
3. Some meso-evolutionary predictions
 - 3.1. The shortest time scale
 - 3.2. Intermediate time scales
 - 3.3. The longest time scale
4. Macro-evolution
 - 4.1. Evolutionary tinkering, tangled maps, and the need for intermediate abstractions
 - 4.2. Phenotypes (and genotypes)
 - 4.3. Internal selection
 - 4.4. A macro-evolutionary prediction
 - 4.5. An Evo-Devo myth
5. Back to meso-evolution
 - 5.1. Justifying adaptive dynamics
 - 5.2. Allopatric speciation
6. How to get on?

1. Introduction

1.1. On micro-, macro- and meso-evolution

The present chapter is complementary to the one by Warren Ewens. It contains almost no equations or genetics (or rather, both are implicitly there, but stay well hidden), fitnesses are not assumed to be constant, but keep changing (in genetical terms, selection is always seen as frequency dependent), the evolutionary path is considered as being shaped by the repeated substitution of novel mutations (as opposed to gene frequency change), and adaptive landscapes depict the fitness of potential mutant types over a space spanned by traits (instead of the mean fitness of a population over a space of gene frequencies). The reason for this difference in emphasis is that my first training is not as a mathematician or population geneticist but as a naturalist, i.e., field oriented taxonomist-ecologist, which shapes the questions that have my interest. As a consequence, within evolutionary biology I am interested primarily in meso-evolution, defined here as evolutionary changes in the values of traits of representative individuals and concomitant patterns of taxonomic diversification. This in contrast to micro-evolution, a term reserved for the changes in gene frequencies on a population dynamical time scale (the topic of Warren Ewens' chapter), and macro-evolution, a term that then can be reserved for large scale changes like anatomical innovations, where one cannot even speak in terms of a fixed set of traits. Thus meso-evolution acts on a time scale above the micro-evolutionary scale of gene substitutions but below the scale on which the intricacies of the developmental process start to have a large influence.

Meso-evolution is far more than micro-evolution writ large, and a similar statement holds for macro- versus meso-evolution. Each of these levels has its own emergent phenomena, and its own explanatory frameworks, which should in the end be based at least in part on idealised summaries of large scale regularities in the outcome of lower level mechanisms, in a similar manner as pressure and temperature can be treated as macroscopic causes, although they themselves are but statistics of the underlying process of molecular motion. And where the results of thermodynamics are predicated on man-made or naturally evolved macroscopic structures confining these motions, so do trait changes result from the micro-evolutionary process of mutant substitutions taking place against the backdrop of a genetic architecture and developmental system as deliverers of the required mutational variation, internal selection caused by the necessity that the machinery of a body stays in concert (think of Cuvier's famous dictum that given one part he could deduce the rest of an organism), and ecological selection due to the interactions of individuals with their conspecifics, resources, predators, parasites and diseases. In this chapter I focus on these encompassing mechanisms rather than on the motion of gene frequencies.

In order to get a clean story I assume time scale separations all over. Not that I believe that such time scale separations hold good even most of the time. In a rigorous sense they only hold good very rarely. However, it looks

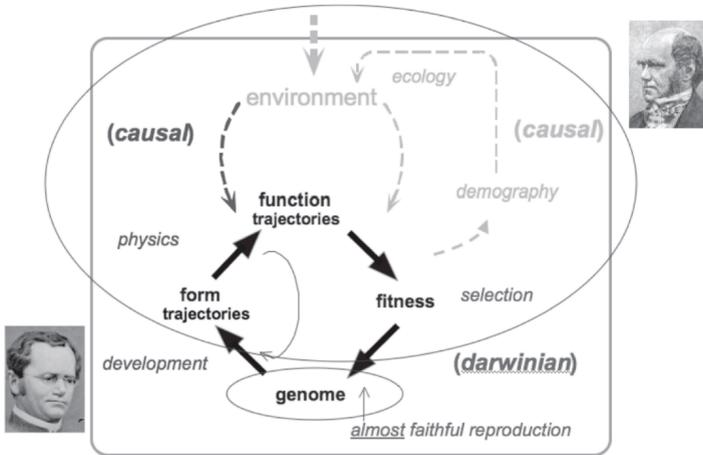


FIGURE 1.1. How the presence of an optimisation principle can be detected from PIPs. Panes a to d are examples of PIPs for models with an optimisation principle. Pane d is not skew symmetric and therefore there is no optimisation principle. Pane e ... In pane f the presence of a rock-scissors-paper trait triple is directly seen to exist. Finally pane h, for a model that has an optimisation principle shows how the transitivity condition can be seen to work through in a direct manner: above resident trait value x_1 there is a certain alternation of plus and minus regions. If we read of the trait values that are selectively neutral relative to x_1 then these trait values should have exactly the same pattern of plus and minus regions above them.

as if arguments based on them lead to fair outcomes in more than a fair fraction of the cases. Moreover, it is only by such time scale arguments that I can easily make the transition from population genetics to the views common among morphologists and taxonomists. As I am aiming at contributing to a postmodern synthesis (c.f. Subsection 1.2), so be it. I have chosen for being wrong in the details, although I believe often close, over being strictly correct but unable to address the larger picture.

1.2. The so-called modern synthesis

Figure 1.1 gives a schematic representation of the “integrated biology” view of life. The small inner loop (fat arrows) can be seen as abstracting a life cycle. The genome produces shapes that change over an individual’s life. Shapes should be interpreted here in a generalised sense, e.g. including the distribution of all sorts of chemicals within a body. What these forms do is called function. What they can do depends on the environment. What they do, together with other aspects of the environment like availability of

resources or density of predators, determines how many descendants they on average contribute to future generations. What they do and how many they are necessarily affects the environment. Their number of descendants if the environment were not to change, measured as the asymptotic average rate of exponential growth, is called fitness (see Subsection 2.2). Only when its fitness is positive a mutant type has a positive probability to invade, and a mutant can only take over if the fitness of its parent type becomes negative when the mutant further grows in numbers. This filtering determines which types will be present in the future. Finally, reproduction needs to be almost faithful in order to get evolution; if it were fully faithful no mutants would occur, if it were too unfaithful, the effects of selection would be swamped by a mutational oozing out over trait space.

For lack of better, Darwin thought of inheritance as a blending of the parental types, which would make evolution impossible (c.f. the chapter by Warren Ewens). Luckily Mendel saved the theory of evolution by natural selection by discovering the faithfully reproducing genes. Seen from first principles the genes are the primary units of evolution, while the usual biological individuals are no more than uneasy coalitions of genes. Luckily, under some simplifying conditions, among which a separation between the time scales of population dynamics and evolution, the Mendelian mechanism fairly often, but certainly not always, allows calculating evolutionary trajectories and outcomes as if it were the individuals that reproduce faithfully (c.f. Subsections 2.2, 2.3 and 2.5).

The synthesis between the Darwinian and Mendelian view was made by the three great theoretical population geneticists: Ronald Fisher, J.B.S. Haldane and Sewall Wright (see the chapter by Warren Ewens). In North America this material reached the less mathematically inclined biologist community through the experimental work of Theodosius Dobzhanski. This let taxonomists like Ernst Mayr and paleontologists like George Gaylord Simpson embrace the so-called population view of evolution. Although espoused by them as the “modern synthesis” (see e.g. [39, 47, 48]), the result was more like a mutual admiration society than a synthesis. The discrepancy between terminology and reality was caused largely by the remaining but blissfully ignored lack of sufficient understanding of both development and population dynamics.

When Mayr grew older, he elevated himself to the role of philosopher and historian of biology, and in this role pushed the simplified picture of integrated biology that you see in Figure 1.2. It was through this ploy that the modern synthesis indeed could be perceived as a synthesis¹. In simple textbook evolutionary scenarios this simplified worldview indeed works remarkably well. However, it unravels when one starts to focus on complicated trait spaces and ecologies as they occur in real life.

¹See [2] for a further discussion, in particular of the political reasons for this purposeful oversimplification.

lower quadrant of Figure 1.1. AD focuses mainly on the shaping of the selective arena by ecological processes, on meso-evolution, and on prediction oriented theory. Its goal is to fill the right upper quadrant of Figure 1.1.

In this chapter I will sketch some mathematical aspects of what directions a postmodern synthesis based on these newer developments might take.

2. A short introduction to adaptive dynamics and its ramifications

2.1. Codifying the ecological principles of meso-evolution

Adaptive dynamics (AD) was developed as an aid for making the transition from micro- to meso-evolution. Meso-evolution proceeds by the selective filtering by the ecology of a continual stream of mutants. AD concentrates on the ecological side of this process, as there are much clearer a priori mathematical structures to be found at that end. The basic theory assumes clonal reproduction, and only a subset of the results extend to the Mendelian case, for monomorphic populations directly (c.f. Subsections 2.3 and 2.5) and for polymorphic populations after appropriate modification (c.f. Subsection 2.8).

One of the immediate consequences of stressing the ecological side of the equation is a strong awareness that fitnesses are not given quantities, but depend both on the traits of an individual and on the environment in which it lives. The ecological feedback loop makes that in the monomorphic and clonal cases necessarily the fitnesses of all types present on an ecological time scale are zero (see Subsection 2.2). Only the fitnesses of potential mutants can be positive or negative. The signs and sizes of these mutant fitnesses determine the direction and speed of evolutionary progress. Evolution corresponds to permanent uphill movement in a fitness landscape that keeps changing so as to keep the resident types exactly at zero. See Figure 2.1.

The main general insight from the mathematical analyses of this picture has been the discovery of a potential mechanism for adaptive speciation that appears with a certain ubiquity in ecological models (see Subsection 2.8). Apart from that, the theory has produced a good number of very effective tools for analysing special families of eco-evolutionary models (Subsections 2.4, 2.5 and 2.7).

Below I will keep exploiting the landscape analogy and refer to zero as sea level, etc. Moreover, to keep the story simple I shall each time initially proceed on the assumption that individuals reproduce clonally.

2.2. Fitness

The ecological perspective. The concept of fitness as a quantitative measure of competitive prowess is a modern invention. Darwin never used the term in this meaning, and neither did population genetics' founding triumvirate (with the exception of [20], otherwise they use terms like selective advantage, see [51]). In population genetics, fitness is generally used for the probability

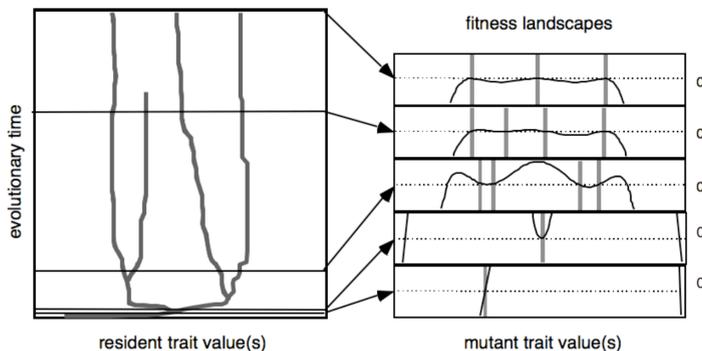


FIGURE 2.1. Left: Evolutionary path simulated on the basis of a population dynamical model, assuming clonal reproduction. Only the traits that are dominantly present in the population are shown. The second ascending branch finishes since the subpopulation under consideration went extinct. Right: The fitness landscapes for five population compositions as these occurred at the indicated times. The vertical bars indicate the types that at that moment were present in the population. At the second selected time the population resided at a branching point (see Subsection 2.8). At the final time the remaining three subpopulations reside at an evolutionarily stable ESC (see Subsection 2.3).

to survive to reproduction. However, this only works for the relatively simple ecological scenarios considered there, where the different life phases are both neatly separated and synchronised. In ecology one has to account for a messier world where populations have age, size, spatial or other structures, and where demographic properties vary with the weather over an individual's life and over the generations.

Let the environment be defined as anything outside an individual that influences its population dynamical behaviour, which by definition consists of impinging on the environment, giving birth, and dying (see e.g. [51, 53, 58, 57]). It is always possible in principle to find a Markovian representation of that behaviour, in terms of a state space, transition probabilities that depend on the course of the environment, and outputs that are either deterministic or occur in a Poisson cluster process with rate and cluster (clutch) size dependent on an individual's state and the condition of the environment at the time. Given the course of the environment, individuals independently move through their state spaces, the population state is a measure over this space, and the expectation of this measure, which is again a measure, moves according to a positive linear evolutionary system. The theory of positive linear systems then tells that generally the expected size of a population in

an ergodic environment will in the long run on average grow or decline exponentially (for details see [51, 19]). This growth rate ρ is what ecologists call fitness. It necessarily is a function of two variables, the type of the individuals Y , parametrised by their traits, and the environment E , to be written as $\rho(Y|E)$. The theory of branching processes moreover tells that when a population is started with a single individual it will, barring some technical conditions, either eventually go extinct or grow exponentially, with the probability of the latter being positive if and only if its fitness is so (see [37, 3, 4, 34]).

In the theory of long-term adaptive evolution one is mainly interested in populations in which the number of individuals exposed to similar environments are sufficiently large that the internal workings of these populations can be modelled in a deterministic manner, with possibly on top an external stochastic driver. In nature populations are necessarily bounded. (Thanks to the above definitions, this can be ascribed to changes in the environment brought about by the growth of those populations.) Hence the population state space is a closed bounded subset of the cone of positive measures over the state space of the individuals, and the state space of a community is the product of the state spaces of the comprising species, plus the state spaces of the dynamics of any inanimate resources. With an infinitesimal amount of noise the states of such communities will approach an “extinction preserving chain attractor” (this is a generalisation that accounts for the fact that extinct populations cannot be resurrected of the concept of Conley-Ruelle or chain attractor; see [38, 33]). With larger amounts of noise the community will in general end up in a stochastic attractor, that is, a stationary distribution of community states. I will throughout assume that this attractor generates an ergodic environment (the exceptions that I have seen constructed all appeared to need biologically pretty exceptional conditions). Let the environment generated by a coalition of clones $C = (X_1, \dots, X_n)$ be written as $E_{\text{attr}}(C)$. A combination of the preceding arguments then leads to the introduction of the *invasion fitness* $\rho(Y|E_{\text{attr}}(C))$ of a new type Y in a C -community.

For ease of exposition I proceed as if $E_{\text{attr}}(C)$ is unique. Most of my statements extend to the general case with only small modifications.

The extension of the previous framework to Mendelian populations turns out to be easier than perhaps expected (although implementing it in concrete cases tends to be horrible). For the community dynamics all one has to do is distinguish individuals according to their genotypes, and incorporate their mating opportunities with different genotypes into the description of the environment (c.f. [11]; this in the case of casual matings, with more extended pair formation it becomes necessary to extend the state space of individuals to keep track of their marriage status). Alleles, of course, reproduce clonally and as such have fitnesses. It is also possible to define a mock fitness of phenotypes by introducing a parallel clonal model with individuals passing through their lives like their Mendelian counterparts and having a reproduction equal to the average of the contributions through the micro- and macro-gametic

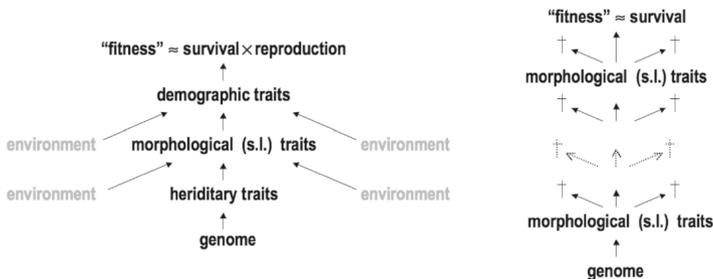


FIGURE 2.2. The Evo-Devo perspective on fitness, with left an ecologically enlightened and right a more narrow perspective. Note that what here is called “fitness” is at best proportional to $\exp(\rho)$.

routes (for humans semen and ova) of those counterparts. With such a definition some essential, but certainly not all, fitness-based deductions for the clonal case go through for Mendelian inheritance. In particular, for genetically homogeneous populations the fitness of a so-called resident, that is, a type that is present in a community dynamical attractor, equals zero (since genetically homogeneous populations breed true and resident populations by definition do not in the long run grow or decline). Moreover, the invasion of a new mutant in a homogeneous population is correctly predicted, as that mutant initially only occurs in heterozygotes that breed true by backcrossing with the homogeneous resident.

The Evo-Devo perspective. The Evo-Devo view of fitness is much closer to the population genetical one. In Evo-Devo people generally only consider whether a developmental program works sufficiently well, that is, produces a functioning organism. The better an organism functions the higher its fitness, with this functioning more often than not being largely independent of the specific environment in which it lives. Moreover, in general a mutant either develops along lines compatible with the environments under consideration or it dies early in its life. Translating these considerations into population dynamics does not necessarily lead to the usual population genetical equations, but it often leads to model formulations supporting a so-called optimisation principle (see [51]). For more refined eco-evolutionary models optimisation principles may be rare, but they occur more readily when one aims for a model that captures only grosser characteristics while glossing over the ecological fine detail. Hence optimality arguments can hold sway in Evo-Devo, where in evolutionary ecology they have been largely replaced by ESS and AD arguments (see Subsections 2.3 to 2.7). Thus, the picture espoused by Evo-Devo researchers tends to be like the ones in Figure 2.2.

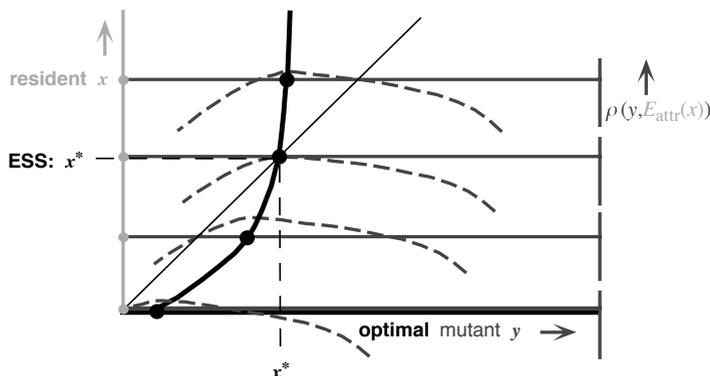


FIGURE 2.3. Scheme for calculating ESSes: For each of the possible resident populations, characterised by a scalar trait, the invasion fitness of all potential mutants is calculated (interrupted curves). The mutant axis is drawn on the same scale as the resident axis. From these fitness curves the optimal strategy for the corresponding resident environment is calculated (fat curve). The ESS is the optimal reply to itself, to be calculated by intersecting the fat curve with the 45° line.

2.3. Evolutionarily statics: ESS theory

Evolution stays put whenever the community produces an environment such that mutants have negative fitness whenever they differ from any of the residents. In the special case where there is but one resident type, we speak of an Evolutionarily Steady Strategy. (The old name Evolutionarily Stable Strategy, introduced by Maynard Smith and Price [46], is a bit of a misnomer, since, as first discovered by Eshel [16] and illustrated in Figure 2.4, ESSes need not be evolutionarily attractive.) In the general case when there may be more than one resident type I will speak of an Evolutionarily Steady Coalition. ESCs are the equilibria of evolution.

One way of calculating ESSes is depicted in Figure 2.3. For each environment as generated by a possible resident the maximum of the invasion fitness landscape $\rho(Y|E_{\text{attr}}(X))$ is calculated. Next one intersects the resulting manifold $Y_{\text{opt}}(X)$ with the linear manifold $Y = X$ to get the ESS $X^* = Y^*$. As necessarily any monomorphic resident has fitness zero, all potential mutants $Y \neq Y^*$ have negative fitness.

The situation for ESCs is a bit more complicated, as there may be so-called genetic constraints. So it may happen, for example, that a trait is under control of a single locus only and at the ESC the heterozygote has a higher fitness than the two homozygotes. The good message is that in the so-called Ideal Free (IF) case, as in the clonal case, all phenotypes comprising an ESC have fitness zero, at least when there is only a single birth state and the

ESC engenders a community dynamical equilibrium³; this IF case is defined by the requirement that there are no genetic constraints whatsoever, that is, mutants can occur that produce any feasible type as heterozygotes in the genetic backgrounds supplied by the resident population. The bad message is that at the present state of knowledge about the genotype to phenotype map there is no way of predicting when genetic constraints may throw a spanner in the works and neither is there an inkling of this becoming feasible in the future.

2.4. Adaptive dynamics I: on traits, PIPs, MIPs and TEPs

Paleontologists and taxonomists are interested in the change of traits on an evolutionary time scale. What are traits to taxonomists are parameters to ecologists. So in AD one is after a dynamics in the parameter space of a community dynamics. The first trick for arriving at such a simple picture is to assume a time scale separation, such that favourable mutants come along singly after a community has relaxed to an attractor. The second trick is to assume clonal reproduction, on the assumption that this way one can find out where the ecology would drive evolution if the latter were not hampered by the constraints of Mendelian genetics⁴.

To get at a purely trait oriented picture, first any reference to the environment should be removed from the expression for invasion fitness:

$$s(Y|C) := \rho(Y|E_{\text{attr}}(C)).$$

(Often this is written as $s_C(Y)$ to emphasize the interpretation as a family of fitness landscapes.)

In this Subsection, I concentrate on scalar traits. I start with the case where there is only a single clonally reproducing resident, $C = x$. The first step in the analysis is plotting a contour plot of $s(y|x)$. Usually this is simplified to plotting only the zero contours, as those are the ones that matter by far the most. The result is customarily called Pairwise Invasibility Plot (PIP). See Figure 2.4.

Note that the diagonal is always a zero contour as residents have fitness zero. The points where some other contour crosses the diagonal are referred to as evolutionarily singular strategies (ess-es). The ESSes are a subset of the ess-es.

Now assume that mutational steps are but small and that in the beginning there is only one resident trait value $x(0)$. Plot this value on the abscissa of the PIP, say the one in the top left panel in Figure 2.4. After some random

³The proof goes by contradiction, see e.g. [56]. It would be most useful if the result could be extended to situations with multiple birth states or non-equilibrium attractors.

⁴The set of ideas described in Subsections 2.4 and 2.7 was first worked out in some detail in [54] with [30] as counterpart geared to a biological rather than a mathematical audience. [10] as well as [54] independently argued that adaptive dynamics should be seen as a limit processes of sequences of stochastic models in which one lets the system size go to infinity and the probability of a mutation per birth event go to zero in such a manner that their product stays bounded, while appropriately rescaling time; a rigorous convergence proof, though thus far only for a particular special case, can be found in [6].

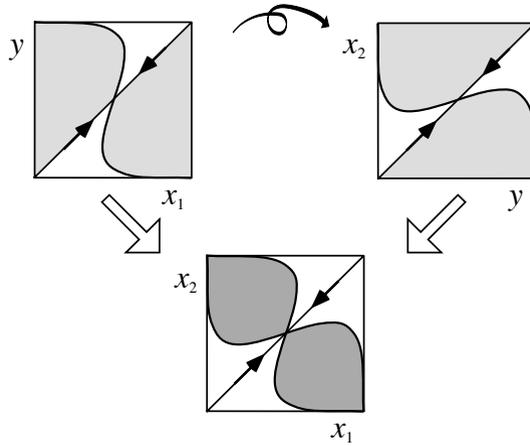


FIGURE 2.5. The construction of a Mutual Invasibility Plot, depicting the set in (trait space)² harbouring protected dimorphisms. Not all polymorphisms occurring in AD are protected, but unprotected polymorphisms have the habit of never lying close to a diagonal (Stefan Geritz, pers com).

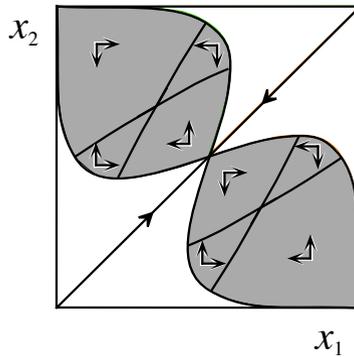


FIGURE 2.6. Trait Evolution Plot, i.e., MIP together with arrows that indicate the direction of the small evolutionary steps that result from the invasion by mutants that differ but little from their progenitor, and adaptive isoclines.

To see how such coexisting pairs of strategies fare it is necessary to consider the set of so-called protected dimorphisms, i.e., pairs of strategies that can mutually invade, to be denoted as (x_1, x_2) . The construction of this set is depicted in the Figure 2.5.

The evolutionary movement of the pair (x_1, x_2) is governed by $s(y|x_1, x_2)$. Under the assumption of small mutational steps a good deal of information

can be extracted from the adaptive isoclines, calculated by setting the selection gradient

$$g_i(x_1, x_2) := \left. \frac{\partial s}{\partial y} (y|x_1, x_2) \right|_{y=x_i}$$

equal to zero. As depicted in Figure 2.6 x_1 will move to the right when g_1 is positive and to the left when it is negative, and x_2 will move up when g_2 is positive and down when it is negative. The likes of Figure 2.6 are customarily referred to as Trait Evolution Plot (TEP).

Subsection 2.7 gives a classification of the possible dynamics near an ess. From that classification it can be seen that the ESS in the left upper PIP in Figure 2.5 also attracts in the dimorphic regime.

2.5. Adaptive dynamics II: the canonical equation

For vectorial traits the geometric constructions exhibited for the scalar case go through in a general sense, but not necessarily with the same consequences. In particular the neat dependence of the dynamical outcomes on no more than the sign of the invasion fitness hinges on the ordering properties of the real line.

The main workhorse in the vectorial case is the so-called Canonical Equation (CE) of AD, a differential equation that captures how the trait vector changes over evolutionary time on the assumption that mutational steps are sufficiently small. See Figure 2.7⁶. The CE moreover adds a quantitative slant to the analysis, by taking account not only of whether a mutant can invade, but also of the probability that it does so. Its equilibrium points are the ess-es mentioned previously.

There are still few results about the quantitative match of the CE to the “real thing”. The main potential problem is that in reality many mutant substitutions will occur in parallel. Luckily, for small mutational steps this tends to affect the environment only in the higher order terms that in the derivation of the CE disappear from sight [50]. A second problem is that in

⁶ The form of the CE given in Figure 2.6 is for the case where the mutation distribution has mean zero and is symmetric around that mean. The CE was derived for ODE population models in [10] at a physicist level of rigour. A mathematically rigorous proof of the implied convergence followed seven years later [5]. These papers also give expressions applicable for more general mutation distributions. The derivation from the basic ingredients shows that the convergence to the CE is not uniform, becoming ever slower near ess-es. The extension to general structured population models was derived by [13], again at a physicist level or rigour; a mathematically rigorous proof for the purely age-dependent case can be found in [49]. The extension to the Mendelian case is made in [52] and [55]. The essential element in the latter extension is that for smooth genotype to phenotype maps, in the absence of any parental effects on gene expression, these maps are locally additive, so that the heterozygote between two different but similar homozygotes has a phenotype that is the average of the parental ones (Andrea Pugliese, pers. com. and [65]). The resident is invaded by heterozygotes, while after take-over there are only homozygotes left. This adds an additional factor 2. In [13] the term N_e in the version CE given in Figure 2.6 is still written as a product of the population size and some life history parameters. It is only recently that Vincent Jansen and I discovered that this product is actually equal to N_e whatever the ecological scenario.

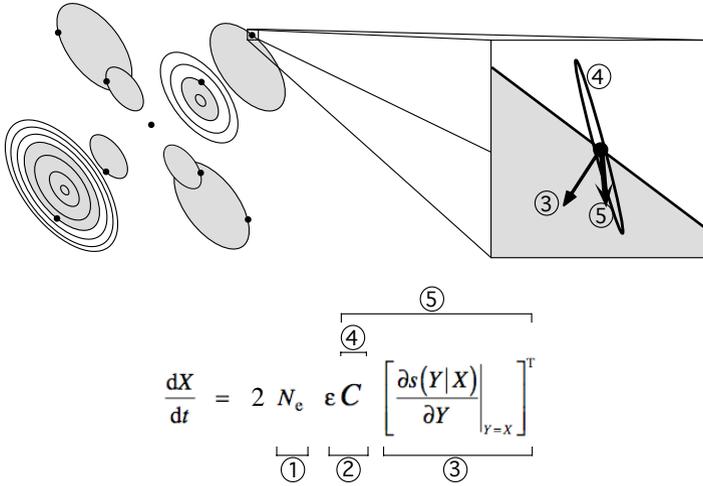


FIGURE 2.7. Upper left: Fitness landscapes for a selected number of residents (black dots; the dot in the center is an ESS (see Subsection 2.3). For two of the landscapes the contour lines are shown, for the other ones only the part is indicated where the fitness landscape is above sea level. Upper right: Enlargement of the coastal area around the resident for one of the fitness landscapes. Arrow ③ is the selection gradient. Ellipse ④ symbolises the probability distribution of mutational steps. Since the mutational steps in different directions are not equally probable an evolutionary movement results according to arrow ⑤. Below: The canonical equation of adaptive dynamics. The speed of evolutionary movement of the trait vector X equals twice the product of ① the effective population size (as defined in population genetics), ② a term summarising the nature of the mutational process, consisting of the mutation probability per birth event and the mutational covariance matrix as locally effective summary of the distribution of the mutational steps, and ③ the selection gradient.

the clonal case the effects of the invading mutants do not add up since a good mutant may be supplanted by an even better one coming from the same parent type. Hence the CE may be supposed to do a better job in the Mendelian case where the substitutions occur in parallel on different loci, which to the required order of approximation should interact additively. Remains that any standing genetic variation enlarges the variance of the offspring number of an allele, which should roughly act to proportionally diminish the effective population size. Unfortunately, this variance is a random variable which appears to show little temporal stability (Jörgen Ripa, pers com).

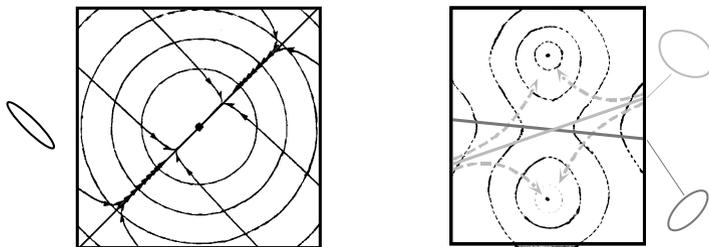


FIGURE 2.8. Two fitness landscapes that are supposed to keep their shape and only to sink when the adaptive trajectory moves uphill (as is the case if and only if the population regulation is through an additional state-independent death rate). Distributions of mutational steps are symbolised by ovals. Left: The shape of the mutation distribution induces a time scale separation between the movement along the diagonal and anti-diagonal direction. Right: The difference in mutation distributions causes a difference in the domains of attraction of the two ESS-es.

A final point is that the more mechanistic detail that has to be brought in for the derivation of the CE for the Mendelian case brings to the fore that the CE is but the first term in a moment expansion, on top of which comes, except under very special assumptions, a similar equation for the change of the mutational covariance matrix, which in turn depends on third mutational moments, etc..

2.6. Links between adaptive dynamics and Evo-Devo

From an AD perspective the link with Evo-Devo is first of all through the mutational covariance matrices. At this point in time Evo-Devo unfortunately has yet little to offer in this area, although there are some promising developments (e.g. [63, 64]). Therefore, at present often the most AD researchers can do is work out how the outcomes of a specific eco-evolutionary model depend on the possible forms of the mutational covariance matrix. The answers from AD thus become Evo-Devo questions: is the mutational covariance matrix for these traits expected to fall within this or that class?

Just to show the importance of the missing Evo-Devo input in AD: mutational covariance matrices have an, often dominating, influence on the time scales of evolution (Figure 2.8, left), the basins of attraction of ess-es (Figure 2.8, right), even to the extent that they often determine whether an ess attracts or not (c.f. [42, 43, 44]).

On a more philosophical level it bears noting that the selection gradient points only in a single direction, while the components of the trait vector orthogonal to that gradient hitchhike with the selectively determined motion thanks to a developmental coupling as expressed in the mutational covariance matrix. The higher the dimension of the trait space the larger the

contribution of development as a determinant of the direction of evolutionary motion. The dimensions of the trait spaces that are routinely considered thus makes for the contrast in attitudes of, for example, behavioural ecologists and morphologists, with the former stressing selection and the latter the developmental options for change. Formulated a bit facetiously, advocates of the absolute supremacy of selection basically think one-dimensionally, while harking on the evolutionary primacy of genes or development is like buying a car for its steering qualities without caring for the motor.

A final point is that I everywhere assume that the trait space for the largest part has the geometry of a manifold. However, in reality that geometry may be much more complicated, as it basically should reflect everything that can be generated by the developmental system. In particular, the trait space may consist of components with different dimension, as when a so-called key innovation adds a new trait that evolution then can seize upon to achieve some rapid progress and often also diversification. At present not much can be said in general about such issues. So I will in this chapter concentrate on conclusions based on the assumption that the trait space at least locally looks Euclidian.

2.7. Adaptive dynamics III: evolutionarily singular strategies

Evolutionarily singular strategies x^* can be calculated by setting the fitness gradient equal to zero. Figure 2.9 shows their classification according to dynamical type for the case of scalar traits⁷.

Devising a good classification for higher dimensional ess-es is an open problem. One of the reasons is that in higher dimensions the attractivity or non-attractivity of a singular point in general depends on the mutational covariance matrix, except in very special cases [42, 43, 44].

2.8. Adaptive speciation

The most interesting ess-es are branching points, where the eco-evolutionary process starts generating diversity. When approaching such points the evolutionary trajectory, although continually moving uphill, still gets itself into a fitness minimum. More precisely, it is overtaken by a fitness minimum. See Figure 2.10. The ecological cause can (by definition) be ascribed to so-called apparent competition, in the form of aversive direct interactions (such as

⁷ The classification is constructed by inserting the ecological consistency conditions $s(x|x) = 0$, $s(x_i|x_1, x_2) = 0$, $i = 1, 2$, $s(y|x_1, x_2) = s(y|x_2, x_1)$ and $s(y|x^*, x^*) = s(y|x^*)$ in the first and second order directional Taylor polynomials of $s(y|x)$ and $s(y|x_2, x_1)$ around the singular point, and analysing what sort of dynamics would result from such simplified invasion fitness functions (see [54, 30]). (In the mutant direction the necessary smoothness conditions can be made part of the modelling assumptions. However, in the resident direction the environment is a derived quantity that in the simplest case is calculated from the community dynamical equilibrium equations. At a point (x^*, x^*) the conditions for the inverse function theorem do not hold good and the map $(x_1, x_2) \mapsto E_{\text{attr}}$ is not a diffeomorphism. However, directional derivatives, and hence directional Taylor polynomials can be shown to exist [13].)

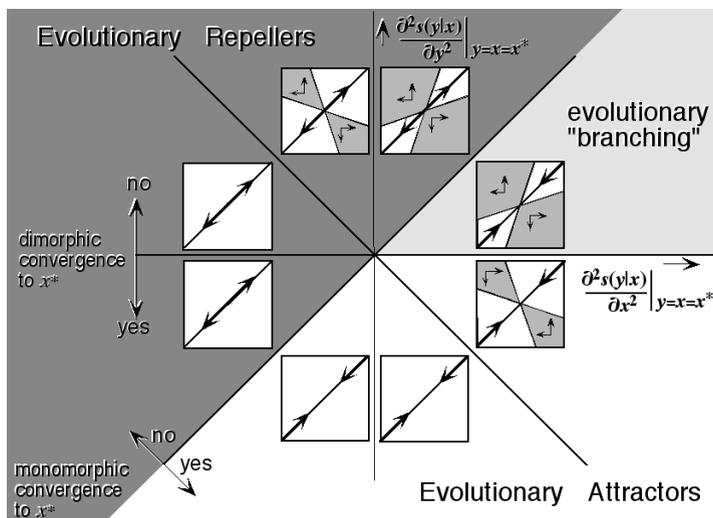


FIGURE 2.9. A classification of the ess-es for scalar traits. The cases in the lower half are all ESSes. The leftmost of these repels, the others attract. The latter ESSes are thus genuine evolutionary attractors. The branching points in the upper rightmost sector attract monomorphically but repel dimorphically.

fighting), competition for resources, having a common predator with a tendency to specialise on the most common types, etc. The following analogy may help intuiting the phenomenon. Somewhere gold has been found. As a result people converge to that special spot. However, after too many diggers have arrived, it becomes attractive to try one's luck at some distance.

The build up of diversity can take very different forms. In the clonal case the population just splits into two as depicted in Figure 2.1. In the Mendelian case the diversification starts with a broadening of the variation in the population. The fitness landscape locally has the shape of a parabola that increases away from x^* . This means that types more on the side have a higher fitness than those in the centre. It therefore pays not to get kids near the centre. The Mendelian mixer has the contrary tendency to produce intermediate kids out of dissimilar parents. Luckily, there are all sorts of mechanisms that may thwart this counterproductive mixing. The most interesting of these is the build up of some mechanism that lets the like extremes mate only among themselves, thus ensuring that the branches become separate genetic units. A very simple mechanism occurs in insects that diversify in their choice of host plants, with mating taking place on those hosts. More complicated mechanisms are mathematically explored in e.g. [8, 28, 12, 59, 61].

My own conviction is that in cases where no automatic mating barrier puts itself in place a build up of other mechanisms engendering assortative

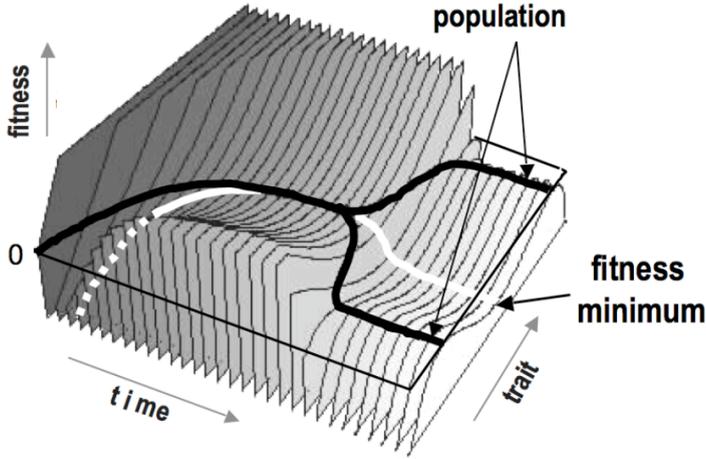


FIGURE 2.10. The development of the fitness landscape during a branching event.

mating is far from unexpected. Present day organisms are not simple particles, but are the product of some three and a half billion years of evolution. During that time their sensory and signalling apparatus has been evolutionarily honed for finding the most advantageous mates. As an example, spider mites, organisms one can barely see with the naked eye, were found (in an artificial arena) to mate assortatively on the basis of the food they had eaten previously [14]. Hence, I expect that there always will be an abundance of template mechanisms. Moreover, these mechanisms, once recruited to the task, will have a tendency to enhance each other in their effect. Therefore, I expect that the available generalised machinery often can rather easily be adapted so as to genetically separate the branches whenever evolution brings the population to a branching point. I should add, though, that most scientists working on the genetics of speciation do not seem to share this view.

3. Some meso-evolutionary predictions

To make meso-evolutionary predictions one has to look at recurring, close to model-independent, features of fitness landscapes. Some results derived by looking at the overall geometric features of well-behaved fitness landscapes are discussed below. I start with cases where there is no need to consider longer-term external environmental drivers, so that all environmental change is due to the community dynamics. (Any environmental fluctuations on a population dynamical time scale, like fluctuations in the day to day and year to year weather, are accounted for in the fitness function.)

3.1. The shortest time scale

1. When the tape of life is recorded in similar locations, say in the follow up of the colonisation of similar lakes after an ice age by the same fish species coming from the sea, the evolutionary trajectories will be similar, at least initially. One common pattern is that it moves to the same ESS. The other common pattern is that it goes to the same branching point. In both cases the trajectories in different localities move at similar speeds, and slow down in a similar manner when the ess is approached. If the latter is a branching point the trajectory lingers there for a while before the branches start to grow apart in a new phase of directional selection⁸. The very first divergence may be in many directions, but after that for a while two branches remain that initially diverge symmetrically, with similar speeds in different locations⁹. (I expect the lingering near the branching point to be more variable in duration since the exact speciation mechanism may depend on specific local opportunities, and therefore will be less repeatable than evolutionary features governed by the gross ecology.)

2. Initially in the phases of directional selection a lot of non-adaptive variability will be incorporated, most of which will be weeded out at later stages. The reason is that initially the fitness landscape looks like a single steep hill, with moreover the resident far removed from the hilltop, so that the local fitness contours have low curvature, as in the left panel of Figure 3.1. As a result all sorts of mutants that are located at some distance from the resident in the almost neutral direction orthogonal to the fitness gradient can invade. Some of the trait changes caused by such mutations may actually be deleterious, but with that deleteriousness compensated by a pleiotropic advantageous change in some other traits. At a later stage, when the residents are closer to a top of the fitness landscape and the fitness contours have become correspondingly more curved, as in the right panel in Figure 3.1, the evolutionary path will become more and more constrained, and in the end all traits will end up at their most advantageous values for the environment realised at that time.

⁸In the clonal case in the small mutational step limit, the speed of diverging from a branching point is third order in the mutational step size, whereas the speed of directional movement is second order in the step size. However, this property is non-robust relative to relaxing the assumption of mutation limitation. For the clonal case this makes the above arguments about the consequences of the fitness landscape shapes less than robust. However, as with the quality of the approximation by the canonical equation, the adaptive dynamics results may without strict mutation limitation actually be saved by Mendelian inheritance as the need for the development of a mechanism for thwarting the Mendelian mixer may be expected to restore the time scale separation between directional selection and branching.

⁹The maximum number of branches that can be present initially equals one plus the dimension of the trait space. However, there are strong mathematical indications that soon only two branches remain that diverge in the direction of the dominant eigenvector of the second partial derivative of the fitness function in the invader direction (Stefan Geritz, pers. com.).

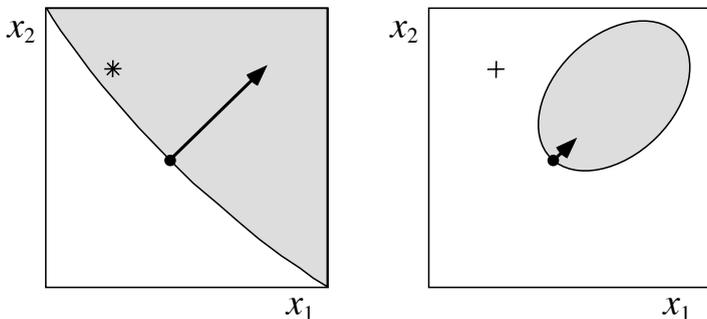


FIGURE 3.1. Left: Local fitness landscape shortly after the colonisation of a new territory, when a lot of slightly deleterious variation may get incorporated. In the example fitness increases with both x_1 and x_2 . Yet, the mutant that is depicted will invade, even though the mutation decreases x_1 , since under the circumstances this decrease is more than compensated by the increase in x_2 . Right: Fitness landscape at a later stage of the adaptive process when such non-adaptive variation gets weeded out again.

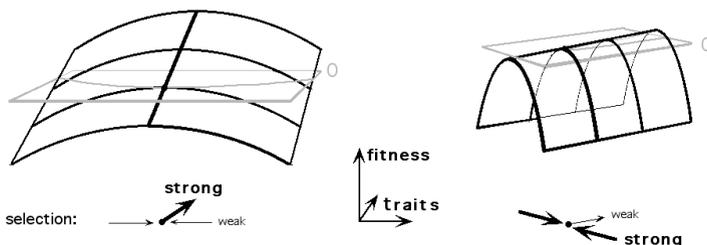


FIGURE 3.2. Left: Representative low dimensional slice of the local fitness landscape shortly after the colonisation of a new territory, when a lot of slightly deleterious variation may get incorporated. Right: Representative low dimensional slice of the local fitness landscape at a later stage of the adaptive process where such variation gets removed again.

Real adaptive processes usually take place in a high dimensional space with different speeds prevailing in different directions. Therefore, the weeding out will in general already start long before any final adaptive stops are reached. Figure 3.2 shows low-dimensional slices of fitness landscapes indicative of the relative balance of directional and stabilising selection in early and later stages of the adaptive process.

3.2. Intermediate time scales

3. When an empty habitat is colonised, initially speciation will occur frequently but the overall speciation rate may be expected to decrease rather quickly. The reason can be seen in Figure 2.1. Initially the fitness landscape tends to be a single big hill, with submerged outer slopes and the resident population somewhere along the shoreline. With an increasing number of diversification events the landscape becomes anchored to zero at more and more points. Hence, at any branching point the steepness of the surrounding upward slopes (as expressed by the second derivative) becomes less and less¹⁰. With a decrease of the fitness slopes the speed at which branching points are reached and left becomes slower. In the Mendelian case the lower local curvature decreases the selection on mechanisms promoting assortative mating, which should slow down the speed of diversification even further (and actually may be expected to dominate that speed except in cases with full mutation limitation).

4. Speciation should be rare in environments that fluctuate on time scales between that of the faster scale of directional evolution and the slower scale of speciation. The reason is depicted in Figure 3.3. Speciation can only start from some very special points in trait space. (In the past this fact was even used as an argument for the improbability of adaptive speciation; e.g. [45, 36]. Although it may seem ecologically unusual for a population to sit at precisely such a special point, one of the lessons from AD is that evolution may actually guide a population towards precisely such a point; see [40].) Slow fluctuations of the physical environment change the picture, in that the branching points will not stay in place but will move in response to the changes in the parameters of the community dynamics. This means that two incipient species will after some while find themselves no longer in the coexistence cone emanating from the branching point, as that cone has moved to another position, and one or the other of them will go extinct, aborting the diversification process. In simulations all this shows up as an abundance of diversification attempts which all turn into dead ends, resulting in an adaptive path with the shape of a closely pruned tree.

3.3. The longest time scales

5. In the nineteen-seventies paleontologists made a point about the common occurrence in the fossil record of so-called punctuated equilibria: short periods of rapid morphological change amidst longish periods of near morphological stasis [15, 32]. This phenomenon was ascribed by the inventors of the term

¹⁰At this moment I have but a heuristic argument for this statement. A rigorous proof should be based on the assumptions that (1) there is a uniform bound on the trait vectors beyond which the fitness landscape disappears below sea level, and (2) the derivatives of the fitness landscape, in particular the third one, are uniformly bounded. As the properties of the fitness landscape reflect only individual level mechanisms these are fair assumptions. (This in contrast to any assumptions about the dependence of the fitness function $s(Y|X)$ on the residents X as the latter dependency is through the community attractor, the properties of which need not depend everywhere smoothly on the resident traits.)

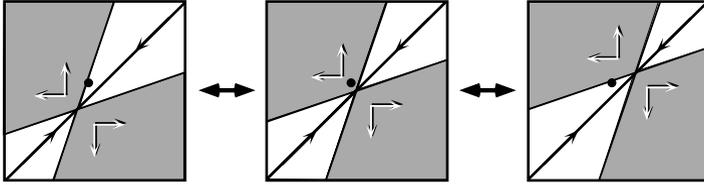


FIGURE 3.3. How a TEP may change as a result of slow fluctuations in the physical environment. As a result, if a species starts to branch (middle panel) the two incipient species may often in a short while find themselves outside the coexistence region (left and right panel).

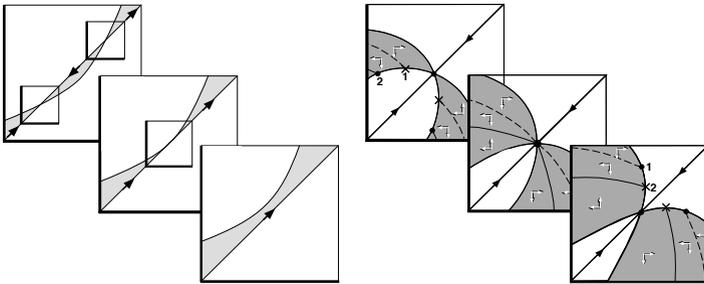


FIGURE 3.4. Robust AD bifurcations that will show up in the fossil record as “punctuated equilibria”. Left: Three PIPs corresponding to a saddle node bifurcation in which an ESS is annihilated by an evolutionary repeller. After a longish period in which the adaptive trajectory tracks the slowly shifting ESS, this ESS abruptly disappears and the trajectory punctuates, i.e., embarks on a much more rapid path to some other attractor. Right: Three TEPs corresponding to a bifurcation of an ESS into a branching point. In this case the adaptive tracking of the ESS stops due to a change in character of the ess, resulting in a punctuation event that starts with speciation.

to the occurrence of “morphological revolutions”. With the paleontologists, I believe in the reality of punctuated equilibria, but I beg to differ with their explanation. I think that the time scales are too short and the morphological disparity that is involved too small for macro-evolutionary mechanisms really to play a role. Moreover, AD predicts the same phenomenon for mathematically well-established ecological reasons.

Given the customary speeds of directional evolution one may expect most slow changes in the fossil record to be due to the evolutionary tracking of slowly changing adaptive equilibria. These slow changes may, however, be punctuated by rare far shorter periods of fast directional evolution.

The change in the ESSes can only be ascribed to slow overall environmental changes, e.g. changes in the climate or in the averaged action of the surrounding biosphere with its millions of simultaneously evolving species. For extended periods these changes will affect the ESSes only in a quantitative manner. However, mathematics tells that such relatively stable regimes will obtain only for so long as the parameter path does not cross any bifurcation points. Figure 3.4 depicts the two common bifurcations of ESSes. The corresponding punctuation events are of two types, “just so”, due to the collision of an ESS and an evolutionary repeller, and coupled to speciation, when an ESS changes into a branching point, in good accordance with the findings of the paleontologists.

4. Macro-evolution

4.1. Evolutionary tinkering, tangled maps, and the need for intermediate abstractions

On a longer, macro-evolutionary, time scale, one has to take into consideration that the overall properties realised during evolution can in principle be realised by very different mechanisms. In general the first mechanism that does a sufficient job inherits the earth. Therefore, analysing which mechanisms should be most easy to realise has considerable predictive power. Moreover, evolution does not necessarily solve problems in the best possible manner. Evolution tinkers; it only optimises under very special circumstances, and then only very locally.

In the short term it may not matter which mechanism realises a certain desirable trait, but in the longer term different mechanisms lead to different mutational covariance matrices and hence to different evolutionary routes. In the language of dynamical systems: the real state space of the evolutionary process is not phenotype space, of which the trait spaces of AD are convenient abstractions, but genotype space. The mutational covariances reflect both the topology of genotype space, as generated by mutational distances, and the genotype to phenotype map generated by the developmental mechanics. This reflection may be not too inadequate locally in genotype space, and therefore locally in evolutionary time. However, for larger scale considerations different approaches are needed, both to delineate the domain of applicability of the simpler framework, and to step beyond its confines.

At the present state of our understanding the knowledge available about the detailed nitty gritty at the molecular level does not seem helpful yet for the questions I have in mind, however interesting it may be in other respects. The reason is the evolved complexity of the developmental process and the resulting tangledness of the genotype to phenotype map.

An indirect proof that the genotype to phenotype map is inexorably tangled is that assuming such tangledness appears the only way to resolve the discrepancy between the domination of adaptive processes as perceived

by ecologists, functional morphologists, and the like, and the relatively satisfactory description that random models give of evolution at the DNA level. This resolution is moreover nicely compatible with the two main failings of the random model: (1) the different speeds of evolutionary change of different pieces of the genome which seem closely related to their functionality just one or a few translation steps away from the genome, but not further, (2) the far higher than Poisson variance in the number of substitutions, which most probably reflect the iteration of selective sweeps (see e.g. [31]).

The previous arguments indicate the need for intermediate abstractions. In the remainder of this chapter I shall discuss some ideas to that effect. Most of these ideas were developed together with Frietson Galis, who is a functional morphologist by training, but over the last years has switched interest to Evo-Devo, where she concentrates precisely on the level that is of most interest in the present context. All the considerations below have to do with so-called internal selection processes (a definition follows below), and once again with the geometry of fitness landscapes. The third player in the game will be the consequences of high dimensionality both of phenotype, but in particular of genotype space. Some of these ideas go back to Sir Ronald Fisher, others are borrowed from Arno Wouters, Günter Wagner, Walter Fontana and Sergey Gavrillets.

4.2. Phenotypes (and genotypes)

Below I will generally leave unspecified whether the fitness landscapes under consideration are over a genotype or a phenotype space. The basic idea is that for most phenotype spaces of interest the genotype to phenotype map comes about as the concatenation of (a great number of) maps between other spaces (c.f. Figure 2.2), coordinates of which can also be chosen as phenotypic coordinates. Below I will argue that the fitness landscape over phenotype space more often than not is very "ridgy". In general, most of the ridginess over a phenotype space may be expected to automatically cascade back to any underlying spaces, although in principle some of it can be cancelled (think of how the chain rule acts) and some other added. Ultimately this ridginess then also cascades back to genotype space.

A related point is that I repeatedly fell back on some implicit smoothness assumptions, and that I will keep doing so. The reason why I think that this is justified is that I believe that the evolutionary changes that I am considering are mostly not so much changes in the coding regions of classical genes as well as in their regulation. Protein coding regions are in general preceded by a large number of relatively short regions where all sorts of regulatory material can dock. Changes in these docking regions, and changes in genes producing regulatory proteins lead to changes in the production rate of the gene product. Genes are more or less active in different parts of the body, at different times during development and under different micro-environmental conditions. So the lowest layer in the cascade of phenotype spaces is a space of gene expression levels as a function of these variables. This is the level from which we may start to think about the smallness of mutational steps.

The influence of any specific regulatory site tends to be rather minor, and most changes in regulatory proteins will have only a minor effect on their affinity for the docking region. When I referred to the smoothness of a genotype to phenotype map in Subsection 2.5 Footnote 6, I referred to the map from this high dimensional vector of gene expression levels to the phenotypic coordinates under consideration.

From the preceding paragraph it has probably already become clear that also the phenotypes of this story may not always be what one may naïvely expect. In keeping with the general ecological view that formed the basis for the ecological definition of fitness, phenotypes should in principle be interpreted as so-called reaction norms (another term is conditional strategies), i.e., maps from micro-environmental conditions to phenotypes in the naïve sense, i.e., characteristics of individuals. These reaction norms supposedly come as families with the phenotypic traits as identifying parameters. Only in the simplest cases these reaction norms are degenerate, taking only a single value, which we then may use as the phenotypic trait, so that the general and naïve sense phenotypes coincide.

4.3. Internal selection

Functional and constructional morphologists usually talk in terms of whether certain mechanisms work properly or not, and discuss evolution as a sequence of mechanisms all of which should work properly, and which only slightly change in every single transformational step. Translated into the language of fitness landscapes, this means that only the properly working mechanisms give fitnesses in the ecologically relevant range, while the improperly working ones give very low fitnesses for all relevant environmental conditions (c.f. [71]). This leads to a picture of narrow, slightly sloping, ridges in a very high dimensional fitness landscape. The slopes on top of the ridges are the domain of ecology, their overall location is largely ecology independent (c.f. Figure 4.1).

As a simple example you may think of human leg length. Few ecologists will ever consider the length of the right and left leg as separate traits. The reason is that these lengths are kept equal by a very strong selection pressure, which keeps in place a developmental system that produces legs of precisely equal length, notwithstanding the fact that during development there is no direct coupling between the processes operating in the two leg primordia. Hence in a trait space spanned by the lengths of the right and left leg ecologists concentrate on just the diagonal.

Please notice that whereas ridges in the familiar three dimensional world are necessarily one dimensional, the trait spaces dealt with in morphology are very high dimensional so that the top of a ridge may be higher dimensional, while away from the ridge the fitness decreases very steeply in a far larger number of orthogonal directions.

A picture similar to that of the functional morphologists emerges from the consideration of developmental processes. The long term conservation of developmental units, think of the phylotypic stage or of homology, can only be

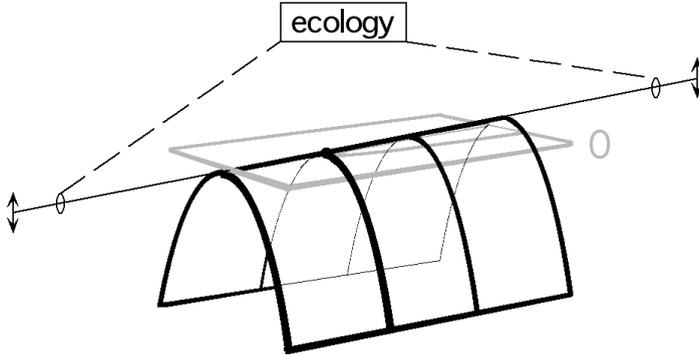


FIGURE 4.1. The interplay between internal and ecological selection. Internal selection sets the large-scale pattern of a deep ocean from which ridges that rise to around sea level. Ecology modulates the elevation on the upper slopes of the ridges.

due to strong stabilising selection, caused by the fact that mutations causing large pattern changes generally have many side effects with dire consequences for fitness (see e.g. [22, 23, 24, 25]). As a result, ecological selection generally acts only on quantitative changes in the shapes and sizes of homologous body parts. As the fitness differences that are the guardians of homology usually manifest themselves already very early in the lives of individuals, in mammals in the womb, people tend to speak here of internal selection. In this chapter I will keep this term, but will use it not in a mechanistic but in a structural sense, as a reference to features of the fitness landscape that, within the confines of a particular argument, can be considered as relatively unchanging, that is, that are roughly the same for all the environments that explicitly or implicitly figure in that argument.

4.4. A macro-evolutionary prediction

Any developmental system that leads to aligning the mutational steps with the direction of the ridge on which the resident is currently sitting, will evolve much faster than a system that is not aligned in such a manner (Figure 4.2). One way in which such a bias can come about is by having the development use cues that are related to the later function of the organ under consideration. In vertebrates, bones, muscles and nerve cells are modelled and/or grow in the embryo depending on their use. Unborn infants move for good reasons, and play is highly functional. As a consequence mammalian morphology evolves much faster than the morphology of e.g. insects, in particular indirectly developing ones (see Table 1). Insects, on the other hand, do far better at the chemical end (think e.g. of DDT resistance), due to their shorter generation times and far larger population sizes, and consequent higher availability of potentially useful mutations. In mammals teeth, which like butterfly

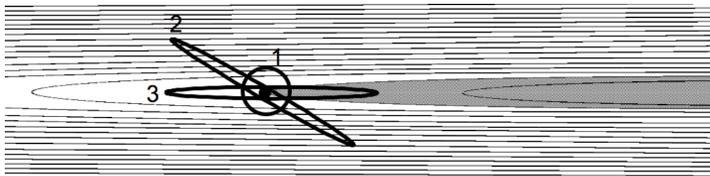


FIGURE 4.2. A steep, slightly sloping, fitness ridge together with three contours of potential unbiased mutation distributions. In distribution 1 all mutational directions are equally probable, distribution 2 is aligned with the ridge, while distribution 3 is similarly elongated, but misaligned. Clearly the corresponding evolutionary rates r_i are ordered as $r_3 > r_1 > r_2$. Effects like these become stronger with a larger number of off-ridge directions.

adults have to develop mainly under ballistic developmental control, are so slow to evolve that they are used to define the higher levels in the Linnean classification.

Dipteran families	179	330
Drosophila subgenera	60	110
Mammalian orders	38	70

4.5. An Evo-Devo myth

The stabilising selection that underlies the long-term conservation of developmental units should have led evolutionarily to a considerable robustness of the developmental process. The selection is indirect, as a side effect of selection for robustness against environmental perturbations. It is, after all, disadvantageous after an unusually cold day to end up with a baby with two heads. The resulting robustness of the developmental system against perturbations will necessarily carry robustness against mutational perturbations in its wake. The inevitably resulting abundance of suspender and belt combinations and the overall tinkering nature of evolution have no doubt made their contribution to the tangledness of the genotype to phenotype map mentioned earlier. In the language of fitness landscapes, this robustness translates into the existence of extensive near neutral sets in genotype space: many mutational steps will have little effect due to the buffering by the developmental control system. The presence of extensive near neutral sets will also make for a mazelike character of the high fitness ridges in this space.

Contrary to naive expectation, however, the robustness of parts of the developmental process cannot be assumed by itself to conserve developmental units, or constrain their evolution. During the long periods of effective evolutionary stasis alluded to in Subsection 3.3, the population necessarily oozes

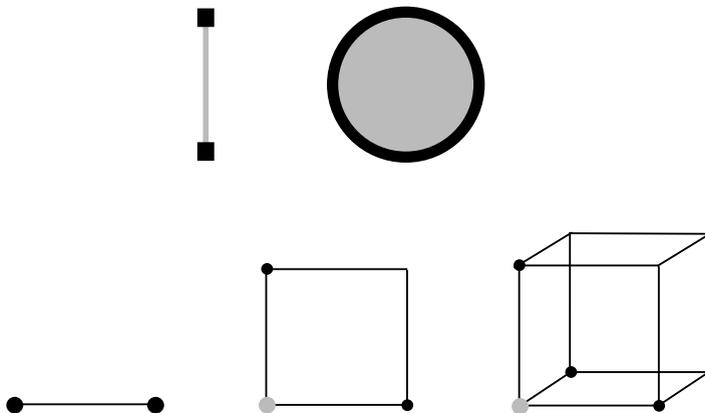


FIGURE 4.3. Top: Balls in 1 and 2 dimensions, with the points that are within a distance ε from the boundary coloured black. Clearly, the fraction of points that is within a ε -distance from the boundary is larger in 2 dimensions. With an increase in the number of dimensions this fraction increases further to go to 1 in the limit, independent of ε . Bottom: a similar construction on the Boolean cube in 1, 2 and 3 dimensions. Here as well, the fraction of points close to the boundary goes to 1 when the dimension goes up.

as a diffuse cloud through the corresponding neutral set. In considering these sets one has to account for the effects of high dimensionality. Genotype space is very high dimensional. For sets in high dimensional spaces in general most points are very close to the boundary (see Figure 4.5). The diffusive oozing together with this geometric effect should make any longer term thwarting of phenotypic change next to impossible.

The upshot is that any long-term conservation can only be due to strong stabilising selection, or in other words, a fitness abyss that keeps the traits confined to narrow high fitness ridges. Hence homology. Evolution thus largely proceeds through the change in properties of homologous elements, the identity of which is conserved by stabilising internal selection.

5. Back to meso-evolution

5.1. Justifying adaptive dynamics

Fisher's old argument that the higher the dimensionality of the trait space the more difficult becomes the final convergence to an adaptive top (see Figure 5.1; see also [21, 60, 70, 68, 69]) seamlessly extends to the movement in a ridgy fitness landscape: the higher the number of orthogonal off-ridge directions, the more rare it is for a mutational step to end up above sea level. By a similar argument (see Figure 5.2) small mutational steps have a far

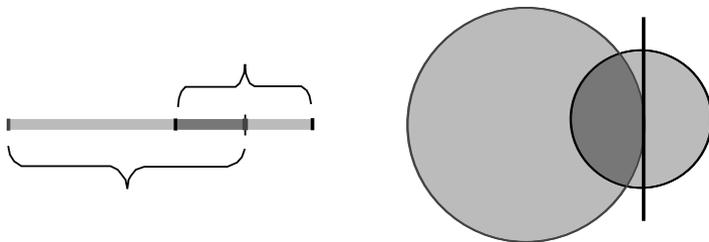


FIGURE 5.1. Left: Two balls in \mathbb{R}^1 , with the centre of the smaller ball on the boundary of the larger ball. The ratio of the volume of their intersection to the volume of the smaller ball is $\frac{1}{2}$. Right: A similar configuration in \mathbb{R}^2 . The volume of the intersection is now a smaller fraction of the volume of the smaller ball. For similar configurations in \mathbb{R}^n this fraction quickly decreases to zero for larger n . Now think of the larger ball as the part above sea level of a fitness hill and of the smaller ball as a mutation distribution. Clearly the fraction of favourable mutants will go to zero with n .

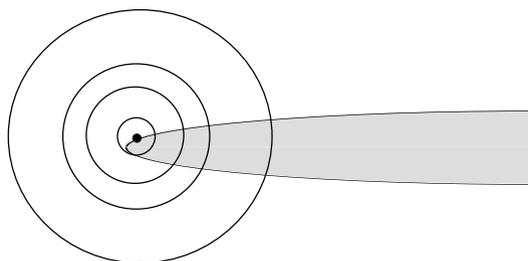


FIGURE 5.2. The shorelines of a fitness ridge and the contour lines of a mutation distribution. On land smaller mutational steps will be overrepresented relative to the larger ones.

higher propensity to end up above sea level than have large ones. Together these two arguments seem to underpin the requirements of AD that mutations in the ecologically relevant directions are scarce and the induced mutational steps generally small.

Notwithstanding the appeal of the above arguments, they contain a biological flaw, the assumed rotational symmetry of the distribution of mutational steps. Real mutation distributions may be expected to show strong correlations between traits. Correlation structures can be represented in terms of principal components. I have not been able to find any good empirical data. However, the general experience with biological data is that almost always patterns are found like the ones shown in Figure 5.3. Figure 5.4 makes clear that the existence of mutational correlations will in general enhance rather

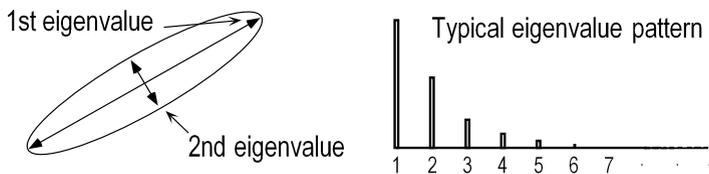


FIGURE 5.3. Contour line of a bivariate distribution, supposedly of mutational steps. The lengths of the two axes of the ellipse, called principal components, are proportional to the square root of the eigenvalues of the mutational covariance matrix. Right: Typical eigenvalue pattern found for large empirical covariance matrices.

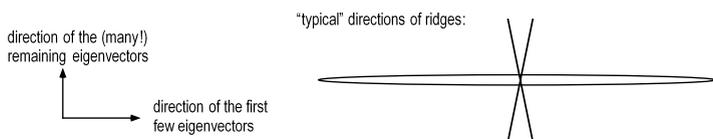


FIGURE 5.4. The mutation distribution will rarely be fully aligned with the fitness ridges. If one takes one's perspective from the mutation distribution and looks at the orientation of the ridges relative to the first few principal axes of this distribution, then, when the number of the dimensions of the trait space is very large and the ridge has a relatively low dimensional top, the ridge will typically extend in a direction of relatively small mutational variation.

than diminish the rareness and relative smallness of the mutational steps that end up above sea level, unless there is a very strong mechanistic link between the direction of the fitness ridges and the first principal axes of the mutation distribution, even stronger than the ones considered in Subsection 4.3¹¹.

The above conclusions seem to underpin nicely the assumptions of AD. Unfortunately, there are empirical observations that appear to contradict these conclusions. Populations brought into the lab always seem to harbour sufficient standing genetic variation to allow quick responses to selection, and so-called Quantitative Trait Loci are often found to underlie the variation in a trait. There are a number of reasons why I believe that these empirical observations may have less bearing on the issue than one might think. First, given the speed of evolution relative to the changes in the overall conditions

¹¹It would be of interest to try to derive the canonical equation of AD, or some similar equation, from such a fitness ridge perspective, instead of from the direct assumption that mutational steps are small. (Two limit procedures come to mind. In the first one the side slopes of the ridge are made steeper and steeper. In the second, probably more appropriate, one the dimension of the trait space and the number of orthogonal off-ridge directions are simultaneously increased.)

of life, populations in the wild are probably most often hanging around some ESS. Moreover, real environments are inextricably noisy. Eco-evolutionary models with environmental noise almost invariably produce ESSes with very flat fitness maxima. This means that after a while a considerable amount of near neutral genetic variation will accumulate, which is exploited first when a population gets artificially selected on. The statics of AD corresponds to standard ESS theory. At ESSes the mutation limitation question is largely moot. Beyond the statics, AD's main interest is in the larger scale features of evolutionary trajectories after the colonisation of new territory or, even grander, a mass extinction. The scale of these features may be expected to require a further mutational supply of variation. Second, environmental fluctuations often lead to the accumulation of a lot of variability in traits directly involved in coping with those fluctuations. (Models incorporating environmental fluctuations often lead to adaptive branching in traits that modulate an individual's reaction to the environmental driver.) In the lab those fluctuations are removed, and selection is also otherwise relaxed but for the traits that are specifically selected for. This means that a lot more variation becomes available for the latter selection than would ever be available in the wild. The same holds for domesticated organisms. The selection of dogs for extreme sizes has had the side effect of producing lots of genetic diseases [26]. Third, AD style theory has shown that in the absence of assortative mating the initial increase of variability after the reaching of a branching point tends in the course of time to get redistributed over a smaller number of loci with increasing relative effect [41, 66]. The end effect will be QTLs, but these are produced through the cumulative effect of small genetic modifications.

5.2. Allopatric speciation

High dimensional ridginess also lies at the base of the usual ideas about so-called allopatric speciation, that is, the origination of reproductive incompatibility as a byproduct of long-term geographic separation. The underlying intuition is that separated populations independently wander around in the high fitness maze. After they have sufficiently diverged, any mixed offspring that occurs when they are secondarily confronted with each other ends up in the abyss. However, the models in this area have all been rigged to get the desired effect. Simple random genotype to phenotype maps followed by a more regular map to fitness almost never appear to do the job (Eke van Batenburg, Carolien de Kovel, pers. com.). Of course, real genotype to phenotype and hence to fitness maps are constrained by earlier evolution as well as by the technical requirements for producing well functioning bodies. This leaves the interesting research question which combinations of ecologies and developmental maps are more and which are less prone to speciation, be it allopatric or sympatric.

Personally I believe that the ridginess of the fitness landscape plays its role largely on a macro-evolutionary scale, while speciation typically is a meso-evolutionary phenomenon, and that therefore the developmental-biology-based intuition is unfounded. This does not mean that allopatric

speciation is a non-phenomenon, as there are good alternatives to the developmentally-based fitness ridge scenario. One such scenario is based in wars of the sexes. For the sake of simplicity I will concentrate on the case of aquatic mass spawners, but similar scenarios can be dreamt up for species with more sophisticated mating systems. Ova and sperm find each other by means of chemical attractants. However, ova besieged by too many sperm are killed. Hence it pays an ovum to produce an attractant mix somewhat at the edge of the population distribution, to which not too many sperm are yet fully adapted. As a result, attractant mixes and receptor capabilities keep wandering in chemo-space. As the movement takes place in a very high dimensional space, separated populations soon become incompatible. The prediction is that the genetic basis for reproductive incompatibilities usually will have to be sought in genes involved in such arms races, rather than in genes for building bodies.

Notwithstanding the somewhat *deus ex machina* character of its fundamental ingredient, for long the dogma has held sway that almost all speciation would be allopatric. Part of the strength of this dogma was due to its promulgation by Ernst Mayr, of modern synthesis fame. Mayr started as an ornithologist, working on bird faunas in the Pacific. I indeed agree, on a priori grounds, that those birds will in almost all cases have speciated allopatrically. The reason is that before on a more remote island a species has gone through the full process of splitting up ecologically as well as genetically, almost certainly an immigrant from another island will have arrived, fulfilling one of the missing ecological roles, and better than the novel local candidate. This new immigrant thus will undercut the ecological basis for the speciation process. On the other hand, birds have complicated sound and colour based mating signals, that may be subject to considerable drift, both as a result of the local circumstances influencing the effectiveness of different variants of the signals, and also as the result of processes with a dynamics similar to that of evolutionary arms races. So diversification in such patterns between different islands would not be unexpected. Moreover, bird taxonomists, in sync with birds, are inclined to attach more importance to dissimilarities in song and coloration than to ecological similarities. However, the story for the snails or palms on those same islands may well be more in line with the adaptive speciation scenario.

An entirely different consideration is that paleontologically song and colour differences hold little interest as opposed to the morphological differences that go with ecological differentiation, in particular since the latter differences may form the fodder for further adaptive radiation. Hence, even when adaptive speciation scenarios do not hold true in the fine details, thinking in such scenarios may still help interpreting longer-term evolutionary patterns.

6. How to get on?

I have now come at the end of my ramblings. I have still one further message that I want to tout: the whole area of macro- and meso-evolutionary modelling is rife with open problems.

On the Evo-Devo side I have hit upon a few general geometrical arguments, in particular ones pertaining to the effect of high dimensionality. The geometrical features of those high dimensional sets are sufficiently counter-intuitive that I had to revise many of my preconceptions. I hope that my discourse has convinced you that such geometries should be the rule rather than the exception. Some open questions are: How can one best characterise high fitness mazes? Are there options for a reduced characterisation of the genotype to phenotype map within a maze? Is it possible to find abstract characterisations that help answering the question posed halfway the previous subsection?

The AD side is much further developed, which does not mean that there are no challenges left. These challenges are more about extending the mathematics than about defining the framework, although, of course, creatively doing the former necessarily asks for at least some smattering of the latter. Some open problems are extending the classification of singular point to higher dimensional trait spaces and developing a full-fledged bifurcation theory for ESSes. Partial results abound, but we are far from having the full picture yet. More at the modelling end there are the question of developing a good background theory for dealing with constraints on the trait space, and the question what biological consistency conditions can be found that constrain the possible geometric structures occurring in AD on top of the ones mentioned in Footnote 7 (already many more are known than I have mentioned!).

I urge you, dear reader, to take up these challenges, and join in further developing this area.

References

- [1] Abrams, P.A. (2001) Modelling the adaptive dynamics of traits involved in inter- and intraspecific interactions: An assessment of three methods. *Ecology Letters* 4: 166-175.
- [2] Amundson, R. (2005) *The Changing Role of the Embryo in Evolutionary Thought: Roots of Evo-Devo*. Cambridge University Press.
- [3] Athreya K.B. & S. Karlin (1971a) On branching processes with random environments: I extinction probabilities. *Ann. Math. Stat.* 42: 1499-1520.
- [4] Athreya K.B. & S. Karlin (1971b) Branching processes with random environments: II limit theorems. *Ann. Math. Stat.* 42: 1843-1858.
- [5] Champagnat, N. (2003) Convergence of adaptive dynamics n-morphic jump processes to the canonical equation and degenerate diffusion approximation. Preprint of the University of Nanterre (Paris 10) No. 03/7.

- [6] Champagnat, N. (2006) A microscopic interpretation for adaptive dynamics trait substitution sequence models. *Stoch. Proc. Appl.* 116: 1127-1160.
- [7] Dercole, F. & S. Rinaldi (2008) *Analysis of Evolutionary Processes: the Adaptive Dynamics Approach and its Applications*. Princeton University Press
- [8] Dieckmann U. & M. Doebeli (1999) On the origin of species by sympatric speciation. *Nature* 400:354-357.
- [9] Dieckmann, U., M. Doebeli, J.A.J. Metz, & D. Tautz eds. (2004) *Adaptive Speciation*. Cambridge University Press.
- [10] Dieckmann, U. & Law, R. (1996) The dynamical theory of coevolution: a derivation from stochastic ecological processes. *J. Math. Biol.* 34: 579-612.
- [11] Diekmann, O., M. Gyllenberg & J.A.J. Metz (2003) Steady State Analysis of Structured Population Models. *Theor. Pop. Biol.* 63:309-338.
- [12] Doebeli M. & U. Dieckmann. 2003. Speciation along environmental gradients. *Nature* 421:259-264.
- [13] Durinx, M., J.A.J. Metz & G. Meszéna (2008) Adaptive dynamics for physiologically structured models. *J. Math. Biol.* 56: 673-742
- [14] Egas, M., M.W. Sabelis, F. Vala, & I. Lesna (2004) Adaptive Speciation in Agricultural Pests. Pp. 249-263 in U.
- [15] Eldredge, N. & S.J. Gould (1972) Punctuated equilibria: an alternative to phyletic gradualism. Pp. 82-115 in T.J.M. Schopf (ed) *Models in Paleobiology*. Freeman, Cooper and Company, San Francisco.
- [16] Eshel, I. (1983) Evolutionary and continuous stability. *J. Theor. Biol.* 103: 99-111.
- [17] Eshel I., M.W. Feldman & A. Bergman (1998) Long-term evolution, short-term evolution, and population genetic theory. *J. theor. Biol.* 191: 280-285.
- [18] Eshel, I. (in press) Short-term and long-term Evolution. In U. Dieckmann & J.A.J. Metz (eds.) *Elements of adaptive dynamics*. Cambridge University Press.
- [19] Ferrière, R., & Gatto, M. (1995). Lyapunov exponents and the mathematics of invasion in oscillatory or chaotic populations. *Theor. Pop. Biol.* 48, 126-171.
- [20] Fisher, R. A. (1922) On the dominance ratio. *Proc. Roy. Soc. Edinburgh* 42: 321-341.
- [21] Fisher, R. A. (1930) *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford.
- [22] Galis, F. & J.A.J. Metz (2001) Testing the vulnerability of the phylotypic stage: on modularity and evolutionary conservation. *J exp Zool (Mol Dev Evol)* 291:195-204.
- [23] Galis, F., J.J.M. van Alphen & J.A.J. Metz (2001) Why do we have five fingers? The evolutionary constraint on digit numbers. *TREE* 16: 637-646.
- [24] Galis, F., T.J.M. van Dooren & J.A.J. Metz (2002) Conservation of the segmented germband stage: modularity and robustness or pleiotropy and stabilizing selection? *TIG* 18: 504-509.
- [25] Galis, F., T.J.M. van Dooren, J.D. Feuth, J.A.J. Metz, A. Witkam, S. Ruinard, M. Steigenga & L.C. Wijnaendts (2006) Extreme selection against homeotic transformations of cervical vertebrae in humans. *Evolution* 60: 2643-2654.

- [26] Galis, F., I. van der Sluijs, T.J.M. van Dooren, J.A.J. Metz & M. Nussbaumer (2007) Do large dogs die young? *J. Exper. Zool. Part B: Mol. Devel, Evol.* 308B: 119-126.
- [27] Geritz S.A.H. (2005) Resident-invader dynamics and the coexistence of similar strategies. *J. Math. Biol.* 50: 67-82
- [28] Geritz S. A. H. & E. Kisdi (2000) Adaptive dynamics in diploid, sexual populations and the evolution of reproductive isolation. *Proc. R. Soc. Lond. B* 267: 1671-1678.
- [29] Geritz, S.A.H, M. Gyllenberg, F.J.A. Jacobs & K. Parvinen (2002) Invasion dynamics and attractor inheritance. *J. Math. Biol.* 44: 548-560.
- [30] Geritz, S.A.H., É Kisdi, G. Meszéna & J.A.J. Metz (1998) Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology* 12, 35-57.
- [31] Gillespie, J.H. (1991) *The causes of molecular evolution.* Oxford University Press.
- [32] Gould, S.J.& N. Eldredge (1977) Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology* 3: 115-151.
- [33] Gyllenberg, M., F.J.A. Jacobs, & J.A.J. Metz (2003) On the concept of attractor in community-dynamical processes II: The case of structured populations. *J. Math. Biol.* 47: 235-248.
- [34] Haccou, P., P. Jagers, & V. Vatutin (2005) *Branching Processes - Variation, Growth, and Extinction of Populations.* Cambridge Studies in Adaptive Dynamics 5. Cambridge University Press.
- [35] Hamilton W.D. (1967). Extraordinary sex ratios. *Science* 156: 477-488.
- [36] Hoekstra, R.F., R. Bijlsma & J. Dolman (1985) Polymorphism from environmental heterogeneity: Models are only robust if the heterozygote is close in fitness to the favoured homozygote in each environment. *Genet. Res. Camb.* 45: 299-314.
- [37] Jagers P. (1975) *Branching Processes with Biological Applications.* Wiley, London
- [38] Jacobs, F.J.A. & J.A.J. Metz (2003) On the concept of attractor in community-dynamical processes I: The case of unstructured populations. *J. Math. Biol.* 47: 222-234.
- [39] Jepsen, G.L., G.G. Simpson & E. Mayr eds. (1949) *Genetics, Paleontology and Evolution.* Princeton University Press.
- [40] Kisdi, É. & S.A.H. Geritz (in press) Evolutionary branching and speciation: Insights from few-locus models. In U. Dieckmann & J.A.J. Metz (eds.) *Elements of adaptive dynamics.* Cambridge University Press.
- [41] Kopp, M. & J. Hermisson (2006) The evolution of genetic architecture under frequency-dependent disruptive selection. *Evolution*, 60: 1537-1550.
- [42] Leimar, O. (2001) Evolutionary change and darwinian demons. *Selection* 2: 65-72.
- [43] Leimar, O. (2005) The evolution of phenotypic polymorphism: randomized strategies versus evolutionary branching. *American Naturalist* 165: 669-681.
- [44] Leimar, O. (in press) Multidimensional convergence stability and the canonical adaptive dynamics. In: Dieckmann, U. & Metz, J.A.J. (eds.) *Elements*

- of Adaptive Dynamics. Cambridge Studies in Adaptive Dynamics. Cambridge University Press.
- [45] Maynard Smith, J. (1966) Sympatric speciation. *Am. Nat.* 100: 637-650
- [46] Maynard Smith, J. & G.R. Price (1973) The logic of animal conflict. *Nature* 246: 15-18.
- [47] Mayr, E. (1982) *The growth of Biological Thought: Diversity, Evolution, and Inheritance*. Belknap Press, Cambridge, Ma.
- [48] Mayr, E. & W.B. Provine (1980) *The Evolutionary Synthesis: Perspectives on the Unification of Biology*. Harvard University Press, Cambridge, Ma.
- [49] Méléard, S & C.V. Tran (2009) Trait substitution sequence process and canonical equation for age-structured populations. *J. Math. Biol.* 58: 881-921.
- [50] Meszéna, G., M. Gyllenberg, F.J.A. Jacobs & J.A.J. Metz (2005) Dynamics of similar populations: the link between population dynamics and evolution. *Phys. Rev. Lett.* 95, 078105(4). [DOI: 10.1103/PhysRevLett.95.078105]
- [51] Metz, J.A.J. (2008) Fitness. Pp. 1599-1612 in S.E. Jørgensen & B.D. Fath (eds.) *Evolutionary Ecology*. Vol. 2 of *Encyclopedia of Ecology*. Elsevier, Oxford.
- [52] Metz J.A.J. (in press) Invasion fitness, canonical equations, and global invasion criteria for Mendelian populations, In: U Dieckmann & JAJ Metz (eds) *Elements of Adaptive Dynamics*. Cambridge University Press.
- [53] Metz, J.A.J. & O. Diekmann (1986) The dynamics of physiologically structured populations. *Lecture Notes in Biomathematics* 68. Springer, Berlin.
- [54] Metz, J.A.J., S.A.H. Geritz, G. Meszéna, F.J.A. Jacobs & J.S. van Heerwaarden (1996) Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. Pp 183-231 in: S.J. van Strien & S.M. Verduyn Lunel (eds.) *Stochastic and spatial structures of dynamical systems*. North-Holland, Amsterdam.
- [55] Metz, J.A.J. & C. de Kovel (in prep) The canonical equation of adaptive dynamics for diploid and haplo-diploid Mendelian populations.
- [56] Metz J.A.J., S.M. Mylius & O. Diekmann (2008) When does evolution optimize? *Evol. Ecol. Res.* 10: 629-654
- [57] Metz, J.A.J., R.M. Nisbet & S.A.H. Geritz (1992) How should we define "fitness" for general ecological scenarios? *TREE* 7: 198-202.
- [58] Metz, J.A.J. & A.M. de Roos (1992) The role of physiologically structured population models within a general individual-based modeling perspective. Pp. 88-111 in: D.L. DeAngelis & L.J. Gross (eds.) *Individual-based models and approaches in ecology*. Chapman & Hall, New York.
- [59] Pennings, P.S., M. Kopp, G. Meszéna, U. Dieckmann & J. Hermisson (2008) An analytically tractable model for competitive speciation. *Am. Nat.* 171: E44E71.
- [60] Rice, S.H. (1990) A geometric model for the evolution of development. *J. theor. Biol.* 143: 319-342.
- [61] Ripa J. 2009. When is sympatric speciation truly adaptive? An analysis of the joint evolution of resource utilization and assortative mating. *Evol. Ecol.* 23: 31-52.
- [62] Rodríguez-Trelles, F., R. Tarrío & F.J. Ayala (2004) Molecular clocks: whence and whither? Pp. 5-26 in P.C.J. Donoghue & M.P. Smith (eds.) *Telling the*

- evolutionary time: molecular clocks and the fossil record. CRC Press, Boca Raton.
- [63] Salazar-Ciudad, I. (2006a) On the origins of morphological disparity and its diverse developmental bases. *Bioessays* 28: 1112-22.
 - [64] Salazar-Ciudad, I. (2006b) Developmental constraints vs. variational properties: How Pattern formation can help to understand evolution and development. *J. Exp. Zool. B: Mol. Dev. Evo.* 306B: 107-125.
 - [65] Van Dooren, T.J.M. (in press) Adaptive dynamics with Mendelian genetics. In: U. Dieckmann & J.A.J. Metz (eds.) *Elements of Adaptive Dynamics*. Cambridge Studies in Adaptive Dynamics. Cambridge University Press.
 - [66] Van Doorn, G.S. & U. Dieckmann (2006) The long-term evolution of multilocus traits under frequency-dependent disruptive selection. *Evolution* 60: 2226-2238.
 - [67] Vincent, T.L. & J.S. Brown (2005) *Evolutionary game theory, natural selection, and Darwinian dynamics*. Cambridge University Press.
 - [68] Waxman, D. (2006) Fisher's geometrical model of evolutionary adaptation — Beyond spherical geometry. *J. Theor. Biol.* 241: 887-895.
 - [69] Waxman, D. (2007) Mean curvature versus normality: A comparison of two approximations of Fisher's geometrical model. *Theor. Pop. Biol.* 71: 3036.
 - [70] Waxman, D. & J.J. Welch (2005) Fisher's microscope and Haldane's ellipse. *Am. Nat.* 166: 447-457.
 - [71] Wouters, A.G. (1999) *Explanation without a cause; verklaren zonder oorzaken te geven. Questiones Infinitae*, Publications of the Zeno Institute of Philosophy, XXIX, Utrecht

J.A.J. (= Hans) Metz

Institute of Biology and Mathematical Institute, Leiden, Netherlands
& Evolution and Ecology Program, IIASA, Laxenburg, Austria
e-mail: j.a.j.metz@biology.leidenuniv.nl