

3

Macroscopic Evolutionary Dynamics

In this chapter we discuss how the maximum entropy methodology of chapter 2 can be applied to simple evolutionary processes. First we describe the formulation of the microscopic equations of motions for such simple evolutionary systems. After that, we discuss which kinds of variables might provide reasonable choices for macroscopic variables in the context of such evolutionary processes. We explicitly construct the macroscopic state spaces and dynamics on the level of the macroscopic states for several simple examples. Additionally, we discuss the qualitative behavior of finite-population dynamics which plays a central role in the rest of the thesis. Finally we discuss how, in the evolutionary context, dynamic symmetry breaking can lead to the appearance or disappearance of dimensions in macroscopic state spaces and discuss its relation to phase transitions.

3.1 Microscopic Description

The maximum entropy methods explained in chapter 2 can only be brought to bear once the microscopic dynamics of the system under study is known and formulated explicitly. For physical systems, such as a fluid or a gas, one might assume that the microscopic equations of motion are formulated in terms of elementary-particle interactions. Although these microscopic laws are known to a certain extent, there is probably no statistical physicist who would take quantum electrodynamics as a starting point for their investigations of the behavior of a gas. Instead, what one typically does is to distill out of the “real” microscopic equations of motion a mathematical model which contains all the ingredients that are thought to be of relevance. For instance, one can model a gas as a collection of spherical particles that attract each other at moderate distances and strongly repel each other at close distances. Although molecular interactions in a real gas are much more complicated than implied by such a model, the model is thought to contain all ingredients that are important for understanding the behavior of real gases.

Although the elementary-particle dynamics of a real gas might be too complicated to take as a starting point for a statistical mechanics approach, in comparison with the microscopic dynamics of “real” evolving populations it is astonishingly simple. There is simply no chance of somehow starting an analytical approach with the actual microscopic dynamics of an evolving population. The microscopic mathematical models for evolving populations that we construct below are not more than caricatures that capture and explain some typical qualitative behaviors of real evolutionary systems.

In this thesis, we restrict ourselves to cases where the microscopic state of the population is given by a list \mathcal{S} of the population’s current genotypes, together with their relative frequencies. Thus, first of all, we assume that every individual in the population experiences the same environment. This is in contrast to spatial models, for instance, where the microscopic state of a population includes the spatial distribution of concentrations of the different genotypes at different spatial locations. Additionally, we assume that the environment is constant. In environments that vary in time and space, the precise state of the environment and its evolution, are also be part of the microscopic description of an evolutionary system. Finally, we assume that the “state” of each individual in the population is simply a function of its genotype. That is, we neglect the possibility of complicated internal dynamics of the individuals affecting their reproductive success—such as in models including developmental processes. In particular, the fitness of an individual in the constant environment is a direct function of its genotype.

Since the members of the population are indistinguishable in our models, only the frequencies of the different genotypes determine the microscopic state \mathcal{S} . Once such a list \mathcal{S} is given, the state of the evolutionary system is completely specified, and the dynamics is a function of \mathcal{S} only. This kind of microscopic description is typical of that found in mathematical population genetics [39, 69]. Additionally, we view the genotypes that can occur as embedded explicitly in a genotype *space*. The idea that the genotypes are embedded in a space of possible genotypes, often the hypercube of all length- L symbol sequences over a finite alphabet \mathcal{A} , is somewhat less common in mathematical population genetics and was first stressed by Eigen [32].

It follows from these restrictions on the formal microscopic definition of an evolutionary system, that the evolutionary dynamics can be generally described as a *Markov chain* with conditional transition probabilities $\Pr(\mathcal{S}'|\mathcal{S})$ that the population at the next time step will be the “microscopic” collection \mathcal{S}' given that its current microscopic state is \mathcal{S} . See Refs. [39] and [109] for this microscopic formulation in the context of mathematical population genetics and the theory of genetic algorithms, respectively. In formulating the microscopic dynamics as a Markov chain, a few assumptions are made implicitly:

1. The dynamics takes place in discrete time steps rather than in continuous time.
2. The microscopic state spaces are *discrete* and *finite*.
3. The Markovian assumption states that the state at the next time step is only dependent on the current state of the system. In particular, it does not depend on an infinite sequence of states of the system at previous time steps. When the next state only depends on a finite set of states in the past, these can all be grouped together in a single effective state, and then the dynamics on these new states will be Markovian again.

The first assumption is not so much a restriction but more a matter of having to choose between discrete and continuous time. In most cases, the dynamics of Markov models in continuous time can be easily mapped to the dynamics of the analogous discrete-time models. In any case, the distinction between discrete and continuous time is typically not considered to be a determining factor for the qualitative dynamical behaviors of the model. Specifically, we are careful in the following to construct our models

3.1 Microscopic Description

such that the discretization of time is not a determining factor in the behavior of the model. For instance, we take the time steps small with respect to the time scale on which the phenomena of interest occur. In particular, we consider populations evolving in discrete generations. The dynamical behaviors that we observe in our models are exactly the same as the dynamical behaviors of analogous continuous-time models.

A similar argument holds with respect to the second assumption. For a constant, homogeneous environment, the microscopic state of the population is only dependent on the frequencies of occurrence of the different genotypes. Since genotype spaces are discrete and finite, so are the population state spaces. Moreover, as will become clear, the analysis presented here can be easily extended to cases where genotype spaces are not of constant size, but are allowed to grow over time—i.e. genomes may grow in length.

The third assumption entails that the dynamics of the system has only finite memory. This does not mean that the system cannot show behavior which has memory over arbitrarily long times. It is possible that an event occurring at a certain point in time has repercussions for the entire future of the system. It is only assumed that the probabilities for the different possible futures of the system are determined only by the current state of the system. That is, the microscopic equations of motion do not exhibit memory effects, but the dynamical behaviors can.

Under these assumptions, the microscopic dynamics of an evolving population can thus be described by a Markov chain. Of course, the general theory of Markov chains can be brought to bear on these systems: this generally involves manipulating the transition matrix $\Pr(S'|S)$ of microscopic transition probabilities. For any reasonable genetic representation, however, there is an enormous number of these microscopic states S and so too of their transition probabilities. For instance, for binary sequences of length L and a population of size M , the number of microscopic states is on the order of $\mathcal{O}(2^{LM})$, which is huge for reasonable sequence lengths and population sizes. This large number of microscopic states makes it almost impossible to concretely analyze the dynamics at this microscopic level. At most, one can use abstract Markov chain theory to obtain results on various kinds of asymptotic properties: such as, in the limit of infinite time, the dynamics reaches a unique fixed-point distribution over the microscopic states. Such results tend to be useless unless one is able to predict a priori how long “asymptotic” is and what this unique distribution looks like.

More practically, a full description of the dynamics on the level of microscopic states S is neither useful nor typically of interest. One is much more likely to be concerned with relatively coarse statistics of the dynamics, such as the evolution of the best and average fitness in the population or the expected waiting times for evolution to produce a genotype of a certain fitness or with certain phenotypic characteristics. It is hard to imagine that one would be interested in a precise description of the evolution of the probability distribution $\Pr(S)$ over all possible lists of genotypes S , unless this description could be directly used to predict coarser statistics of interest. However, the huge size of the space of microscopic states S is precisely what prohibits such a direct derivation of the distribution $\Pr(S)$. The result is that quantitative mathematical analysis faces the task of finding a macroscopic description of the microscopic evolutionary dynamics that is simple enough to be tractable numerically or analytically and that, moreover, facilitates predicting the quantities of interest. Additionally, one would hope that such

a macroscopic description of the dynamics would provide insight into the qualitative mechanisms by which certain observed dynamical behaviors at the macroscopic level.

3.2 Evolutionary Macrostates

Thus, we are interested in describing the dynamics of an evolving population on a relatively coarse and macroscopic level. The first step in constructing such a description is to choose a set of macroscopic variables. As already pointed out in chapter 2 there is, as of yet, no general algorithm by which to choose a “suitable” set of macroscopic variables.¹ The macroscopic description should be capable of predicting the statistics of interest, but it should also be simple enough to allow for analyzing the dynamics in terms of these variables.

As a first step, one might attempt to remove all degrees of freedom in the dynamics that do not play a role in determining the statistics of interest. For example, there might be *symmetries* in the microscopic dynamics that can be *factored out*. A symmetry consists of a set of transformations that leave the microscopic dynamics invariant. More formally, let a transformation map each microscopic state \mathcal{S} to a microscopic state $\mathcal{T}(\mathcal{S})$. If we have for all \mathcal{S} and \mathcal{S}' that $\Pr(\mathcal{S}'|\mathcal{S}) = \Pr(\mathcal{T}(\mathcal{S}')|\mathcal{T}(\mathcal{S}))$, then this transformation is a symmetry of the microscopic dynamics. It should be easy to see that we can *group* together all microstates that are related via such symmetry transformations and describe the dynamics on the level of these grouped states. For instance, if the genotypes consist of sequences over a binary alphabet and the dynamics is symmetric under flipping of the n th bit in all sequences of the population, then we can group together the pairs of states $(\mathcal{S}, \mathcal{T}(\mathcal{S}))$ into effective states and describe the dynamics on the level of these effective states. If there are groups of transformations $\mathcal{T}_1, \mathcal{T}_2$ and so on, the grouping may lead to a large reduction in the number of states in the microscopic phase space. By finding all such symmetries, the microscopic dynamics may be reduced to a minimal number of degrees of freedom. This could be potentially very helpful in analysis.

There are two problems with this formal approach however. First, the symmetries of the microscopic dynamics do not necessarily respect the statistics of interest. For instance, we might be interested in the evolution of the average fitness in the population, but the average fitness of state \mathcal{S} is not necessarily the same as the average fitness of the state $\mathcal{T}(\mathcal{S})$. Second, microscopic symmetries are typically not very abundant except in the simplest cases.

Although the microscopic dynamics may not harbor many microscopic symmetries, it may still contain many symmetries with respect to some macroscopic statistic. If we are only interested in, say, the dynamics of the average fitness in the population, we may find a set of transformations that do not form a microscopic symmetry, but that keep the dynamics of the average fitness invariant. Formally, we consider the probabilities $\Pr(f(t)|\mathcal{S})$ that, given a current population state \mathcal{S} , the average fitness $\langle f \rangle$ follows the function $f(t)$ for all future times t . That is, given the microscopic state \mathcal{S} , $\Pr(f(t)|\mathcal{S})$ denotes the probability that the average fitness will be $f(1)$ at the next time step, $f(2)$

¹One might guess that the lack of such an algorithm is simply caused by the absence of a clear definition of “suitable” in most circumstances.

at the time step after that, and so on. If two states \mathcal{S} and \mathcal{S}' , have the same probabilities $\Pr(f(t)|\mathcal{S}) = \Pr(f(t)|\mathcal{S}')$ for *all* possible average fitness futures $f(t)$, then the microscopic states \mathcal{S} and \mathcal{S}' are *equivalent* with respect to the dynamics of the average fitness. This idea of equivalencing microscopic states with respect to the probabilities of possible futures is one of the key concepts of the computational mechanics approach [21, 23, 25, 41] to natural complexity. All states that are equivalent with respect to the probability of possible fitness futures are grouped into one equivalence class C which is called a *causal state*. One may use the microscopic equations of motion $\Pr(\mathcal{S}'|\mathcal{S})$ to construct the dynamics $\Pr(C'|C)$ on the level of these causal states. The important feature to note is that it is not necessary that the collections of microscopic states that form the causal states be related to symmetries of the microscopic dynamics. The causal states encode symmetries on the level of the macroscopic statistic of interest, such as average fitness. This makes this approach much more powerful in reducing the dimensionality of the state spaces than looking for actual symmetries of the microscopic dynamics.

Still, one may typically find, as one can easily imagine, that an enormous number of (causal) states remain, even with respect to the dynamics of average fitness only. Bolder choices for the macroscopic variables are then necessary to reduce the description of the dynamics to proportions that facilitate analysis. At this point, however, we have exhausted all mathematically rigorous possibilities of reducing the complexity of the macroscopic evolutionary dynamics. If we *do* want to predict the dynamics of the average fitness, the causal states form the smallest set of states that is capable of describing the dynamics of the average fitness completely [23]. The only way of further reducing the size of the state space is by giving up some of the accuracy in predicting the dynamics of the average fitness for all possible microscopic states.

There is much to be gained by this, however. For the cases studied here, it proved possible to reduce the description to a small number of macroscopic variables and, still, capture most of the dynamical behaviors on the level of the fitness that occur in the population. The main reason for this effectiveness is that in order to get reasonably good predictions, one does not need to find a set of variables that describes the average fitness dynamics *exactly*, one only has to find a set of macroscopic variables that describes the average fitness dynamics in typical situations with a reasonable accuracy. Once a macroscopic description is chosen, there might be microscopic states \mathcal{S} for which this macroscopic description breaks down, but if such microscopic states are very unlikely to occur in practice, they will minimally influence the accuracy of the predictions.

What kinds of macroscopic variables would be suitable in the evolutionary context? As noted earlier, this question does not have a general answer, but there are some insights to guide us. On a very intuitive level, the Neo-Darwinian paradigm of biological evolution suggests a natural decomposition of the evolutionary dynamics into a “selection” part and a “genetic diversification” part. Simply stated, the selection is thought of as an ordering force that installs information about the environment into the population by letting “adapted” individuals survive and reproduce, and letting maladapted individuals die. Selection acts on the level of the phenotypes, or even more abstractly, on the degree of the individuals’ adaption to their environment. In contrast, genetic diversification is viewed as an *independent* and largely disordering force that acts on the level of the genotypes. Roughly speaking, one can argue that different genotypes with the same level of adaptation to the environment are treated symmetrically *on average* by the evol-

utionary dynamics. Selection by definition does not distinguish between equal fitness individuals. Since genetic diversification is randomizing to a certain extent, it does not distinguish between equal fitness-individuals on average.

3.2.1 Neutrality and the Macroscopic State Spaces

As noted, we consider cases where the fitness of individuals is a direct function of their genotype. In choosing a set of macroscopic variables we are guided by a single key feature of the fitness functions that are studied here. That key feature is that there are large *degeneracies* in the map from genotype to fitness. In other words, there are many more genotypes than distinct fitness values.

Additionally, these fitness functions typically give rise to *neutral subbasins*. Neutral subbasins are sets of iso-fitness genotypes that are mutually connected through paths of single genetic diversification steps—such as point mutations. The occurrence of such neutral subbasins is, to a certain extent, due to the high dimensionality of genotype spaces. The fact that each genotype has many single mutant neighbors makes it likely that at least one of its single mutant neighbors is a *neutral* neighbor—i.e. a genotype with the same fitness. Viewed in a slightly different way, the key feature of the fitness functions studied in this thesis is that they possess *local symmetries* with respect to fitness. That is, for any genotype there always are some local genetic diversification moves that leave the fitness unchanged. This ensures, in particular, that sets of iso-fitness genotypes form connected components (subbasins) under the local genetic operators. In this way, the genotype space decomposes into a relatively small set of neutral subbasins that are entangled with each other in complicated ways².




Figure 3.1: Caricature of a neutral subbasin architecture in genotype space.

²In the theory of molecular evolution, these neutral subbasins are generally referred to as *neutral networks*.

Figure 3.1 shows a caricature of such an architecture of neutral subbasins. The neutral subbasins are indicated as the larger volumes while the connections between them are indicated as tubes. The representation of the connections between the subbasins is somewhat misleading in that, given the high-dimensionality of genotype spaces, different subbasins are often nested inside other subbasins. The figure does, however, nicely convey the conceptual idea of genotype space as an architecture of entangled neutral subbasins.

In light of the above observations, it is natural to choose the neutral subbasins as macroscopic variables for studying the evolutionary dynamics. More specifically, we will take the proportions of the population in each of these neutral subbasins as macroscopic variables that describe the state of an evolving population. Assume that there are N neutral subbasins. In most cases that we consider, there is only one neutral subbasin for each fitness value. We then describe the state of the population at any time t by a fitness or neutral subbasin distribution $\vec{P}(t) = (P_1(t), \dots, P_N(t))$, where the components $P_i(t)$ denote the *proportion* of the population in each of the neutral subbasins or fitness classes i at time t . The vector $\vec{P}(t)$ constitutes the set of macroscopic variables to which we will apply the maximum entropy techniques that were discussed in chapter 2.

When there are N different neutral subbasins in genotype space, the macroscopic state vector \vec{P} has N components. Since the components P_i give the proportions of individuals in each neutral subbasin, $P_i \geq 0$, and we have the normalization condition

$$\sum_{i=1}^N P_i = 1. \quad (3.1)$$

The set of all such vectors \vec{P} forms the *macroscopic state space* Λ

$$\Lambda = \{ \vec{P} \in \mathbb{R}_+^N \mid \sum_{i=1}^N P_i = 1 \}. \quad (3.2)$$

Strictly speaking, only infinite populations are allowed anywhere in this space Λ . For an infinite population, the proportions P_i can take on any value between 0 and 1. For a finite population of size M , in contrast, the values of P_i can only be multiples of $1/M$. That is, P_i can equal $0, 1/M, 2/M$, and so on, but not any of the intermediate values. Thus, for finite populations, the state space is a discrete subset of Λ . In particular, the state space Λ_M for a population of size M is a lattice embedded in the state space Λ with a lattice spacing of $1/M$:

$$\Lambda_M = \{ \vec{P} = \frac{\vec{n}}{M}, \vec{n} \in \mathbb{N}^N \mid \sum_{i=1}^N n_i = M \}. \quad (3.3)$$

The discrete nature of the finite-population state space plays an important role in the qualitative behavior of the evolutionary dynamics for finite populations.

3.3 Infinite-Population Dynamics

First, we will construct the evolutionary dynamics on the level of the macroscopic variables in the limit of infinite, or very large, populations. This large population-size limit

is analogous to the thermodynamic limit in statistical mechanics. We start by specifying explicitly the microscopic evolutionary dynamics in this limit. For an infinite population, the microscopic states \mathcal{S} are given by density distributions over genotype space. In other words, each state \mathcal{S} is a list of the relative proportions of all possible genotypes. For infinite populations, the dynamics on the level of these states \mathcal{S} is *deterministic*. See, for instance, [32, 33] for a formulation of such deterministic dynamics for well mixed populations of self-replicating molecules.

In principle, the microscopic dynamics is determined explicitly in terms of parameters, such as selection coefficients, and mutation and crossover rates. The compound actions of selection and the genetic operators deterministically map each genotype distribution \mathcal{S} to a genotype distribution $\mathcal{S}' = g(\mathcal{S})$ at the next generation. Similarly, we denote by $g_t(\mathcal{S})$ the microscopic genotype distribution at time t , given that the population had a genotype distribution \mathcal{S} at $t = 0$. From this deterministic microscopic dynamics on the level of the genotypes, we construct the dynamics on the level of the fitness distribution \vec{P} . To this end, we will follow the maximum entropy method described in section 2.4.

We want to make a prediction for the fitness distribution $\vec{P}(t + 1)$ given that we have information about the fitness distributions $\vec{P}(\tau)$, at all previous time steps, $\tau = 0, 1, \dots, t$. Thus, our prediction for $\vec{P}(t + 1)$ is based on the information that the fitness distribution was $\vec{P}(0)$ at time 0, $\vec{P}(1)$ at time 1, and so on. Note that we assume that the fitness distribution took on these values *exactly*, as opposed to only knowing the *average* values of the fitness distribution. The maximum entropy approach then tells us to assign equal weight to all initial genotype distributions \mathcal{S} that are consistent with all “measured” macroscopic states $\vec{P}(0)$ through $\vec{P}(t)$. Denote by C the set of all initial genotype distributions that are consistent with the entire sequence of fitness distributions $\vec{P}(0)$, $\vec{P}(1)$, and so on. Then we have that the expected fitness distribution $\langle \vec{P}(t + 1) \rangle$ at time $t + 1$ is:

$$\langle \vec{P}(t + 1) \rangle = \sum_{\mathcal{S} \in C} \frac{\vec{P}[g_{t+1}(\mathcal{S})]}{|C|}, \quad (3.4)$$

where by $\vec{P}[\mathcal{S}]$ we denote the fitness distribution of the genotype distribution \mathcal{S} , and $|C|$ is the size of the set of consistent genotype distributions. Obviously, the determining component here is the set C . This set generally depends on all previous fitness distributions $\vec{P}(t)$. In other words, we cannot generally predict the fitness distribution $\vec{P}(t + 1)$ from the current fitness distribution $\vec{P}(t)$ alone.

3.3.1 Memoryless Approximation

From a mathematical point of view, we can only predict $\vec{P}(t + 1)$ from the current state $\vec{P}(t)$ if the microscopic dynamics is exactly symmetric with respect to the fitness distribution in the following sense: If two microstates have the same fitness distribution, then they always lead to equal fitness-distributions at the next time step. Formally, if for all \mathcal{S} and \mathcal{S}' with $\vec{P}[\mathcal{S}] = \vec{P}[\mathcal{S}']$ we have that $\vec{P}[g(\mathcal{S})] = \vec{P}[g(\mathcal{S}')]$. In those cases, we really only need to know the fitness distribution $\vec{P}(t)$ to predict the distribution $\vec{P}(t + 1)$. Such exact symmetries do not typically occur. However, we may still be able to

3.3 Infinite-Population Dynamics

accurately predict the dynamics on the level of fitness distributions using a memoryless approximation.

To see how accurate prediction may still occur, we denote by C_1 the set of genotype distributions \mathcal{S} that have the current fitness distribution $\vec{P}(t)$ as their fitness distribution, i.e. $\vec{P}[\mathcal{S}] = \vec{P}(t)$. Using only this current fitness distribution as our information on which to base our prediction, rather than *all* previous fitness distributions, we have

$$\langle \vec{P}(t+1) \rangle_1 = \sum_{\mathcal{S} \in C_1} \frac{\vec{P}[g(\mathcal{S})]}{|C_1|}. \quad (3.5)$$

The 1 in the subscript of the prediction $\langle \vec{P}(t+1) \rangle_1$ denotes that we have only used the current fitness distribution $\vec{P}(t)$ and the corresponding set of microscopic states C_1 in our prediction.

If, for most cases, this prediction $\langle \vec{P}(t+1) \rangle_1$ is close to the prediction $\langle \vec{P}(t+1) \rangle$ which is obtained by including all previous fitness distributions, one obtains accurate predictions for the dynamics on the level of fitness distributions by only taking the current fitness distribution into account. In other words, the values of the fitness distributions at previous times do not contain much information with respect to the future of the fitness distributions. In this thesis we use this memoryless approximation to the dynamics of the fitness distribution without explicitly attempting to prove that it leads to accurate predictions. That our memoryless approximation leads to accurate predictions simply follows from comparing these predictions with data obtained from simulations of the actual population dynamics.

In summary, at each time t we use the maximum entropy distribution over genotype distributions, conditioned on $\vec{P}(t)$ only, to predict $\vec{P}(t+1)$. The maximum entropy distribution is uniform over the set C_1 of genotype distributions that are consistent with the current fitness distribution $\vec{P}(t)$.

3.3.2 The Generation Operator

The memoryless maximum-entropy approximation to the dynamics of fitness distributions is implemented by constructing a *generation operator* \mathbf{G} that takes the current fitness distribution $\vec{P}(t)$ and maps it to the fitness distribution $\vec{P}(t+1)$ at the next time step. Formally,

$$\vec{P}(t+1) = \mathbf{G}[\vec{P}(t)]. \quad (3.6)$$

We shall focus on simple evolutionary dynamics, which only involve selection and mutation. We decompose the generation operator \mathbf{G} into a selection operator \mathbf{S} and a mutation operator \mathbf{M} that account for the effects of selection and mutation on the fitness distribution respectively.

The selection operator is easy to construct since its effects on the fitness distribution depend only on the current fitness distribution. One of the most common forms of selection is fitness-proportionate selection: the expected number of offspring that an individual with fitness f produces in the next generation is proportional to f . If we denote by f_i the fitness of the genotypes in neutral subbasin i and by P_i^{sel} the proportion

of individuals in subbasin i after selection, we have:

$$P_i^{\text{sel}} = \frac{f_i P_i}{\langle f \rangle} \equiv \left(\mathbf{S}[\vec{P}] \right)_i. \quad (3.7)$$

To calculate the effects of mutation on the fitness distribution we must explicitly use the maximum entropy method. We need to calculate the probabilities M_{ij} that a genotype in neutral subbasin j will mutate to a genotype in neutral subbasin i . As explained in the previous section, the maximum entropy distribution is uniform over all genotype distributions that are consistent with the fitness distribution. This is equivalent to assuming that a single individual in neutral subbasin j is equally likely to be any of the genotypes in neutral subbasin j . If we denote by V_j the set of genotypes in neutral subbasin j and by $T_\mu(s \rightarrow s')$ the probability that mutation transforms genotype s into genotype s' , we have

$$M_{ij} = \sum_{s \in V_j, s' \in V_i} \frac{T_\mu(s \rightarrow s')}{|V_j|}. \quad (3.8)$$

In this thesis we consider genotypes that are bit strings of some fixed length L . If we denote by $d(s, s')$ the Hamming distance between genotypes s and s' , the mutation probabilities for a uniform mutation rate μ per bit become

$$T_\mu(s \rightarrow s') = (1 - \mu)^L \left(\frac{\mu}{1 - \mu} \right)^{d(s, s')}. \quad (3.9)$$

Finally, the generation operator is the product of the selection and mutation operators. If P_j^{sel} is the proportion of the population in the neutral subbasin j after selection, then $\sum_j M_{ij} P_j^{\text{sel}}$ is the proportion of the population in the neutral subbasin i after selection and mutation. More formally, we have for the equations of motion of the fitness distribution $\vec{P}(t)$:

$$\vec{P}(t+1) = \mathbf{M} \cdot \left(\mathbf{S}[\vec{P}(t)] \right) \equiv \mathbf{G}[\vec{P}(t)]. \quad (3.10)$$

With these, the infinite-population dynamics has been explicitly constructed on the level of neutral subbasins or, equivalently, on the level of fitness distributions. For each current fitness distribution \vec{P} , acting with the generation operator \mathbf{G} gives us the fitness distribution at the next generation.

Viewed in a slightly different way, we can focus on the *change* $d\vec{P} = \mathbf{G}[\vec{P}] - \vec{P}$ over one generation when the population currently has fitness distribution \vec{P} . By considering this change $d\vec{P}$ for each point \vec{P} in the macroscopic state space, we get a sense of the “force” that is generating the flow of populations through the state space of fitness distributions.

An example of such a flow through the macroscopic state space is shown in figure 3.2. The figure illustrates the flow $d\vec{P}$ for a simple fitness function, over the space of all 2^{30} bit strings of length $L = 30$, that contains only 4 neutral subbasins. The subbasins are denoted 0, 1, 2, and 3 and have respective fitnesses $f_0 = 0$, $f_1 = 1$, $f_2 = 2$, and $f_3 = 3$. The three-dimensional state space Λ forms a simplex in 4 dimensions, by normalization equation (3.1). The component P_3 is determined by the others, i.e. $P_3 = 1 - P_0 - P_1 - P_2$. The population evolves under fitness proportionate selection and

3.3 Infinite-Population Dynamics

a uniform mutation rate $\mu = 0.005$ per bit per reproduction. For a subset of the possible fitness distributions \vec{P} in Λ the arrows show the change $d\vec{P}$ over one generation. In other words, the arrows point from \vec{P} to $\mathbf{G}[\vec{P}]$.




Figure 3.2: Fitness distribution flow $d\vec{P}$ in the state space simplex Λ for a fitness function which contains 4 neutral subbasins and for a population evolving under fitness proportionate selection and mutation. The arrows indicate $d\vec{P}$. Fixed points of the flow are shown as large balls. The grey ball corresponds to the stable, asymptotic fixed point in the interior of the simplex Λ . The white balls indicate the locations of the unstable fixed points that are outside the simplex. The latter do not represent valid populations, but nonetheless they affect the dynamics of allowed populations within the simplex by slowing down (short arrows) the flow near them. This figure was taken from [24].

Once these arrows are known for all \vec{P} , analyzing the infinite-population dynamics becomes a more or less standard problem in dynamical systems. All techniques that have been developed in dynamical systems theory [113] can then be brought to bear on the analysis of the population dynamics. For instance, we have also indicated as large balls in figure 3.2 the fixed points of the generation operator \mathbf{G} where the flow $d\vec{P} = 0$. Of course, one can then perform linear perturbation analysis to determine the stability of these fixed points. The white balls indicate the locations of the unstable hyperbolic fixed points whereas the gray ball indicates the location of the asymptotically stable fixed

point. In fact, the fixed point corresponding to the gray ball is the only fixed point that is located *inside* the state space simplex Λ . The other fixed points are located just outside this simplex. The population can therefore never reach these fixed points. However, as we explain in the next section, these fixed points still play an important role in the finite-population dynamics.

For infinite populations, all possible genotypes are always present in the population. Additionally, there is always a nonzero probability for any genotype s' to be generated by any other genotype s , i.e. $T_\mu(s \rightarrow s') > 0$. Even though some of these transition probabilities may be extremely small, in an infinite population such transitions still occur infinitely often. For finite populations of any reasonable size, most of these very unlikely mutations will not occur and this may drastically alter the dynamics. In the following section we consider how the infinite-population dynamics constructed above can be used and altered to obtain the dynamics for finite populations.

3.4 Finite Population Dynamics

As pointed out above, we mainly focus on the dynamics of finite populations of a constant size in this thesis. For populations of size M the state space Λ_M is a lattice embedded in the infinite-population state space Λ with a lattice spacing, $1/M$, that is inversely proportional to the population size. Apart from the fact that the population can only take on points of the discrete lattice Λ_M at any point in time, the largest difference between the finite-population dynamics and the infinite-population dynamics is that the dynamics is no longer *deterministic* on the level of fitness distributions. In two different realizations of the dynamics, the same fitness distribution $\vec{P}(t)$ may give rise to different future trajectories through the space of fitness distributions.

There are two ways to deal with this situation. We could restrict ourselves to predicting the *average dynamics*, averaged over many realizations of the process, together with variances and maybe even including higher moments. This is the approach taken by Shapiro, Prügel-Bennett, and Rattray in describing the dynamics of genetic algorithms [116, 117, 118, 119]. They also make use of the maximum entropy methods from statistical mechanics and focus on “fitness” as a macroscopic variable as well. More specifically, they use the average evolution of the first cumulants of the fitness distribution as macroscopic variables. This approach works well when the dynamics in each of the realizations of the dynamics fluctuates around the *average* dynamics over many realizations. That is, if the average dynamics is *typical* for any realization of the process. In this thesis, a somewhat different approach is taken.

When the current fitness distribution $\vec{P}(t)$ is known, we simply restrict ourselves to predicting the *probability distribution* of the fitness distributions $\vec{P}(t+1)$ occurring at time $t+1$. That is, we do not assume that the average dynamics is necessarily representative for the typical dynamics in each particular realization. For almost all cases studied in this thesis, the dynamics varies between realizations in such a way, that the average dynamics is indeed not representative for any of the realizations.

Thus, we want to construct the probabilities $\Pr[\vec{Q}|\vec{P}]$ that the current fitness distribution $\vec{P} \in \Lambda_M$ leads to a distribution $\vec{Q} \in \Lambda_M$ at the next generation. The probability distribution $\Pr[\vec{Q}|\vec{P}]$ generally depends on the way selection is implemented. In this

thesis, we focus on fitness-proportionate selection in which a new generation of individuals is created by selecting M times, with replacement, a random individual from the current population. The probability for each individual to be selected is proportional to its fitness. After that, all M selected individuals are mutated. The mutated copy of each selected individual is placed in the next generation. Another way one can think of implementing fitness proportionate selection is that each individual in the current population creates a large number of copies of itself as “potential offspring”. The number of potential offspring that each individual creates is proportional to its fitness. From this large pool of potential offspring, M individuals are selected at random and then mutated. These mutated individuals then form the next generation. Fitness-proportionate selection is equivalent to selection in continuous-time models where an individual’s reproduction rate is proportional to its fitness, and a global dilution flux ensures that the population remains roughly constant in size—such as in the Eigen model of molecular evolution [32, 33].

It is easy to see that for fitness-proportionate selection, each of the M offspring in the next generation has a probability P_i^{sel} to be the offspring of an individual in the neutral subbasin i , with P_i^{sel} given by equation (3.7). Individuals that are offspring of an individual in neutral subbasin j have a probability M_{ij} to occur in the neutral subbasin i after mutation has taken place. Therefore, the probability that a randomly chosen individual in the next generation is type i is given by the i th component $\mathbf{G}_i[\vec{P}]$ of the generation operator acting on the current fitness distribution \vec{P} . Finally, since each of the M individuals in the next generation are the result of *independent* selection and mutation events, it follows that the distribution \vec{Q} at the next time, is given by a multinomial sample of size M of the distribution $\mathbf{G}[\vec{P}]$. If we define $Q_i = n_i/M$ we have:

$$\Pr[\vec{Q}|\vec{P}] = M! \prod_{i=1}^N \frac{(\mathbf{G}_i[\vec{P}])^{n_i}}{n_i!}. \quad (3.11)$$

Thus, the *expected* fitness distribution at the next generation is still $\mathbf{G}[\vec{P}]$. In the limit of infinite populations, this distribution is always exactly realized. For finite populations, however, the distribution $\mathbf{G}[\vec{P}]$ is typically not exactly realized: different distributions \vec{Q} occur for different runs of the evolutionary dynamics. Moreover, It is generally not *possible* that the expected distribution $\mathbf{G}[\vec{P}]$ is realized, since the components of $\mathbf{G}[\vec{P}]$ are unlikely to be multiples of $1/M$.

The finite-population dynamics as constructed from the infinite-population dynamics $\vec{P} \rightarrow \mathbf{G}[\vec{P}]$ and the multinomial sampling over the lattice Λ_M are illustrated in figure 3.3. Allowed finite-population fitness distributions $\vec{Q} \in \Lambda_M$ are shown as the large dots. The arrow points from the current fitness distribution \vec{P} to the expected fitness distribution $\mathbf{G}[\vec{P}]$ at the next generation. This expected fitness distribution is shown as a small dot and does not typically fall on one the allowed finite-population distributions of Λ_M . The bars over the large dots indicate the multinomial distribution $\Pr[\vec{Q}|\vec{P}]$. The variance of the multinomial distribution around the expected distribution $\mathbf{G}[\vec{P}]$ is proportional to $1/M$.

In analyzing the finite-population dynamics, we in general do not attempt to “iterate” the stochastic dynamics of equation (3.11) to obtain the stochastic population dynamics




Figure 3.3: Illustration of the stochastic dynamics over one generation. The current fitness distribution \vec{P} is mapped by the generation operator \mathbf{G} to the *expected* fitness distribution $\mathbf{G}[\vec{P}]$ at the next generation, which is indicated by the small dot. The actual fitness distribution \vec{Q} at the next generation is given by a multinomial sample of size M with distribution (3.11), which is indicated by the bars over the large dots. The large dots indicate allowed finite population fitness distributions $\vec{Q} \in \Lambda_M$. Note that the expected distribution $\mathbf{G}[\vec{P}]$ is typically not located on one of the points of Λ_M .

over arbitrary lengths of time. Formally, of course, this would give a rigorous view of the different dynamical “trajectories” that a finite population may follow with more or less probability. However, such an approach is not practical simply because it does not appear to be tractable analytically. Instead, we use the infinite-population dynamics to identify where in state space the “interesting” regions are and to get a rough sense of what regions of state space are likely to be visited by the finite population. We for instance find that the finite population dynamics spends most of its time close to unstable hyperbolic fixed points of \mathbf{G} and short transition times in “tubes” connecting the regions close to these unstable hyperbolic fixed points. We then analyze the finite population dynamics more explicitly in those specific areas. In particular, we approximate the local finite-population dynamics in a region of state space using diffusion equation approximations analogous to those introduced in mathematical population genetics by Kimura [90].

3.5 Metastability and Phase Space Unfolding

Obvious candidates for the locations of interesting state space regions are the neighborhoods of the fixed points of the generation operator \mathbf{G} . It turns out, however, that for the dynamics studied in this thesis there is typically only a single fixed point of \mathbf{G} located inside the state space Λ . This fixed point gives the asymptotically stable fitness distribution towards which the population evolves in the limit of long times. Since the other fixed points of \mathbf{G} lie outside the state space, one would generally conclude that the population simply cannot reach these points. However, we also find that these fixed

3.5 Metastability and Phase Space Unfolding

points are typically located *very close* to the state space Λ . As can be seen from the small arrows in the neighborhood of the unstable fixed points (white balls) in figure 3.2, the flow can become very small in the neighborhood of these unstable fixed points.

This observation turns out to be of great importance for the qualitative dynamics of finite populations. The essential point is that a finite population can only take on fitness distributions that are points of the discrete lattice Λ_M . If the expected flow dP_i in direction i is small compared to the lattice spacing $1/M$, the population is most likely not to move in direction i . A large population can still be carried by a small flow, but if the population gets small—such that the lattice spacing $1/M$ is large with respect to the flow—the population stops moving, even if there is no fixed point locally. Only after a long time will the population make the “jump” to the next lattice point stochastically. If the flow at this lattice point is much larger, the population may then “take off”, moving rapidly away from the neighborhood of small flow. In this way, finite populations induce metastability in the absence of fixed points. The fixed points outside Λ play a prominent role, since they indicate where the flow is smallest and, thus, where metastability is likely to occur for small populations.




Figure 3.4: Illustration of the population dynamics on the level of average and best fitness (a), and fitness distributions (b) for the same fitness function and evolutionary parameters as in figure 3.2. In (b) the fitness distribution at each generation is indicated by a dot. The arrows indicate the direction of the flow of the fitness distribution over time. The times at which different metastable states are first reached are indicated as well. (Figure taken from [24])

This mechanism is illustrated in figure 3.4. The figure shows the dynamics of the average fitness and best fitness (a) and the fitness distribution (b) for a single run of the evolutionary population dynamics corresponding to the flow of figure 3.2 with a population size of $M = 250$. The evolution of average fitness in figure (a) shows four *epochs*, corresponding to time intervals of constant best fitness in the population. In figure 3.4(b) the fitness distribution at each generation is indicated by a dot. We see that up to generation $t = 20$ the fitness distribution is located in the lower left corner

($P_0 = 1$) of the simplex. Between $t = 20$ and $t = 60$ it fluctuates around an equilibrium in the lower right corner, which corresponds to one of the unstable fixed points in figure 3.2. At time $t = 60$ the fitness distribution suddenly starts to move upward and reaches a new equilibrium point in the top of the simplex around $t = 70$. This equilibrium fitness distribution corresponds to one of the unstable fixed points in figure 3.2 as well. The fitness distribution fluctuates around this point until approximately $t = 170$, at which point it starts to move downward to the asymptotically stable fixed point (the gray ball in figure 3.2). It reaches this fixed point around $t = 200$.

Note that the consecutive epochs are associated with increasing dimensionalities of their metastable fitness distributions. That is, the first metastable distribution occurs at $P_0 = 1$ and has dimensionality 0. The second metastable distribution occurs on the line $P_0 + P_1 = 1$ and therefore has dimensionality 1. That is, as the fitness distribution fluctuates around the fixed point during this epoch, it remains on the *line* $P_0 + P_1 = 1$. The third metastable distribution is located in the plane $P_0 + P_1 + P_2 = 1$, and the final asymptotically stable fitness distribution has dimensionality 3. In this way, the succession of metastable states through which a population evolves is associated with *unfolding dimensions* of the macroscopic state space.

3.5.1 Unfolding Dimensions

Intuitively, this type of qualitatively behavior is caused by the enormous variance in the relative sizes of the neutral subbasins in genotype space. Neutral subbasins that correspond to sequences of low fitness are typically large while neutral subbasins of genotypes with high fitness are small. The genotype space is dominated by genotypes of low fitness. When the evolution is seeded with individuals occurring at one or more randomly chosen genotypes, the population is most likely to contain low-fitness genotypes only. In the example of figure 3.4, for instance, the population initially only contains sequences of fitness zero. That is, the fitness distribution is $P_0 = 1$; a zero-dimensional fitness distribution. The population is located at the unstable fixed point in the lower left corner of figure 3.2. For populations that are not too large, the population remains in this corner for some period of time. The flow components dP_1 and dP_2 are not zero, but they are very small compared to $1/M$. In the language of neutral subbasins, the flows dP_1 and dP_2 are small since it is unlikely that any individual in the population will leave a mutant offspring of fitness 1 or 2 in the next generation. Of course, mutations induce the population to explore new parts of genotype space, but since the subbasin of fitness zero genotypes dominates genotype space, it generally takes many generations before an individual embarks on a sequence with fitness 1, 2, or 3. To be more precise, if $dP_1 = 0.05/M$, this can be interpreted as meaning that for a population of size M , on average 0.05 sequences of fitness 1 will be created in the next generation. It will thus take on the order of 20 generations before *one* sequence of fitness 1 is discovered.

When this has happened, the component P_1 jumps from from $P_1 = 0$ to $P_1 = 1/M$. Typically, selection then quickly expands the population of sequences with fitness 1 until an equilibrium between fitness-0 and fitness-1 sequences is established in the population. The fitness distribution is located on the line $P_0 + P_1 = 1$ near the unstable fixed point on the right in figure 3.2. This fitness distribution is 1-dimensional in the sense that it is described by the proportions of two components that sum to one. Through the

3.5 Metastability and Phase Space Unfolding

discovery of sequences of fitness 1 a new dimension has been added to the macroscopic state space. In figure 3.4, the population fluctuates around this metastable state in the time period between $t = 20$ and $t = 60$.

This scenario repeats itself. Under mutation, the population moves through the neutral subbasins with fitness 0 and 1, but since genotypes of fitness 2 are even more rare, they take a longer time to be discovered by mutation. Mutation has to move the population through most parts of the neutral subbasins in genotype space before a fitness-2 sequence is discovered. When a fitness-2 sequence is discovered it quickly spreads. The population moves into the plane $P_0 + P_1 + P_2 = 1$ and will stabilize in this plane around the location of the upper unstable fixed point in figure 3.2. This happens between $t = 60$ and $t = 70$ in figure 3.4. The dimensionality of the state space has become 2-dimensional at this point. Another dimension has been unfolded by the dynamics. Finally, when sequences of fitness 3 are discovered, the population moves to the asymptotic fixed point in the interior, indicated by the gray ball. This occurs between times $t = 170$ and $t = 200$ in figure 3.4.




Figure 3.5: Illustration of the unfolding of macroscopic state space through epochal evolutionary dynamics. While, in genotype space, the population drifts through neutral subbasins of iso-fitness sequences, the distribution of fitness or phenotypes in the population is constant. Every time a new and better adapted phenotype is discovered that spreads through the population (an innovation), a new dimension is added to the macroscopic state space.

This scenario of macroscopic state space unfolding one dimension at a time is il-

illustrated schematically by figure 3.5. The population starts out with individuals of only 1 macroscopic type which is indicated as a dot. Once individuals of a new, and better adapted phenotype have been discovered, this fixed point becomes unstable and the population moves to a new equilibrium which involves a mixture of individuals of both the old and the new macroscopic type. The equilibrium distribution is now a point on a line. Further innovations to new macroscopic types move the population into a plane, and into three dimensional space. Each time a new macroscopic type is discovered the old fixed point becomes unstable and the population moves to a fixed point that contains more independent components. In this way, the macroscopic state of the population increases its dimensionality each time a better adapted macroscopic type is discovered. Of course, this unfolding can go on to successfully higher dimensions.

The scenario of incremental unfolding of macroscopic dimensions is potentially very general. In constant selective environments, the discovery of genotypes that confer a substantially higher fitness on individuals tends to be very rare. Often, such changes coincide with the discovery of new functionality on the level of the phenotypes or new adaptations to the environment. In this sense, these *innovations* add a new degree-of-freedom to the population dynamical system on the level of macroscopic states. This is reflected formally by the unfolding of a new dimension in state space.

It is important to note here that such new macroscopic dimensions are by no means uniquely predetermined. There may be many different macroscopic types with an increased adaptive value that may be accessed from the current state of the population. Depending on which innovative type is discovered first (an essentially stochastic process) different new macroscopic types may “freeze in”. These occurrences are generally referred to as *frozen accidents*. The occurrence of frozen accidents is directly formalized in our analytical framework by the fact that at any point in time there may be many different but *mutually exclusive* dimensions that may be unfolded.

Moreover, once a new macroscopic dimension has unfolded, the stage may be set, so to speak, for further macroscopic dimensions that may now be unfolded. In other words, the potential unfolding of new macroscopic dimensions is contingent on the unfolding of previous macroscopic dimensions. This phenomenon, which we generally refer to as *historical contingency*, also fits naturally within our picture of unfolding macroscopic dimensions. The current macroscopic types that occur in the population determine what the genotypic potential for neutral variations is. This space of neutral variants determines which new macroscopic types may unfold next.

3.5.2 Unfolding and Phase Transitions

During a metastable period when the population state is located around a fixed point in the macroscopic state space, random genetic diversification mechanisms lead the population to explore phenotypically neutral variations of the current macroscopic states. This exploration of neutral variants continues until one of these variants turns out to have adaptive value. At the moment this happens, a symmetry of the dynamics is broken.

Remember that our statistical dynamics approach assumes a maximum entropy distribution over the subbasin of mutually neutral variants in genotype space. This is equivalent to assuming that the microscopic dynamics is *symmetric* with respect to all the neutral variants, as explained in section 3.3. While the population explores the space

3.5 Metastability and Phase Space Unfolding

of neutral variants, these neutral variants effectively act as if they are symmetric with respect to the dynamics on the level of the macrostates. Simply put, the dynamics on the level of the macroscopic states is not affected by the genetic exploration of neutral variants. To give a simple example: an individual's fitness may be independent of the content of letters 11 through 23 of its genotype, *except* when these 13 letters form the combination AATGGTCATACGT. In this case, the small segment of the genome changes from a disfunctional pseudogene into a gene with a novel and adaptive functionality. The dynamics *will* be effectively symmetric with respect to the content of letters 11 through 23 as long as no individual has hit the above “jackpot” combination. When this happens, the microscopic symmetry will be broken and a new macroscopic dimension will unfold.

The above discussion makes it clear that there is a strong connection between the concept of a phase transition from statistical mechanics and the unfolding of new macroscopic dimensions through evolutionary innovations: a microscopic symmetry of the dynamics is broken and a new macroscopic variable appears. The situation is, however, not entirely identical. In the statistical mechanical examples of chapter 2, a phase transition is induced by the change in an *external* control parameter. In the evolutionary case, the symmetry breaking occurs *dynamically* through a process that is endogenous to the system. Moreover, in the evolutionary process, a symmetry is based on a lack of information. The dynamics appears symmetric with respect to the different genotypic variants as long as all variants that have been explored have been of a macroscopic type already present in the current population. As long as the exploration of new genotypes through mutations only encounters neutral (or deleterious) genotypes, the dynamics on the level of macroscopic variables is invariant under this genotypic exploration—i.e. as if such genotypic variations form a symmetry with respect to the macroscopic dynamics. When the “jackpot” genetic combination is hit for the first time, the dynamics *discovers* that the symmetry was not complete: a very special genotypic combination did not give a neutral or deleterious variant, but gave something new on a macroscopic level. This transition is then accompanied by the appearance of a new “order parameter”.

