Course 11

EVOLUTIONARY DYNAMICS

Daniel S. Fisher

Lyman Laboratory, Harvard University
Cambridge, MA 02138, USA

J.-P. Bouchaud, M. Mézard and J. Dalibard, eds.
Les Houches, Session LXXXV, 2006
Complex Systems

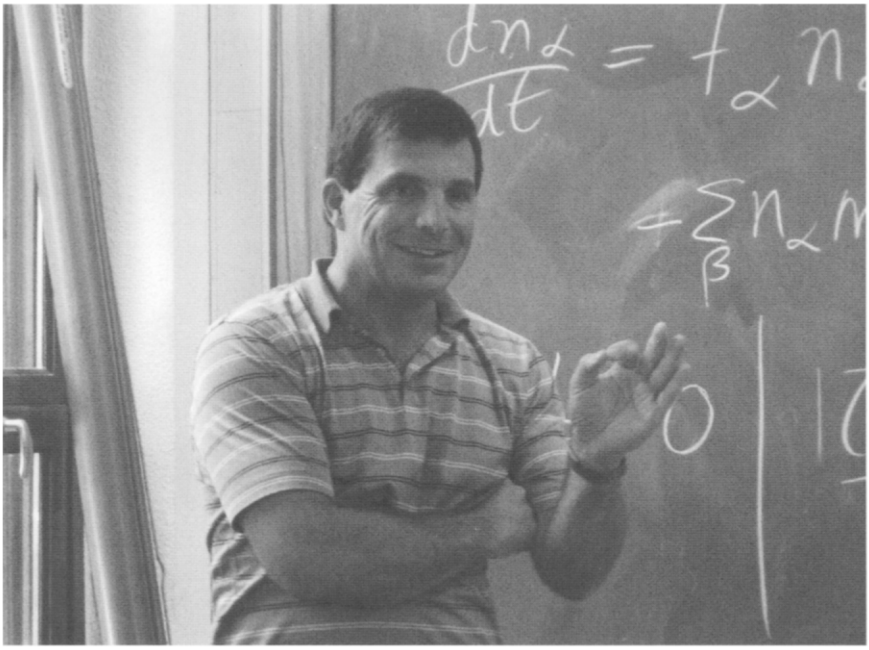

Contents

1. Introduction and Questions

The basic laws of evolution have been known for more than a century: heritable variation (Mendel), selection (Darwin's survival of the fittest) and random mutations and sexual recombination to produce the variation. The *genome* of an organism stores the primary information, but selection acts on the *phenome*: the collection of its properties, behavior, etc. With these laws, there is no basic puzzle: given *enough* time *anything* can evolve. Even a cell could arise spontaneously from an extremely rare fluctuation. But somewhere between such absurdly improbable events and small, fast, evolutionary changes in microbes that can be directly observed, are evolutionary processes that can occur on a broad spectrum of time scales: from days in the laboratory to billions of years.

The fossil record and the diversity of existing species illustrates the type of phenomes that can evolve on million and billion year time scales, and recent sequencing data provides the associated genomes. The understanding of *phylogeny* — how organisms are related to and descended from others — is impressive. Yet the understanding of the *dynamics* of evolution, and even what sets the time scales is very poor. Indeed, one could make a good case that the most fundamental puzzles about evolution are the *quantitative* ones, most basically: How is evolution of complex functions, body plans, etc. so fast? Of course, this depends on what one means by "fast": Compared to what expectations?

These lectures start with general questions to motivate and set the stage. The focus then shifts to analysis of some of the simplest aspects of evolutionary dynamics before coming back to broader issues at the end.

1.1. Difficulties

A classic problem, going back to Darwin, is the evolution of an eye. The general view among biologists emphasizes the long times involved. This is expressed quantitatively by Richard Dawkins, one of the great expositors of evolution, who says that an eye could not evolve in a thousand generations, maybe not in a million generations, but clearly could evolve in a billion generations. But where do these numbers come from? The time scales are known from the fossil record, but understanding the ages of fossils comes solely from geology, radiochemistry and physics. Thus Dawkins' and other such statements rely on knowing the an-

swer: there is no understanding from biology or evolutionary theory of the time scales. Recently, there has become available a crude measure of the time scales from biology: information from rates of neutral mutations — ones that do not change proteins – provides estimates of time scales for evolutionary history consistent to within an order of magnitude or better with dating of fossils. But such neutral changes, *ipso facto*, are not evolutionary, thus again they only provide a clock.

A way I like to phrase the primary question to evolutionary biologists is to ask what their reaction would be if they learned that the physical scientists had messed up and really life was 10^{100} years old instead of about 10^{10} years, or, for that matter, $10^{10^{10}}$ years. Would they expect that there would have been much greater diversification? Or evolution of completely new abilities of cells or whole organisms? Some have honestly said that they would not know how to react: the lack of understanding of such quantitative issues is poor enough that they would not have even rough expectations. Indeed, it is not even clear that the total time available — although this is what tends to get emphasized — is the most important quantity. Since the number of evolutionary "experiments" is roughly proportional to the total number of organisms that have ever lived, perhaps this is a better parameter. So if life on earth was a thousand times younger, but the earth a thousand times larger would as complex and diverse organisms have evolved? Or, conversely, if it were a thousand times older but a thousand times smaller (with a similar diversity of environments)? To avoid the issue of things that may have only happened once (for which talking about probabilities is seriously problematic), one can best ask these questions about the time since the first cells. Or, at a later stage, since the origins of multicellular life.

Quantitative questions about evolution on geological time scales are surely very hard to answer. But one can ask similar questions on much shorter time scales and for more modest evolutionary changes. Crudely, what combinations of parameters — and other features — determine what types of evolutionary processes can occur? This, of course, depends both on the biology of the organisms as they exist now, and on the evolutionary history which gave rise to them.

Theodosius Dobzhansky's famous dictum is that "nothing in biology makes sense except in the light of evolution". In contrast to physics, in biology a simpler explanation is no more likely to be right — unless it is simpler in evolutionary rather than functional terms. Thus evolution is the only Ockham's razor in biology — but it is hardly ever used quantitatively. The difficulty of the field and its state of development account for this. Most evolutionary theory — including quantitative modeling — deals with phenomes. Yet these are controlled by genomes which are completely different beasts. And the mapping between

genomes and phenomes is extremely complicated containing all the richness of biology and ecology. Population genetics focuses on genomic changes, but mostly either specific small changes or statistical analyses of widespread nearly-neutral variation with information on the phenotypic effects lost.

It is often said that genomic sequence data is "like reading the lab notebooks of nature". But this is an extremely misleading analogy. By the time species (or even clearly unidentifiable strains) have diverged, the number of genetic differences is so large that one cannot extract the significant changes, nor whether these arose as a long series of small phenotypic changes or by a few drastic changes via "hopeful monsters". Thus a much better analogy is that genomes are like indexes of successful textbooks — indiscriminate indexes that randomly mix useful and useless entries — with hints of the original evolutionary "ideas" very hard to extract.

The focus on phenomic evolution can give rise to major misconceptions. Often the genetic bases of phenotypic changes are only considered implicitly: as giving rise to the variability of phenotypic traits on which selection can act. A striking example is a recent estimate of the time to evolve a vertebrate eye. [1,2] This "pessimistic-at-every-step" estimate of a few hundred thousand generations sounds encouraging, but it has a fatal flaw: nowhere does the population size or the mutation rate, or indeed, the genome at all, enter the analysis. All that is considered is selection on *assumed* phenotypic variation, and this is assumed to be sufficient for changes in quantitative traits equivalent to hundreds of standard deviations. Yet this calculation is cited by Richard Dawkins as "stilling Darwin's shudder"! [3] I, for one, find this highly disturbing, especially as it comes together with the dismissal of questions about the lack of quantitative understanding of evolution as an "argument from personal incredulity".

1.2. Prospects

Thanks to the enormous advances in molecular and cell biology, the ability to observe and manipulate — genetically, chemically, and physically — organisms in the laboratory, and the explosion of DNA sequencing technology, we are presented for the first time with the opportunity to greatly expand our understanding of the dynamics of evolution. The goal is to take it from a largely-historical field — which some argue it is intrinsically [4] — to a more fully developed field of science.

1.2.1. Experiments

In order to make progress, a broad spectrum of laboratory experiments are crucial: this necessitates focusing on microbes. Bacteria — more accurately eubacteria and archaea, two very different groups that are often lumped together —

have limited morphological diversity but tremendous catabolic, metabolic, and sensory diversity and in almost any environment some can live. Concomitantly, the genetic diversity of bacteria is enormous (by some measures, the genetic diversity just among *E. coli* is greater than that among all vertebrates). Because many species of bacteria can be grown in the lab and population sizes are large, evolution of a variety of functions and ecologies of interacting species can be studied in the laboratory.

Experimental evolution has, thus far, not been a large field. But there have been a variety of experiments over the years and considerably more recently. Evolutionary experiments on multicellular organisms are primarily selective breeding: sexual recombination of genes in the existing gene pool, sometimes together with a limited number of mutational changes, can be selected on to produce remarkable variation. A the opposite extreme are viruses which have high mutation rates, can evolve rapidly, and have small enough genomes that resequencing of multiple strains can be done efficiently. Laboratory evolution of phages, viruses that infect bacteria, is a growing field. But bacteria have the greatest potential for laboratory evolution: they have short generation times and high population densities, tremendous natural diversity, and they can be manipulated genetically in many ways — including adding genes and selectively mutating parts of their genomes. And sequencing costs are now becoming low enough to re-sequence the whole genomes of evolved strains. Combinations of evolving and engineering bacteria as well as individual proteins are being used to develop useful bacterial functions, such as to clean up environmental waste or manufacture particular chemicals.

Laboratory evolution of bacteria with the goal of understanding evolutionary dynamics is beginning to burgeon — if perhaps less so thus far than one might hope. Such experiments go back to Leo Szilard's chemostats in the early 1950s. [5] And the Delbruck-Luria experiments to directly observe the effects of new mutations that occur in the lab probe the most basic evolutionary process. [6] In recent years, Richard Lenski and his collaborators have taken the lead with a spectrum of experiments, mostly with *E. coli* — molecular and cellular biologists' favorite bacterium. Their primary experiment, evolving *E. coli* to grow better in low glucose, has gone on for almost twenty years and 40,000 generations. [7] A wide range of interesting results have come from this one experiment. Some of these could not have even been found without the leaps in biological methods and knowledge that took place while the experiment was being carried out. Unfortunately, there is not the space here for even a cursory review of these or other experiments.

For interpreting and guiding laboratory experiments, theory has already played an important role. But much more is needed, both for this, and more generally to develop a broader understanding of evolution.

1.2.2. Types of theory

There are three general types of theory that are needed to understand evolutionary dynamics. First is phenomenological theory which starts from a mapping — or statistical aspects of it — between some set of genetic changes and the corresponding phenotypic changes: specifically, the effects of these on fitness in some defined set of contexts, including interactions between organisms. For such approaches, the biology is assumed given and the focus is the evolutionary dynamics that this drives. A second type of theory incorporates — and strives to inform — some understanding of aspects of the molecular and cellular biology: for example, evolution of signaling pathways or metabolic or regulatory networks. A third type of theory is abstract modeling: formulation and analysis of simple models that incorporate a few essential features with the goal of developing concrete understanding of these, and — especially crucial for evolution — how to extrapolate over broad ranges of parameters. Such models need have no connection to biology: understanding evolutionary processes in far simpler contexts — e.g. "genetic" algorithms in computer science — should enable a focus on aspects that are well beyond our ability to even model in biology. The hope, of course, is applicability — at least of the gains in understanding — beyond the specific models. Analogies of all these types of theory have played crucial roles in physics, especially condensed matter physics. In evolutionary dynamics, phenomenology has dominated, molecular-interaction based theory is just now developing, and instructive abstract modeling is almost nonexistent.

Simulations can also play a role in understanding evolutionary dynamics, especially for exploring different scenarios and processes. But there are fundamental and ubiquitous difficulties with simulations of many interacting components (here individuals, genes, etc). These are particularly problematic for evolution because of its crucial dependence on rare events and the very broad spectrum of time scales involved. In the absence of a good theoretical framework, it is impossible to extrapolate reliably from the ranges of parameters that can be studied in simulations to other much larger or smaller parameters — even if no new qualitative features arise in the more realistic regimes. Furthermore, as soon as there are more than a few features and parameters in a model, it is hard to infer which are the essential aspects on which some observed behavior depends. Thus most simulations are analogous to macroscopic evolutionary theory and yield little useful quantitative information. It is too easy to find evolution in simulations — but too hard to learn much beyond the specific model.

The focus of these lecture notes is almost exclusively phenomenological theory — as we shall see, even this rapidly becomes difficult. But we conclude with a few comments about broader needs: for abstract modeling at the one end and for incorporation of biological architecture and organization — molecular and cellular — at the other.

1.3. Numbers

Before trying to develop any quantitative theory, we start with some numbers. At this point, it is not at all clear which numbers are important: understanding this is one of the long term goals.

1.3.1. Genomes, genetic changes, and genetic differences

First, some typical sizes of genomes: the number of genes, and the size of the genome in base pairs: bp. Note that these vary substantially within groups of organisms, and in some species can be much larger than the sizes given here.

Small viruses: 10^4 bp. 10's of genes

Bacteria: 10^{6-7} bp. 500 to thousands of genes

Budding yeast: 2×10^7 bp. 6,000 genes

Humans: 2×10^9 bp. about 25,000 genes

Viruses cannot reproduce on their own — and thus are not really alive: they are basically parasitic bits of genetic material. Bacteria and archaea are *prokaryotes* which do not have nuclei. They normally reproduce asexually but can exchange DNA. [8] Budding yeast (wine-making yeast), one of the best studied laboratory organisms, is a single celled *eukaryote*: it has a nucleus and other organelles similar to all animals. Yeast can reproduce either sexually or asexually. In humans, a good fraction of the genes have a clear homolog — a common ancestor — in the yeast genome. It is remarkable that with only a factor of four more genes than yeast, all the complexity of higher animals can exist. Some of the non-protein-coding parts of genomes are involved in gene regulation which is successively more complicated going from bacteria to yeast to multicellular organisms. Yet unlike in the single celled organisms, most of the much larger genomes of vertebrates — and the even larger ones of plants — has no known function. Thus what the essential size of the information in these genomes is is unclear — quite possibly an order of magnitude smaller than their total size.

Mutation rates are remarkably small. The simplest mutations are single base mutations: changing, for example, from an A to a G. The rates of a subset of these are found directly from observations, although mutation rates can vary substantially throughout a single genome because of the local context and other factors. There are many other types of mutational changes: insertions and deletions, duplications (including of whole genes and even whole genomes), transposable elements that move around the genome, etc. Far less is known about the rates of these, but cumulatively the number that occur is in a similar range to the total number of point mutations that occur.

Point mutation rates in units of per base pair per generation and per genome per generation give a sense of the numbers. For viruses, mutation rates are very high, as much as 10^{-4} per bp, and 10^{0} per genome. Bacteria replicate their DNA with remarkable accuracy with point mutation rates reported as low as $10^{-9\frac{1}{2}}$ per bp, although an order of magnitude higher may be more typical. These correspond to rates of any error at all in the whole genome of 10^{-2} or less per cell division! Humans reproduce DNA less accurately than bacteria about 10^{-8} mutations per bp corresponding to 10^{1} mutations per genome — although the comparison is somewhat unfair as it takes many cell divisions for a human egg to produce another human egg.

The magnitude of *genetic differences* between individuals within a species and between species are also instructive. Human genomes differ from each other by about one part in a thousand, chimpanzee and human by about a part in a hundred, and human and mouse by about fifteen percent.

Bacterial genomes vary enormously. And in many bacterial phyla — the highest level classification — no organisms are known! Their existence is inferred from ribosomal RNA (rRNA) sequences which differ considerably from those of previously known phyla. [Ribosomes, the most basic machine crucial to all life, convert DNA sequences, via messenger RNA, to proteins. Their functional core is itself RNA that is coded for by genomic DNA.] As bacteria normally reproduce asexually, species are not well defined. But in recent years similarities in rRNA have been used to loosely define bacterial species: e.g., if these differ by less than about 3%. But even with a tighter definition of 1% — comparable to human-mouse differences in rRNA sequences — a single species of bacteria can have widely varying sizes and contents of their genomes, with a core of half or so genes in common, and the others completely different.

Beyond mutational changes within an individual organism's DNA, genomes can change by recombination of DNA and other mechanisms of DNA transfer between organisms. Most species of eukaryotes reproduce sexually at least some of the time: in some cases, always, in others — such as yeast — only occasionally with many generations of asexual reproduction in between. Bacteria, while normally asexual, have various mechanisms for acquiring DNA from other bacteria, both from members of the same species, and from unrelated species. [8] Little is known about the rates of these processes except in particular circumstances.

1.3.2. Populations and generations

Mutation and recombination provide the genomic variability and thus the kind of experiments nature can perform. But the number of such experiments is determined by population sizes and the numbers of generations available.

In a human body, there are of order 10^{14} cells. But a human is host to an order of magnitude larger number of bacteria. World wide, the number of cells is even more dominated by bacteria.

Total number of bacteria: Good estimates are difficult as many of the environments in which bacteria live, especially deep into the earth, are hard to sample. A recent upper-range estimate of the total number in all environments is 10^{31}. [9]

Number of bacterial generations: The time between cell divisions in bacteria varies widely. The conventional figure for *E. coli* is twenty minutes: 10^3 seconds. But this is in optimal conditions in the lab. In human guts, the turnover time, and hence average division time, is a few days, so 10^{5-6} seconds is more realistic. In other environments, divisions may be far less frequent, even many years: $> 10^8$ seconds. If we take an optimistic value of 10^5 seconds for a mean generation time, then in the few billion years since the first bacteria, the average number of bacterial generations is 10^{12} — roughly the evolutionary time in dimensionless units.

Total number of bacterial cell divisions: From the above estimates, the total number of cell divisions since the first bacteria is of order 10^{43} — although this may well be an overestimate by a few orders of magnitude. For those who like natural logs, this is about e^{100}, an easy number to remember.

Vertebrates: I do not know of estimates of the total number of vertebrates, but 10^{15} is likely an overestimate. Even with ten generations per year, this would give less than 10^{25} total vertebrate births ever. Since chimpanzee and man diverged, there have been perhaps of order 10^{12} individuals. Thus any given base pair has mutated only about a total of 10^4 times in all these in individuals together. In one lineage, only about 1% of the base pairs have mutated — 10^6 generations at a rate of 10^{-8} per generation — the observed differences between humans and chimps.

1.3.3. Explorations of genome space

The total size of genome space is enormous: even with a few megabases of DNA for bacteria: of order 10^{10^6} possible sequences. A drastic overestimate of the number of sequences ever explored by nature (since early cells) is from the total amount of DNA ever produced $10^{43} \times 10^7 \sim 10^{50}$. Assuming this was completely random, it would still provide less than the number of possible sequences of 90 nucleotides (four types). Thus only sequences of at most 30 amino acids — not much larger than a single functional domain of a protein and too small to be considered a protein on its own — could have been fully explored. And the actual extent of the exploration of sequence space is surely far less.

A better way to think of these numbers is in terms of the size of steps that can be taken in genome space. Because mutation rates are low, complex mutations, in which K changes happen in one cell division, are very rare. But beyond single mutations that increase fitness or are at worst neutral in the present environment, efficient exploration of genome space would seem to require genetic changes that involve downhill steps. If K changes are needed for the new genome to be fitter, a crude estimate of the rate of this process is K mutations in the same generation: i.e., mutation rate to the Kth power. Even with point mutation rates of 10^{-6} per bp — which bacteria can often not survive — the maximum K for a multiple-point mutation that could ever have occurred in any bacterium is about $K = 7$. In a more concrete context: for all *E. coli* in all humans ever, at rates of 10^{-9} per bp, all possible three-point mutations could have taken place, but almost none of the possible four-point mutations. Of course, these may be large underestimates because the multi-point mutations need not happen in one cell division if the intermediaries are not lethal: we discuss this point later.

How to think quantitatively about the effects of sexual recombination is less clear. If this acts primarily to move around already evolved genes, then one needs to consider recombinations at this level and consider how efficiently gene-combination space is explored. Or perhaps protein-domain space is better to consider.

1.3.4. What numbers matter?

Whether one thinks of 10^{12} generations, or even e^{100} cell divisions ever, as enormous enough numbers to obviate the need for quantitative thinking about evolutionary dynamics, depends, perhaps, on one's background. But, as we shall see, even in the simplest idealized situations, it is not known which combinations of parameters are most important for determining the evolutionary potential. My own belief is that the current lack of understanding of evolutionary dynamics is high enough that, except for knowing answers from nature, one cannot have concrete expectations. And from nature we only know about long time scales with natural mutational processes and population sizes. To understand evolution on shorter time scales in the lab, especially if genetic changes and selection can be made far more efficient than in nature, surely requires far better quantitative understanding.

2. Analysis of phenomenological models

The primary focus of these lectures is phenomenological theory with the mapping of possible genetic changes to fitness assumed. How this is determined by the biology, we only discuss briefly at the end. Due to both time limitations and

the difficulty of the problems, we only consider some of the simplest situations. And we focus on asexual reproduction which is much easier to analyze. The introductory aspects are well known [10, 11], although I hope the way they are discussed here will provide additional insights: these are needed for even the slightly more complicated situations that are discussed later.

2.1. *General formulation*

We are interested in the dynamics of interacting populations of *asexual* organisms which reproduce or die and can mutate to other genotypes. Defining the population with genome α to be n_α, we consider the simplest situation in which there is no spatial structure. For the aspects we are interested in the details of the cell division and death processes do not matter much. A simple model is to consider these to be continuous time processes with birth rate, B_α, death rate D_α, and mutation rate from genome β to genome α, $M_{\alpha\beta}$. Because of the stochasticity, we need to study the joint probability distribution of all the $\{n_\alpha\}$. This changes with time:

$$\text{Prob}\big[n_\alpha(t+dt)_\alpha(t)+1\big] = dt\left[B_\alpha n_\alpha - n_\alpha \sum_\beta M_{\beta\alpha} + \sum_\beta M_{\alpha\beta} n_\beta\right] \tag{2.1}$$

$$\text{Prob}\big[n_\alpha(t+dt)\big]_\alpha(t)-1] = dt\; D_\alpha n_\alpha. \tag{2.2}$$

The ecology enters in the dependences of the birth and death rates on the environment and the other organisms:

$$B_\alpha = B_\alpha\big(n_\alpha, \{n_\beta\}, t\big) \tag{2.3}$$

and similarly D_α, with the explicit time dependence from changes in the environment. More realistically the populations and environment also depend on spatial location and the mobility of the organisms is then also be important. But even without this, the system is complicated enough. The dynamical evolution equations are very general but, like the many body Schrodinger equation in physics, almost totally useless!

The simplest ecology is when all individuals of all species are competing for the same resources: i.e., with competition only with the total population

$$N(t) \equiv \sum_\alpha n_\alpha(t). \tag{2.4}$$

Except for initial transients, in a constant environment this competition can be taken into account by ensuring that for each birth there is a death and *vice versa* to keep N constant at the carrying capacity of the environment. This is a simple enough situation to analyze various aspects of, and will be the focus of these lectures.

2.2. *Deterministic approximation*

If, as is usually the case with microbes, the populations are large, it is tempting to approximate the dynamics as *deterministic*. In this limit the populations can be treated as continuous variables with

$$\frac{dn_\alpha}{dt} \approx \Phi_\alpha n_\alpha + \sum_\beta \left[M_{\alpha\beta} n_\beta - M_{\beta\alpha} n_\alpha \right] \tag{2.5}$$

with $\Phi_\alpha = B_\alpha - D_\alpha$ the growth (or decay) rate of population α now reflecting both birth and death processes.

With the simplest competition, N can be kept constant by taking

$$\Phi_\alpha\left(n_\alpha, \{n_\beta\}\right) = \phi_\alpha - \frac{1}{N}\sum_\beta \phi_\beta n_\beta = \phi_\alpha - \bar{\phi}(t) \tag{2.6}$$

with ϕ_α the (constant) "fitness" of organisms α in this environment — how fast a population n_α, would grow in the absence of any competition. The organisms compete only with the mean fitness, $\bar{\phi}(t)$, of the population. In this simplest ecology, $\{\phi_\alpha\}$, together with connections between genomes given by the non-zero elements of the mutation matrix, can be thought of as a "fitness landscape".

The rate of change in the mean fitness of the population is simple if the mutation rates are small enough that selection dominates. From Eq. (2.5) this is found to depend only on the variance of the fitness within the population:

$$\frac{d\bar{\phi}}{dt} \approx \mathrm{var}[\phi], \tag{2.7}$$

a general result that is valid when the effects of mutations can be neglected: it is known as the "fundamental theorem of natural selection". But a crucial question is then: What determines the variance? Its dynamics will be controlled by the third cumulant, whose dynamics is controlled by the fourth cumulant, etc. And if mutations really can be neglected, the fittest individuals will take over the population and the evolution soon stop. Nevertheless, on short time scales selection on existing variance in a population will increase the fitness at a rate proportional to the variance. This is often used to estimate evolution rates. But, disturbingly, it is often *assumed* that such variance can be maintained and continue to be selected on even though for this to happen *requires* mutations. An extreme example of the dangers of such an assumption was discussed in the introduction.

We now consider more generally the deterministic approximation to the dynamics with mutations included. With the simplest competition, $\bar{\phi}(t)$ plays the role of a Lagrange multiplier and the dynamics, Eq. (2.5) is effectively linear.

Thus at long times, the behavior will be determined by the largest eigenvalue of the fitness-plus-mutation matrix with a steady state represented by the largest eigenvector eventually being reached.

In situations in which there is an optimum genome — the highest peak in the fitness landscape — the equilibrium population will be distributed among this optimum, say $\alpha = 0$, and genomes a few mutations away from it which are nearly as fit. For example, the relative population size of a deleterious mutant, β, with fitness lower by $\delta_\beta = \phi_0 - \phi_\beta$ which is produced from the optimum genome at a rate $M_{\beta 0}$ will have an equilibrium population $n_\beta \approx n_0 M_{\beta 0}/\delta_\beta$. [Here and later, we ignore "back" mutations, here $\beta \to 0$. This can be justified in many, although by no means all, situations.] An equilibrium distribution around a fitness maximum is known as an Eigen quasispecies — Eigen rather than eigen after Manfred Eigen — and is a useful concept in situations with high mutation rates, especially viruses. But if — as is always the case in nature — fitter genomes exist but are far away from the genomes of the existing populations — i.e. requiring many mutations to reach — then the deterministic approximation will (almost always) give complete nonsense!

In practice, or course, any fitness maximum is only a local maximum: there will always be fitter genomes further away. To illustrate the effects of this, it is useful to analyze a simple example. Consider an effectively one dimensional landscape in which the genomes are labeled by an integer x with a peak at $x = 0$ with $\phi_0 = 0$, a valley of width W of depth δ, i.e. $\phi_x = -\delta$ for $x = 1, 2 \ldots W$, separating it from a higher peak at $x = K \equiv W + 1$ with $\phi_K = s > 0$. Mutations occur at rate m that take x to $x + 1$ for $x \leq W$. Again we ignore the effects of back mutations that decrease x: here this is justified if $m \ll \delta$, as we assume. Population losses from the "forward" mutations are equivalent to shifting all fitnesses by m, which we can hence also ignore. If the population of size N initially all has genome $x = 0$, the dynamics is straightforward to analyze in the deterministic approximation. One finds

$$n_x = N\left(\frac{m}{\delta}\right)^x \left[1 - e^{-\delta t}\left(\sum_{y=0}^{x-1} \frac{(\delta t)^y}{y!}\right)\right] \tag{2.8}$$

for $x \leq W$. There are two regimes: for $t \ll x/\delta$, $n_\delta \approx N e^{-\delta t}(mt)^x/x!$ while for $t \gg x/\delta$, the steady state distribution in the valley is achieved: $n_x \approx N(m/\delta)^x$. The steady state thus "propagates" at speed $1/\delta$. Analysis of the population at the peak, n_K, shows that already at times, $t > \tau_{\text{nuc}} \approx W/(\delta + s)$, n_K grows exponentially at rate s. The time τ_{nuc} is the dominant time at which the mutations to the fittest genome appear: the apparent nucleation time (determined by a balance between the exponentially small rate of mutations to the fittest population and its exponential growth). But something is worrying: τ_{nuc} does not depend

on the mutation rate! A "sweep" time $\tau_{sw} \approx \frac{W}{s} \ln[(\delta + s)/m]$ after τ_{nuc}, the fittest population will take over and become the dominant population. This does depend weakly on m but only because the fitter population at $t = \tau_{\text{nuc}}$ is of order $N[m/(\delta + s)]^K$ because of the K mutations needed to reach it from the original genome: after it has been reached, the fitter population grows exponentially.

2.3. *Failures of deterministic approximations*

Are the above results reasonable? That depends on the question. For times shorter than $\tau_{\text{nuc}} + \tau_{\text{sw}}$, the population is still dominated by the original unmutated population. Thus $\bar{\phi}(t)$ has changed very little and ignoring the competition entirely should be legitimate. In this case, averaging the stochastic dynamical equations will result in exactly the deterministic approximation which would then correctly yield $\langle n_K(t) \rangle$ and the time at which this average becomes of order N would be correctly given by the above analysis. But is this the right question? We must examine how the average arises. From the above analysis we see that at the time τ_{nuc} when the last mutations occur that dominate the later behavior of the fittest population, this average population $\langle n_K \rangle$ is a very small fraction of the total. Indeed, unless $N > [\delta/m]^K$ — which is huge for small mutation rates and broad valleys — there is, on average, much less than one fit individual at the time τ_{nuc}. But this must mean that something has gone wrong with the analysis — or that the average is a very misleading quantity. The latter is indeed the case unless the population sizes are really enormous, but how enormous will depend on the specific context. In general, as we shall see, taking averages is very misleading for evolutionary dynamics.

There is a very important lesson from the failures of the deterministic approximation. In physics, when the number of constituents (atoms, etc.) is large, thermodynamics and other approaches that focus on average quantities with small fluctuations around these are good. The basic starting points for theoretical treatments are often called mean field theories, and fluctuations can usually be added systematically to these. More generally, the enormous power of the renormalization group framework relies on —- and leads to understanding of — the simplifications that occur with large numbers and long time scales. In a broad range of contexts, this justifies the ideas of universality: dependence of many features of large systems on only a few aspects of the microscopic structure and interactions. But evolutionary dynamics has a crucial component that is very different. Mutations that arise initially in *one* individual can take over the whole population. And individuals that migrate to a new environment can give rise to new populations. Thus rare events are essential. Much of the difficulty in understanding evolutionary dynamics — even when the genome-to-phenome mapping is known — comes from the interplay between such rare individual events and the

approximately deterministic dynamics of populations once they become large. Conventional "mean-field" like approaches are thus doomed to failure.

To go beyond the dangerous deterministic approximation, and to gain some intuition we turn to the simplest evolutionary system: a fitter population being fed by stochastic mutations at a fixed rate Nm from an original population (this is equivalent to $W = 0$ in the above). But first we have to understand individual populations with stochastic dynamics. The remainder of this section is all standard material but we emphasize heuristic arguments rather than exact results to gain intuition which is needed for more complicated situations for which exact analysis is too difficult.

2.4. *Single population*

A single population without mutations is the simplest context in which to understand the stochastic dynamics of the birth and death processes. For now we ignore any limits on the population size. [The specific continuous-time model we study is somewhat different than the conventional ones with discrete generations. [10, 11] Nevertheless the key aspects of the behavior are the same; we comment later on the differences.]

Define $p_n(t)$ as the probability that there are n individuals at time t. Then

$$\frac{dp_n}{dt} = B(n-1)p_{n-1} + D(n+1)p_{n+1} - (B+D)np_n \tag{2.9}$$

so that $B - D$ is the growth rate of the mean population:

$$\frac{d\langle n\rangle}{dt} = (B-D)\langle n\rangle. \tag{2.10}$$

The variance grows proportional to $\langle n\rangle$

$$\frac{\mathrm{var}[n]}{dt} = (B+D)\langle n\rangle \tag{2.11}$$

so that the effective diffusion constant of the population distribution around its mean is $\frac{B+D}{2}\langle n\rangle$. This should be expected given the stochastic birth and death rates proportional to n so that variations around the mean after one generation are of order $\sqrt{n}$.

Before analyzing the dynamics of the distribution p_n in detail, it is instructive to try making some heuristic arguments to see if we can guess the behavior. For more complicated situations, we will have to rely on such arguments, so we had better get some practice! It is convenient to work in units of generations defined,

for convenience, by setting the death rate $D = 1$. The net growth (or decay) rate of the population per generation is defined as

$$r \equiv \frac{B - D}{D}. \tag{2.12}$$

Consider a large initial population of size $n_0 \gg 1$. There are then two characteristic time scales. The first is the obvious one: $1/|r|$ for growth or decay of the mean. The second is associated with fluctuations, referred to in the context of evolutionary dynamics as *drift*. Drift is simplest in the *neutral* case: i.e. the mean growth rate, r, is zero. As long as the deviations from n_0 are small compared to n_0, we can approximate the fluctuations around the mean by a diffusion constant n_0 so that

$$n(t) \approx n_0 \pm \mathcal{O}\left(\sqrt{n_0 t}\right) \tag{2.13}$$

valid until $t \sim n_0$ by which point the variations will have either decreased or increased in magnitude due to the accumulated changes in n. Thus the characteristic time scale for the fluctuations to substantially change the population is of order n_0 generations. If the population happens to decrease by, say, a factor of two over this time interval, then the time scale for another factor of two decrease, if it occurs, will be of order $n_0/2$, the next factor of two decrease in $n_0/4$ generations, etc. From this we can guess that the characteristic time in which the population can fluctuate away entirely — i.e. become extinct — is of order n_0. Although this argument is sloppy, we will see that it is basically correct: a population, n_0, with no mean growth or decay is likely to die out in of order n_0 generations with substantial probability. If it has not died out, then it is likely to have increased: this must be the case as the mean population is unchanged. The above argument suggests that, if it has not died out, the population will typically have increased in n_0 generations by of order a factor of two.

For a population with a small mean growth rate, $r \ll 1$, the time scale for fluctuations to drive the population to zero will be shorter than that for it to grow on average if $n_0 \ll 1/r$. Thus we expect that there is a characteristic population size of order $1/|r|$: for $n \gg 1/|r|$ the population is very likely to grow (or decay if $r < 0$), while for $n \ll 1/|r|$ its behavior is dominated by fluctuations and it might go extinct.

What if the initial population is $n_0 = 1$? Again, first consider the neutral case. In spite of this neutrality, after a time of order unity there is a substantial chance that the population will have died out. If it has not, it is likely to be larger, say of order 2. A time of order 2 later, there is again a good chance that is will have died out, if it has not done so, then it will be of order 4, etc. If it survives up to time t, then we can guess that $n(t)$ will be of order t, but with a broad distribution around

this, presumably with width of order t. What we cannot get from this argument is the probability that the population will survive to a time t. But this argument does suggest that the conditional probability, q, of surviving for a further time t, given that it has not until time t, is roughly independent of t. If the conditional survival probability from t to $2t$ is roughly independent of that from, say, $16t$ to $32t$, then the probability of survival up to a long time t is roughly the product of $\log_2 t$ factors of q. This suggests that the probability of survival is of order $e^{-\kappa \ln t} = 1/t^{\kappa}$ with some exponent $\kappa \sim \ln(1/q)$. Such behavior is reminiscent of a conventional one-dimensional diffusion process with an absorbing boundary at zero. Indeed, the behavior could have been guessed by considering the variable $y \equiv 2\sqrt{n}$ which fluctuates with an effective diffusion coefficient of unity [since $n/(\sqrt{n})^2 = \mathcal{O}(1)$]. The survival probability for a time t of a such a random walk in y is of order $1/\sqrt{t}$ suggesting $\kappa_{RW} = \frac{1}{2}$.. This argument is tempting but wrong!

But we can get the right answer for the survival probability using some additional information. We know that the average $\langle n(t) \rangle = 1$ for all t. From the above discussion we expect that $n(t)$ is either of order t, or exactly zero. This is only consistent if the probability of it being non-zero is of order $1/t$: i.e. $\kappa = 1$. Note that we have here combined heuristic arguments with an exact result (which by itself was somewhat misleading) to obtain a concrete prediction.

Let us now analyze the behavior more carefully. The evolution of the probabilities can be analyzed exactly from Eq. (2.9) by using the generating function

$$Q(z) \equiv \sum_n z^n p_n, \tag{2.14}$$

but the results are somewhat messy. For the simplest case, starting with a population $n_0 = 1$, the solution can be guessed:

$$p_0 = 1 - h(t) \tag{2.15}$$

and

$$p_n = h(t)[(1 - a(t)]\,[a(t)]^{n-1} \quad \text{for } n \neq 0 \tag{2.16}$$

with

$$a = \frac{(1+r)\left(1 - e^{-rt}\right)}{1 + r - e^{-rt}} \tag{2.17}$$

and the survival probability

$$h = \frac{r}{1 + r - e^{-rt}}. \tag{2.18}$$

If r is negative, the survival probability decays to zero exponentially as expected, but with a small coefficient that is not obvious: $h \approx (-r)e^{-(-r)t}$ for $r < 0$. For $r > 0$, the population has a non-zero chance of surviving forever: for $t \to \infty$, $h \to r/(1+r)$. *If* it survives, the population typically grows exponentially: the *conditional mean* given non-extinction is

$$\langle n|n \neq 0\rangle = \frac{1}{1-a} = \frac{(1+r)e^{rt}-1}{r}: \tag{2.19}$$

note the prefactor which is large, $\approx 1/r$, for small growth rate. We call the probability that the lineage descended from one individual survives to grow exponentially, the *establishment probability*, ϵ. For this simple dynamics, $\epsilon = r/(1+r)$. The result Eq. (2.19) should be contrasted with the mean population $\langle n\rangle = e^{rt} = h\langle n|n \neq 0\rangle$ in which there are two — misleadingly — canceling factors of r: arguing on the basis of the overall mean population is thus very dangerous! It is the (much larger) conditional mean that reflects the typical size of the population given that it survives. But even if it does survive, the distribution of n is still broad: it is exponentially distributed with, e.g. standard deviation around the mean that is of order the mean itself. But it does have a "typical" scale which is well characterized by the conditional mean.

For the neutral case, the survival probability is

$$h = \frac{1}{1+t} \tag{2.20}$$

so that as we guessed, the exponent $\kappa = 1$ determines its decay The conditional mean population if it survives is

$$\langle n|n \neq 0\rangle = 1+t \tag{2.21}$$

again confirming the heuristic argument given above. We shall see shortly why the random walk analogy, which suggested $\kappa_{RW} = \frac{1}{2}$. fails.

We can use the results for the neutral case to better understand the behavior in the growing case with small r. For a time, $\tau_{\rm est}$, of order $1/r$, the fluctuations dominate the dynamics. Through this time, the survival probability is only of order $r \sim 1/\tau_{\rm est}$. But if the population survives this long, it is likely to be of size of order $\tau_{\rm est} \sim 1/r$ (as from Eq. 2.19). Then the deterministic growth takes over and and the population starts to grow exponentially. Thus $\langle n|n \neq 0\rangle \sim \frac{1}{r}e^{rt}$ at later times. This process is the establishment of the population: it has survived long enough — for $t \sim \tau_{\rm est}$ — to become large enough, $n \sim 1/r$, to grow exponentially. Eventually, of course, such exponential growth must slow down and the population saturate. But we have been ignoring, so far, any limits on the population size: we will discuss these shortly.

Once we have the results for $n_0 = 1$, we can find the behavior for general n_0 by observing that each of the n_0 initial lineages behaves independently and $n(t)$ is the sum of the sizes of each of these lineages. Thus Prob$[n(t)|n_0]$ is the convolution of Prob$[n(t)|n(0) = 1]$ with itself n_0 times. In the neutral case for large n_0, we can check the heuristic argument that after a time of order n_0 there is a substantial chance that the population has died out. The probability that a particular one of the lineages survives is of order $1/t$. Thus the probability *none* survive is $\approx (1 - 1/t)^{n_0} \approx e^{-n_0/t}$ which indeed starts to grow substantially when $t \sim n_0$ — as expected. For $r > 0$, the survival probability for infinite time is similarly found to be large for $n_0 \gg 1/r$: this supports the basic picture that once a population reaches $\sim 1/r$ it is likely to become established and grow exponentially from then on.

At this point, we should pause and ask which of our results are general and which are specific to the detailed model of the dynamics. We have analyzed the case in which reproduction and death are continuous time processes that occur at some rates, B and D, respectively. In many situations, a more realistic model is discrete generations with a distribution of number of offspring in the next generation, and death of the parent. The neutral case corresponds to a mean number of offspring being unity. Slow growth corresponds to mean number of offspring of $1 + r$, resulting in a growth in the mean population of r per generation. But the fluctuations with this dynamics is somewhat different: in particular, it will depend on the variance in the number of offspring when the population is large, and, when it is small, can depend on the whole distribution. Nevertheless, the overall behavior is very similar: the diffusion of the population is proportional to n for large n, in the neutral case the survival probability to a long time t is proportional to $1/t$, and the probability of establishment of the population starting with one individual is, for small r, proportional to r. But the coefficients of these depend on the details of the dynamics. In particular, the establishment probability for small r is in general cr with the constant c depending on the reproduction and death processes.

2.5. *Continuous n diffusion approximation*

If the growth or decay rates are small and one is primarily interested in the dynamics of populations over many generations, the behavior simplifies. Most populations will be either relatively large or zero: as we have seen, if n is small, it is likely to become either zero or substantially larger in a relatively short time. It is thus natural to try and approximate the population as a continuous variable — but with zero playing a special role. Surprisingly, this is a good approximation. We can guess the appropriate stochastic Langevin equation, but have to be careful

what it means:

$$\frac{dn}{dt} = rn + \sqrt{n}\eta(t) \tag{2.22}$$

with η gaussian white noise with covariance $\langle \eta(t)\eta(t')\rangle = 2\delta(t-t')$. [In general, the strength of the noise will depend on details of the birth and death processes, we choose the value that corresponds to the continuous time dynamics analyzed above.] The correct interpretation of Eq. (2.22) is the Ito one with $n(t+dt)(t) + dt[rn(t) + \sqrt{n(t)}\eta(t)]$. This means that the probability density $p(n,t)$ satisfies

$$\frac{\partial p}{\partial t} = -\frac{\partial(rnp)}{\partial n} + \frac{\partial^2(np)}{\partial n^2} \tag{2.23}$$

with the n inside *both* derivatives in the diffusive-like term. As can be checked, this is necessary for the mean $\langle n\rangle$ to grow proportional to $\langle n\rangle$.

To see how the naive argument for the survival probability from diffusion of the variable $y = 2\sqrt{n}$ goes wrong, we need to contrast the Ito convention with the less-physical Stratonovich one often used by physicists. In the Stratonovich convention, the noise effectively acts at time $t + dt/2$ so that in our case its coefficient is $\sqrt{[n(t+dt)+n(t)]/2}$. The diffusion coefficient proportional to n then appears in the form: $\frac{\partial}{\partial n}\left[n\frac{\partial p}{\partial n}\right]$, But this would yield $d\langle n\rangle/dt = r\langle n\rangle + 1$ which is not correct. In the Stratonovich convention, variables can be changed straightforwardly and a diffusion equation for n with diffusion coefficient n is equivalent to that for y with diffusion coefficient unity. But with the Ito convention, one must be careful changing variables: if $q(y,t)$ is the probability density of y, then $\frac{\partial q}{\partial t} = -\frac{\partial}{\partial y}\{[\frac{ry}{2} - \frac{1}{y}]q\} + \frac{\partial^2 q}{\partial y^2}$ which corresponds to an extra term that drives y to zero. If this term had not been there, the survival probability would have decayed as $1/\sqrt{t}$. But in its presence, the pushing of y towards zero has a comparable effect to the stochastic parts for any y: it changes the survival probability to $1/t$.

A simple solution for $p(n,t)$ can be found which is analogous to that we found above for the discrete-n distribution. For the neutral case this is

$$p = \frac{1}{t^2}e^{-n/t} + \left(1 - \frac{1}{t}\right)\delta(n). \tag{2.24}$$

If we change t to $t+1$, this roughly corresponds to starting with n around 1, although we cannot take this solution — or the continuous n approximation — seriously for $n = \mathcal{O}(1)$. We leave it as an exercise to find the corresponding solution for $r \neq 0$ of the form

$$p = h\gamma e^{-\gamma n} + (1-h)\delta(n) \tag{2.25}$$

with $d\gamma/dt = -r\gamma - \gamma^2$ and the survival probability $h = \exp(-\int_0^t dt'\gamma)$ and to show that it yields the same behavior as the proper discrete analysis in the appropriate regimes of r, n, and t. This, and the general solution, can be found by Laplace transforming Eq. (2.23) in n to λ: the resulting first order partial differential equations in t and λ can be solved by method of characteristics. The Laplace transform is analogous to the generating function, $Q(z)$, for discrete n: using the latter results in a similar PDE.

2.6. *Problems with averages*

It is, perhaps, somewhat surprising that the continuous n approximation works so well. As discussed above, the problem with the deterministic approximation appears, naively, to do with the role of fractional individuals which are also present in the continuous approximation. But the latter does include fluctuations whose form, being proportional to $\sqrt{n}$, is related to the crucial role of zero populations. In the continuous approximation, in contrast to the deterministic approximation, strictly-zero populations exist and non-zero populations smaller than unity are rare. This is the primary reason that average populations are often a very poor characterization of the distribution. As we have seen in the simple situation of a single population with a slow mean growth rate, the average is completely dominated by rare instances in which the population is anomalously large — the population is usually zero — and in those rare cases it is much larger than would be guessed from the average. Thus in our earlier example of evolution in a fitness landscape with a broad valley separating two peaks, we can guess that, typically, after time $\tau_{\text{nuc}} + \tau_{\text{sw}}$ when the average of the fittest population becomes of order N, in actuality it will almost always be zero. This means that in the very rare cases in which it is non-zero it must be much larger than its average. When this occurs, the total population will have already become much large than N and we are no longer justified in ignoring the fixed total N constraint: the role of the mean fitness of the population, $\bar{\phi}(t)$, becomes very important. But since this dependence makes the birth and death rates depend on the $\{n_\beta\}$, averaging of the dynamical evolution equations can no longer be done straightforwardly: indeed, it becomes very tricky, involving all higher moments — and not very useful.

2.7. *Mutant population and selective sweeps*

Thus far, we have only analyzed simple population dynamics without selection or mutations or even limited total resources: we have allowed n to become arbitrarily large. If there are limited resources so that the birth and death rates depend on the population size, then the dynamics is already much harder to analyze exactly. A more interesting situation is two populations that are competing for total resources. If the competition results in the total population size being fixed, this

is equivalent to a particular form of population-size dependence of the birth and death rates.

Consider a population of *fixed* total size N, that consists of one population of size n — which for future purposes we will call the mutant population — with birth rate b and another — conventionally called the "wild-type" population — of size $N-n$ with birth rate B. To keep N fixed, for each birth one randomly chosen individual dies. With probability n/N this will be a mutant: thus n will increase only if the birth is a mutant and the death a wild-type, while it will decrease only if the birth is a wild-type (which occurs at overall rate $B(N-n)$) and the death a mutant. Thus the effective birth rate for the mutants is $b(1-n/N)$ and the effective death rate is $B(1-n/N)$. Even in the continuous approximation, the dynamics now becomes much harder to analyze exactly (special functions, etc. are required). But using a combination of the intuition gained for the simplest case, and a few results that are easy to derive, we can understand as much as we want.

First, consider the neutral case in which both populations have the same birth rate $b = B = 1$ (and hence also the same death rate). In this case, the individual lineages are all equivalent, we hence need consider only one of them. There are two possibilities: either this particular lineage will die out, or it will fluctuate large enough that all the others will die out: i.e. it will reach N after some time: *fixation* of this lineage. For $n \ll N$, the dynamics are similar to the case we have already studied: the probability of surviving until a time t is about $1/t$, and the typical population if it does survive is of order t. This should be valid until $n \sim (N/2)$ which will occur with probability of order $1/N$ and in time $t = \mathcal{O}(N)$. Once $n = N/2$ it has, by symmetry, a fifty percent chance of fixing. Thus we expect the probability of fixation is of order $1/N$ and, if it does fix, this will take a time of order N.

A simple argument shows that the fixation probability is exactly $1/N$: in the neutral case, each of the N original individuals has an equal chance of fixing; since only one can fix, the probability of a particular one doing so is $1/N$. This immediately implies that if the initial mutant population is n_0, the probability it fixes is n_0/N. Unless $N - n_0$ is much less than N — i.e. that the wild-type population is a small minority — the average time to fixation will be of order N. If $N - n_0$ is small, then the typical time to fixation will be of order $N - n_0$. But the average fixation time in this regime is more subtle. It is equivalent to the average time to extinction of an initially small mutant population, given that it goes extinct, which it is very likely to do. Since the probability of going extinct in a time interval dt is $dt\, n_0/t^2$ for $t \gg 1/n_0$, the mean time to extinction would be infinite but for the upper bound on n. Cutting off the integral over t at $t = \mathcal{O}(N)$ (corresponding to $n \sim N$) yields a conditional mean extinction time of $n_0 \ln N$. Thus the conditional mean fixation time for $N - n_0 \ll N$ is $(N - n_0) \ln N$

by the same argument. Note that results for mean fixation times can be found exactly from a recursion relation that connects different values of n_0, without fully analyzing the dynamics. But, as is becoming a pattern, the average fixation time is often not a very useful characterization of the behavior. Thus one learns rather less from these exact results than from the heuristic arguments!

We now consider a mutant population that has a small *selective advantage*, s, over the wild-type: i.e. s is its differential growth rate per generation of the wild-type, $s = (b - B)/B$. If $s \ll 1/N$, the mean growth will be swamped by the fluctuations and the behavior is essentially that of the neutral case. But for $s \gg 1/N$, the selection is strong enough that a sufficiently large mutant population will rise to fixation: this is called a *selective sweep*. But with a single individual initially, the mutant population is likely to die out. Yet with probability s It will reach a size of order $1/s$ and become established. If it does become established, say in a time τ_{est}, then it will grow roughly deterministically as

$$n \approx \frac{1}{s} e^{s(t-\tau_{\text{est}})} \tag{2.26}$$

until it becomes large enough to become a majority of the population: this will take a time, $\tau_{\text{sw}} \approx \frac{1}{s} \ln Ns$, the *sweep time* (although it will take twice as long for the full sweep to when the wild-type population disappears completely). After the mutant population has reached $N/2$, the wild-type population will decrease exponentially with $N - n \sim e^{-st}$ until it reaches of order $1/s$ and fluctuations take over to drive it to extinction a time of order $1/s$ later. In the intermediate regime the dynamics is close to deterministic with the "logistic" form $n \approx \nu e^{st}/[1 + \nu e^{st}/N]$ with, in the case of a single individual at time zero, the appropriate $\nu \approx e^{-s\tau_{\text{est}}}/s$. It is convenient to *define* τ_{est} in terms of the population in the deterministic regime pretending that it had been deterministic back to a time τ_{est} at which it was $1/s$. The fluctuations at early times, plus the relatively weak fluctuations for times after τ_{est}, can then all be incorporated into τ_{est} which is thus a stochastic quantity with a distribution of values of order $1/s$. The advantages of speaking in term of a stochastic τ_{est} is that this correctly matches together the fluctuation regime for small n with the deterministic regime in which the fluctuations are negligible.

The analysis we have sketched here, in particular the use of τ_{est}, is an example of a *matched asymptotic expansion* in which the existence of a small parameter, here $1/Ns$, enables different regimes to be handled separately and matched together. Here the regimes are $n \ll N$ for which the non-linearities from the dependence of the growth rate on n can be ignored, and $n \gg 1/s$ for which the fluctuations can be ignored. Understanding the scales involved — e.g. here the time scale $1/s$ in the small n regime — and being able to separate the regimes is essential for understanding more complicated situations such as those we consider

later. If the regimes and scales are understood, then for many purposes, cruder methods of matching are sufficient. In this context one could pretend that there is a strict separation of regimes: $n < 1/s$ with neutral drift only, $1/s < n < N/2$ with deterministic exponential growth of n, and $N/2 < n < N$ with deterministic exponential decay of $N - n$. The only things one would really miss are numerical factors of order unity. But since in biology all equations are wrong (in contrast to what is often pretended in physics!), such errors are likely to be less significant than those we have made in writing down the model (e.g. the particular form of the stochastic birth and death processes).

Before proceeding, we briefly consider the case of a deleterious mutant with $s = -\delta$ negative. In this case, the dynamics is stochastic until a time of order $1/\delta$ and the chances of surviving that long are of order δ. After that, even the lucky survivors, which will have reached population sizes of order $1/\delta$, will tend to die out, with the probability of survival decaying exponentially at later times. We shall need these results when we consider deleterious intermediaries.

2.8. Mutation and selection

We now turn to the combination of mutation and selection. The simplest situation is to start with a single population of size N which mutates at rate m per generation to a fitter mutant population, n, with selective advantage s which competes with the original population so as to keep the total population size fixed at N. This is just the situation analyzed above with the addition of the mutations. We focus on the case of s small but N large enough that $Ns \gg 1$.

There are several important time scales. The first is the growth time: $1/s$. The second is the time to drift from a single individual to $n \sim 1/s$: this is also of order $1/s$. The third is the sweep time: $\tau_{\text{sw}} = \frac{1}{s} \ln Ns$ which is substantially longer. The fourth is the typical time between mutations: $1/(Nm)$ which for now we will assume is long. But there is a fifth, less obvious, time scale which is more important: the time for the mutant population to become established. The basic process is simple. Each new mutant has a probability of order s of surviving drift. Eventually one of these will definitely become established. As the establishment probability is $\epsilon \approx s$, of order $1/s$ new mutants are needed for one to establish and eventually one will. The stochastic establishment time is thus $\tau_{\text{est}} \sim 1/(Nms)$ which is much longer than $1/s$ if $Nm \ll 1$. Therefore when the mutations are limiting, $\tau_{\text{sw}} \ll \tau_{\text{est}}$ and the stochastic establishment dominates the total time until the mutant population fixes.

If we tried to ignore the constant total population constraint — which should be alright at short times, — then the average mutant population would be

$$\langle n \rangle \approx \frac{Nm}{s}[e^{st} - 1]. \tag{2.27}$$

If Nm is small, $\langle n \rangle$ becomes of order $1/s$ at a time $\frac{1}{s} \ln Nm$ — but this is *not* one of the characteristic time scales. How does this happen? Consider the distribution of establishment times. As these are typically of order $1/Nms$, define a stochastic variable $\beta \equiv Nms\tau_{\text{est}}$ which is typically of order unity. After establishment, the mutant population is

$$n \approx \frac{1}{s} e^{st - s\tau_{\text{est}}} = \frac{1}{s} e^{st} e^{-\beta/Nm}. \tag{2.28}$$

For atypically small β, the probability density $\rho_{\text{est}}(\beta) \approx 1$ since the establishment is effectively a Poisson process with rate Nms. Thus $\langle e^{-\beta/Nm} \rangle \approx Nm$ being dominated by anomalously small $\beta = \mathcal{O}(Nm) \ll 1$. For $Nm \ll 1$, we see that the average reflects very atypical instances..

If the overall mutation production rate, Nm, is *large*, then $\langle n \rangle$ becomes $1/s$ when $t \approx 1/(Nms)$ still the typical establishment time. In this regime, $\tau_{\text{est}} \ll \tau_{\text{sw}}$ so that many mutant lineages become established before any sweeps to fixation. The dynamics is thus close to deterministic with the time to half-fixation of the mutant population $\frac{1}{s} \ln(s/m)$ as suggested by the deterministic approximation.

All the results we have discussed thus far are well known, although, as we have seen, the readily calculable quantities can be very misleading. Armed with some understanding and experience with heuristic arguments, we now turn to more interesting situations.

3. Acquisition of multiple beneficial mutations

In most environments there are likely to be many potentially beneficial mutations available and by acquiring multiple such mutations, the fitness can continue to increase. How fast does this happen? Surprisingly, even in the simplest possible model, this is not an easy question to answer and, in spite of much literature on the subject, the correct behavior has only been derived very recently — and then by statistical physicists. In order to make progress, the heuristic understandings of the simple situations we have already discussed are invaluable.

The simplest model of multiple beneficial mutations is a *staircase model*: a fitness "landscape" that consists of a long regular staircase with each step representing a single beneficial mutation that increases the fitness by the *same* small amount, s. The effects of the mutations are considered *additive*, so that acquiring x of them increases the fitness, ϕ, by sx. The competition is for total resources which keeps the total population fixed at N. The mean growth rate of the subpopulation with x mutations is $s[x - \bar{x}(t)]$ with $\bar{\phi} = s\bar{x}(t)$ the average fitness of the population at that time — not the average over all histories, but that of the

particular populations at time t: $\bar{x}(t)$ is thus a stochastic variable whose dynamics we are particularly interested in. Being physicists, we can consider the staircase to be infinitely long, with the beneficial mutations occurring at rate m and never being depleted. There are then just three parameters, N, m, and s.

We are interested in the mean *speed of evolution*

$$v \equiv \frac{d}{dt}\langle\bar{\phi}\rangle \tag{3.1}$$

— assuming this exists — and more generally in the dynamics of the fitness distribution within the population. How do these depend on the parameters? From the discussion of a single mutant population arising and fixing, and of the continuous n approximation discussed earlier, we can guess that with s small the parameters will enter in combinations such as Ns and Nm.

We will focus on the regime in which the there is *strong selection*

$$s \gg \frac{1}{N} \tag{3.2}$$

and the mutation rate is small relative to the selection:

$$m \ll s. \tag{3.3}$$

This is applicable in almost all contexts for single-celled organisms and for all but small populations of multicellular organisms: in very small populations, ($Ns < 1$), drift can dominate over selection for weakly beneficial mutations. The analysis we outline here was done in collaboration with Michael Desai [12]. A different regime, $m \gg s$, can obtain for viruses and for almost neutral mutations more generally; the staircase-model in this regime has been studied by Rouzine et al [13].

3.1. Deterministic approximation?

For very large population sizes, we can hope to use the deterministic approximation. This is straightforward to analyze. Starting from a single population of size N at $x = 0$, the subpopulation with x mutations, n_x, is found to be

$$n_x(t) = N\frac{[b(t)]^x e^{-b(t)}}{x!} \tag{3.4}$$

with mean number of mutations,

$$\bar{x}(t) = b(t) = \frac{m}{s}[e^{st} - 1] \tag{3.5}$$

so that the evolution speed is

$$v_{\text{det}} = mse^{st}. \tag{3.6}$$

The evolution is thus exponentially accelerating in the deterministic approximation. Concomitantly, the distribution is getting broader and broader with standard deviation of the fitness $s\sqrt{b(t)}$. But we should be highly suspicious: from our earlier discussion of evolution from one fitness peak to a higher one via an intervening valley, we can guess that the dominant mutations that give rise to the population at a large x will arise from an exponentially — or smaller — population at $x - 1$. Thus near the "front" of the fitness distribution the discreteness of the individuals is crucial: again, $\langle n_x \rangle$ is very misleading when it is less than unity. No matter how large N is, we will eventually — actually very quickly given the exponentially growing speed — run into this problem.

Thus we are faced with a situation in which there is *no* well-defined speed in the limit of large N: everything must be controlled by fluctuations (except perhaps at early times, although even then the deterministic approximation is dangerous).

3.2. Successional sweeps: modest population sizes

To proceed, we first consider the simplest regime in which the population is not very large. We can then use results we have already obtained for mutations and selective sweeps. If the total mutation production rate, Nm, is small, the dynamics is mutation limited. From an initially monoclonal population with $x = 0$, (i.e. $n_{x=0}(t = 0) = N$), mutations will occur to $x = 1$. After a stochastic establishment time, τ_{est}, of order $1/Nms$, one of the mutants will become established. It will then sweep to dominate the population in a time $\tau_{\text{sw}} \approx \ln(Ns)/s$. If

$$Nm \ll \frac{1}{\ln Ns}, \tag{3.7}$$

$\tau_{\text{sw}} \ll \tau_{\text{est}}$ so that the establishment process will dominate the time for the population to increase its mean fitness by s. Such a sweep is fast enough that it is unlikely there will be further mutations established, either from the original population or from the new population with $x = 1$, until the sweep is essentially complete and the mean fitness becomes $\bar{\phi} \approx s$. The process will then begin again with an establishment and sweep of an $x = 2$ mutant after which $\bar{\phi} \approx 2s$, etc.. In each round, the mutant offspring will succeed their parents: we thus call this process *successional mutations*. In this regime, the distribution of the subpopulations will usually be concentrated almost entirely at a single value of x.

But occasionally, in mid-sweep, it will be bimodal concentrated on two successive values of x. The average speed of evolution in the successional mutations regime is given by

$$v \approx \frac{s}{\langle \tau_{\text{est}} \rangle} \approx Nms^2 \tag{3.8}$$

since the establishment is a approximately a Poisson process with rate Nms: i.e. the probability of an establishment in the interval $(t, t + dt)$ is simply $dt\, Nmse^{-Nmst}$. [Note that more generally, as discussed earlier, the establishment probability of a single mutant will be cs rather than s with c depending on the birth and death processes. The result for the speed will thus in general be modified by a multiplicative factor of c.]

3.3. Multiple mutations in large populations

If the population is large enough that there are many new beneficial mutations each generation, $Nm \gg 1$, then the behavior is very different. For such large populations the first establishment time is much less than the sweep time for a single mutant. This means that after establishment of the first mutant, but before it can sweep, there will be further establishments of other mutations from the original population. And, more importantly, there will be new mutations from the already fitter $x = 1$ population. These double mutants will be fitter than the single mutants and can out-compete them. But before they fix, they can themselves give rise to even-fitter triple mutants, etc. Eventually, one of the mutant populations will takeover and become the majority population. But by then there will already be mutants with several more mutations that are destined to fix.

We make the *Ansatz* that there is a roughly steady state distribution of the populations $\{n_x\}$ around some mean value $\bar{x}(t)$ which advances step by step at a mean speed v, with $\bar{x}(t)$ most of the time an integer, and $n_{\bar{x}(t)} \approx N$ dominating the population at time t. At any time, there will be a fittest mutant in the population, at some $x = \bar{x}(t) + q(t)$: we define q as the *lead* of these fittest mutants. They are fitter than the average members of the population by qs and thus their population, once they have become established, will grow as e^{qst} until the mean fitness of the population increases. We assume that the fitness advantage of each mutation is small enough that $qs \ll 1$.

With small mutation rate, the lead population will become established and be growing exponentially before the next fitter mutants establish. We take time zero as the time at which the next-most fit population became established, and we label the populations by $x - \bar{x}$ rather than x. For simplicity, consider the situation at which the mean fitness increases by s at the same time zero. Then we have for

some time interval,

$$n_{q-1} \approx \frac{e^{(q-1)st}}{qs} \tag{3.9}$$

(qs rather than $(q-1)s$ in the denominator because this population became established while it had lead qs). The rate of mutations into n_q is $mn_{q-1}(t)$. As each new mutant has a chance qs of becoming established, we expect that one of them will become established when

$$\int_0^\tau mn_{q-1}(t)dt \sim \frac{1}{qs}. \tag{3.10}$$

Assuming this takes a time long compared to $1/qs$, the integral is $e^{(q-1)s\tau}/[q(q-1)s^2]$ which means that the time, τ_q, for establishment of the new lead population is

$$\tau_q \approx \frac{1}{(q-1)s} \ln(s/m). \tag{3.11}$$

Indeed, because of the exponential increase in the rate of mutations, if no mutant has become established by τ_q one is very likely to be in another, smaller, time of $1/(q-1)s$ or so later. Thus the variations in τ_q are small compared to τ_q itself by a factor of $1/\ln(s/m)$: we assume that s/m is very large so that $\ln(s/m)$ is itself a relatively large parameter.

A more detailed analysis shows that many additional similar mutant populations will be established soon after the first. Although these start growing later and are each typically substantially smaller than the first-established, collectively they decrease the effective establishment time by about $\ln(q-1)/(q-1)s$ canceling a factor of $q-1$ that would have appeared inside the logarithm had we used the first establishment alone. In practice, such corrections of order unity are comparable to the errors we are making from the approximations, in particular by assuming $\ln(s/m)$ is large. In particular, changing the establishment probability from qs to cqs to reflect different birth and death processes would result in similar small corrections.

We have derived the time for the front of the distribution to advance one step. For consistency, this must also be the time in which the mean fitness advances by s. Thus

$$v \approx \frac{s}{\langle \tau_q \rangle} \tag{3.12}$$

in terms of the mean τ_q (although, as noted, τ_q does not vary much: it is typically close to its mean).

But we now need to find q. After the lead population has become established, it will grow essentially deterministically with mutations into it from less fit populations no longer playing a substantial role. Thus the lead population proceeds from establishment by mutation to growth by selection. Under the conditions assumed above that the mean $\bar{x}$ advances by one around the same time as the front advances — i.e. the new lead becomes established — the (soon-to-be) second-fittest population grows at rate $(q-1)s$ for a time $\tau_q = s/v$, after which $\bar{x}$ advances and it grows more slowly at rate $(q-2)s$ for a further time interval s/v. Its fitness advantage over the mean decreases step by step until it becomes the dominant population. This takes a total time

$$\tau_{\rm sw} \approx (q-1)\tau_q \approx \frac{\ln(s/m)}{s}, \tag{3.13}$$

the steady-state sweep time for new mutations. During this time, the formerly-lead population has grown to a size that, for consistency, must be about N so that:

$$\frac{1}{qs}\exp\left[\frac{q(q-1)s^2}{2v}\right] \sim N \tag{3.14}$$

yielding, after plugging in for τ_q, the lead

$$q \approx \frac{2\ln(Ns)}{\ln(s/m)} \tag{3.15}$$

(ignoring the factor of q inside $\ln Ns$ which is in any case comparable to other factors we are ignoring).

The speed of evolution is obtained from the consistency condition that the lead population sweeps to become the dominant population of size N, yielding

$$v \approx s^2 \frac{2\ln(Ns) - \ln(s/m)}{\ln^2(s/m)}. \tag{3.16}$$

Several aspects of these results are important to note. Most dramatically, the dependence of v on N has gone from linear in the successional mutations regime to logarithmic at higher populations. Almost all the mutations that occur in the large populations are wasted: only those occurring near the front of the distribution — on the already fittest multiple mutants — are important, the others are destined to die out after being out-competed by the fitter ones. The primary role of the bulk of the population is to lower the mean fitness. Away from the front, mutations have little effects on the dynamics: selection completely dominates. Thus the steady state distribution and its speed are determined by the balance between mutations at the front and selection in the bulk of the distribution. This is why

the overall production rate of mutations, Nm, does not enter: m enters rather in the combination s/m and the behavior is only logarithmically dependent on the mutation rate.

Because selection dominates most of the distribution, the evolution speed is very well approximated by the the variance of the fitness, the general result mentioned earlier:

$$v = \frac{d\bar{\phi}}{dt} \approx \mathrm{var}[\phi]. \tag{3.17}$$

But as we have seen, what really matters is the front of the distribution which is many standard deviations away from the mean: the variance is not a very useful characterization of the distribution. It is the balance between the mutational dynamics at the front — a tiny fraction of the population — and selection in the bulk that determines both the steady state distribution (including the variance) and the speed. On the simple fitness staircase we are considering, the distribution of fitness is close to gaussian many standard deviations away from the mean, indeed, until the sub-populations are of order $1/qs$. [Note that this is rather unusual: distributions tend to be gaussian (if at all) only near their mean — the *central* in the "central limit theorem" — with tails of different forms.]

The evolution of the population distribution is, of course, not really steady. But the nature of the dynamics at the front implies that it does not fluctuate much. A proper analysis of the fluctuations is rather complicated, but can be done along similar lines to the above heuristic analysis. [12]

We briefly mention several minor caveats about the above results. Strictly speaking, Eq. (3.16) is only valid for particular values of, say, $\ln(Ns)$ for which the lead advances at the same time as the mean advances. In general, the dependence on q, and hence on the two logarithmic factors is more complicated. In addition, for very large N, v becomes comparable to s^2 and larger. In this regime, significant new establishments continue in the second-fittest population while establishments of the new fittest population are occurring. The mean fitness in this regime advances more smoothly, and there are several sub-populations around $\bar{x}$ that contribute to the total as the standard deviation of the distribution is larger than s. But Eq. (3.16) is a very good approximation for the whole regime with $Nm \gg 1$.

Crossover between regimes The multiple-mutations analysis is valid when $Nm \gg 1$. The border of its validity, $N \approx 1/m$, corresponds to $q = 2$: this is when the fittest, sweeping, population produces new mutants soon before it become the majority population. For not-much larger N, there thus appears a small population two steps above the mean. For smaller N, the condition for validity of the successional mutations regime is that $Nm \ll \ln(s/m)$. Between these regimes there is a crossover that is straightforward to work out. Especially

as this occurs only in a narrow range of N, we ignore this here and refer the reader to reference [12].

3.4. Beyond the staircase model

The staircase model we have been discussing is very unrealistic. It has several key simplifications: first, that all beneficial mutations have the same fitness advantage; second, that there are no deleterious mutations; third, that the effects of the beneficial mutations are additive; and fourth, that there is an infinite supply of beneficial mutations so that they are not depleted. But one of the advantages of starting with such a simple model is that additional effects can be added and understood one by one. We briefly discuss relaxing each of the assumptions and some of the additional features that can then occur.

3.4.1. Distribution of beneficial mutations

One could argue that when all mutations confer the same selective advantage and their effects are additive, a rough result for the speed could be guessed without any calculations: with important mutations acting only in the small front of the distribution where the populations are small, logarithmic dependences on population size and mutation rates could have been anticipated. As the basic scale of the speed is s^2, this suggests v will be equal to s^2 times logarithmic factors. But when there is a distribution of fitness increments, one can no longer make such an argument: what s would one use for the basic scale of the speed?

In reality, different beneficial mutations will give rise to different increases in the fitness. If there are many possible such mutations, each individually with a very low rate, these can be modeled by a distribution of mutation rates: $\mu(s)ds$ for mutations with fitness increments in the range $(s, s + ds)$. As mutations of large effect are likely to be fewer in number — of if more complicated mutational processes (see later) are involved, would occur at much lower rates — we expect $\mu(s)$ to fall-off with increasing s. Which range of s is most important for the evolution?

If the population size is sufficiently small that a mutation arises, becomes established, and fixes before others can become established, the evolution occurring via the one-by-one establishment of a succession of mutants with a distribution of strengths. Given an establishment probability of s and a fitness increment of s, the speed of evolution is

$$v \approx N \int_0^\infty s^2 \mu(s) ds \tag{3.18}$$

so that in this successional mutations regime it is the mean-square s that controls the behavior.

But when the population size is larger — roughly when the total beneficial mutation production rate, $N \int_0^\infty \mu(s)ds$, is large — then new mutants can establish before earlier ones fix. This can have large effects. For example, if a mutation, A, with s_A becomes established, its population will grow exponentially. But if, before it takes over the population, another more beneficial mutation, B, occurs with $s_B > s_A$, B can out-compete A even though it arose later. If this occurs, then the mutation A is wasted. This process is known as *clonal interference* between the different mutant lineages. [14] As mutant populations with larger s grow exponentially faster than those with smaller s, this interference suggests that the evolution will be dominated by mutations with anomalously large s. But if s is too large, the mutations will be so rare that other smaller ones will arise and fix first. Thus there should be some dominant range of s not too large and not too small. Various authors have considered this effect and tried to estimate the dominant s and the resulting behavior.

But there is a crucial complication: while mutation A may be out-competed by a stronger mutation B, the population with A could itself produce a mutation C: if $s_A + s_C > s_B$, the double-mutant AC can out-compete B. Indeed, when clonal interference occurs, such double mutants will also: if the original population can produce many new further mutants before earlier ones fix, a mutant population can produce some double mutants before it fixes. Thus whenever clonal interference is important, *multiple mutations* are likely to be important as well.

Analyzing the interplay between multiple-mutation and clonal interference is well beyond the scope of these lectures and is still only partially understood. But it can be shown that a simple approximation works rather well. For each s, acting alone, the speed, v_s, can be estimated from the constant-s model using an effective mutation rate

$$m_s \sim s\mu(s). \tag{3.19}$$

Then v_s is maximized to find the most effective s: the strength $\tilde{s}$ of these *predominant mutants* gives

$$v \approx \max_s v_s = v_{\tilde{s}}. \tag{3.20}$$

This predominant mutants approximation turns out to be surprisingly good. As long as the distribution $\mu(s)$ falls off sufficiently rapidly (faster than a simple exponential), the predominant s is roughly independent of N for large N. Although the predominant mutants approximation suggests that a broad range around $\tilde{s}$ is likely to contribute, it turns out that this is not the case: the important range of s around $\tilde{s}$ is narrow compared to $\tilde{s}$. [12] But $\tilde{s}$ does depend weakly on the overall mutation rate. The simplest to consider is increasing the rates of all types

of mutations uniformly: i.e., multiplying $\mu(s)$ by a factor of g. This results in a decrease in $\tilde{s}$ and a somewhat weaker dependence of v on g than in the simple staircase model. But this is hard to distinguish from various other effects (noted below). Thus detailed predictions of the effects of increasing mutation rates are not robust. Yet the weak logarithmic dependence on the mutation rate in the multiple mutations regime is robust, and is in striking contrast to the linear dependence in the simple successional mutations regime.

3.4.2. Deleterious mutations and optimal mutation rate

Most mutations are not beneficial: far more are deleterious. Thus the distribution of mutation rates, $\mu(s)$, should have most of its weight at negative s. In the absence of beneficial mutations there will be an equilibrium distribution of deleterious mutations present in the population: as discussed earlier, the deterministic approximation is usually good in this case. When beneficial mutations are present and there is continual evolution, the deleterious mutations still play a role, in particular altering the shape of the fitness distribution and slowing down the evolution somewhat. These effects are small unless the deleterious mutation rate is rather large. But if it is large, the mean fitness can actually decline. This phenomena, known as Muller's ratchet, occurs if the fittest genome in the population disappears because of deleterious mutations, then the next fittest, and so on. In this situation, even without beneficial mutations, the deterministic approximation fails and fluctuations — most crucially of the fittest remaining population — dominate.

We thus see that there are two competing effects of increasing the overall mutation rate: more beneficial mutations increases the speed of evolution but more deleterious mutations decreases the speed. This raises an important general question: what is the optimal overall mutation rate? It is not at all clear how to frame this question in any general situation. But in the specific context of continual evolution with a distribution of beneficial and deleterious mutations whose effects are additive, it can be addressed. Specifically: if the overall mutation rate is increased by a factor of g, by $\mu(s) \to g\mu(s)$, the speed of evolution changes. For small g, v increases with g — first linearly, then logarithmically. But for large enough g, the deleterious mutations start to dominate and v decreases. [13] This implies that there is an optimum g which depends on the population size and the distribution $\mu(s)$.

3.4.3. Interactions between mutations

The effects of different mutations are generally not additive. Specifically: the selective advantages (or losses) of a mutation A, a mutation B and the double

mutant AB are not simply related:

$$s_{AB} \neq s_A + s_B. \tag{3.21}$$

Such *interactions* between the effects of mutations is known as *epistasis*. It surely plays crucial roles for long term evolution — including speciation via of separated sexual populations. But more simply, in asexual populations, interactions between mutations would appear to invalidate the scenario for the acquisition of multiple beneficial mutations that we have been analyzing. One effect of interactions is conditionally beneficial (or deleterious) mutations for which a first beneficial mutation, A, changes whether or not a second mutation B is beneficial, or how beneficial it is. But as long as $s_{AB} > s_A$, the second mutation can add to the first, whether or not it would have been beneficial on its own. Similarly, mutation A could eliminate the potential of an otherwise-beneficial mutation B. An important example of this is mutations that are in some sense in the same class: if any of a number of different mutations results in the same phenotypic changes with a second mutation in the same class giving no further effect, then this class of mutations can be considered as one type of mutation with a rate that is the total rate of all the mutations in the class.

How do these various forms of epistasis affect the dynamics of asexual evolution via acquisition of multiple beneficial mutations ? What is needed for the scenario we have analyzed to obtain is *not* that the effects of mutations are actually additive. The crucial feature is that there are a large number of beneficial mutations always available with the *distribution* of their selective advantages, $\mu(s)$, roughly independent of earlier mutations — even though which mutations are available depends on the past history. If this is the case, then the scenario we have analyzed is a good approximation to the dynamics and our quantitative results should be applicable.

But there are other effects of interactions between mutations that do not play a role in the uphill climbs we have discussed but could nevertheless be important: *deleterious intermediaries*. For example, if s_A and s_B are both negative, but s_{AB} is positive: two mutations are then needed to produce the beneficial combination with the first step downhill in fitness. As this requires two mutations, for any particular such *two-hit* process, the rate will be very small. But a crucial question then arises: how many such potentially beneficial two-hit processes are there relative to the number of beneficial single mutations? And how small are the rates? We return to these questions later.

3.4.4. Depletion of beneficial mutations

In a constant environment with only the simplest genome-independent competition which does not change as the organisms evolve, one would expect there to be

locally optimal genomes whose fitness cannot be increased by single mutations — i.e., fitness peaks. If such a peak is reached, there will be no more beneficial mutations available — except more complicated processes with deleterious intermediaries such as the two-hit process discussed above. Before a peak is reached, the supply of beneficial mutations is likely to decrease and the rate of increase of the fitness slow down. If the effects of beneficial mutations are additive, they will simply be depleted, although how long this takes depends on whether there are a modest number of available beneficial mutations — or classes of such mutations — with relatively high rates, or many more available mutations but each with much lower rates. With interactions between mutations the situation is more complicated: if on average each beneficial mutation acquired enables one other to become available, the evolution can continue — unless an unlucky route that ends in a local fitness maximum is taken. And if two-hit processes with deleterious intermediaries can occur, the chances of becoming stuck at a local fitness maximum is far lower. Understanding the possible behaviors even with a constant environment requires far more knowledge of local fitness landscapes. And these depend on many aspects of the biological architecture as well as particulars of the past history and the type of selective pressures in the current environment.

3.5. Experiments on the speed of asexual evolution

To test the basic results of the theory outlined above for acquisition of multiple beneficial mutations, Michael Desai undertook experiments on asexual evolution of budding yeast in Andrew Murrray's lab. [15] The goal was to investigate the dependence of the evolutionary dynamics on the mutation rate and population size. The environmental conditions used were low glucose in which the yeast divided about 70% as fast as in high glucose: these conditions should cause sufficiently broad stresses that many potentially beneficial mutations are available. To eliminate all but the simplest competition, the yeast were kept in exponential growth phase at low enough densities that interactions between them were not important. The selection was simple: at regular intervals all but a small fraction of the cells were discarded, so that lineages that divide faster yield a larger fraction after the next dilution: this is roughly equivalent to keeping a fixed effective population size, N (approximately the geometrical mean of the time-dependent population size). Three different population sizes, with N spanning a factor of 2500, were used and two strains with different mutation rates, differing by about a factor of ten: the higher rate was a "mutator" strain which had one of its mutation repair mechanisms knocked out. After each of the populations evolved for some time, it was mixed in with a marked unevolved strain and the difference between their fitnesses measured by direct competition. At the end of 500 generations, some of the evolved populations were sampled and the fitnesses of

96 individuals from each measured to obtain the fitnesss distribution within the population. We will not reproduce the results here, but summarize some of the salient features. [15]

Even in the largest population of mutators, there is no sign of depletion of the supply of beneficial mutations: the rate of increase, v, of the mean fitness stays roughly constant over the 500 generations. The dependence of the evolution speed on population size is far weaker than linear — as would have been the case in the successional mutations regime — and consistent with logarithmic in N. The dependence of v on the mutation rate is also much weaker than linear. These suggest that the multiple mutations scenario does indeed apply in this experimental context. An important check on this is provided by the fitness distributions: in contrast to the usually-monoclonal, sometimes-bimodal behavior expected in the successional mutations regime, in the large populations the distributions have a single peak with a substantial width consistent with expectations.

All of the data can be well-fitted by the simple staircase model with a single value of s — about 2% per mutation -- and two values of beneficial mutation rate m differing by the expected factor of ten. With these parameters, the expected lead of the largest mutator populations is about $q = 4$ so that quadruple mutants above the mean sweep together — far faster than individual 2% mutants could on their own. Deviations of the measured speeds from predictions are within the ranges expected from fluctuations. There is some excess width to the fitness distributions of the mutator populations which can reasonably be attributed to deleterious mutations ignored in the simple model. Other scenarios for the evolution, in particular by a series of small successional mutations or by one or two large ones, are ruled out (except perhaps by appealing to fortuitous circumstances).

That the data appear to be described so well by the highly overly simplified staircase model is *a priori* surprising. But in light of the discussion above about distributions of beneficial mutations, it is reasonable to expect that there is a characteristic strength, $\tilde{s}$, of beneficial mutations that dominates the dynamics. If the probability of mutations with $s > \tilde{s}$ falls off rapidly, then the continual evolution will happen via multiple mutations of the predominant size $\tilde{s}$ with $\tilde{s}$ depending only weakly on $\ln N$. Some dependence of $\tilde{s}$ on the overall mutation rate is expected, but with a factor of ten difference of mutation rates between the normal and mutator strain this probably has little effect beyond that predicted by the simple model with fixed s. In any case, as noted above, increasing the overall mutation rate also changes other aspects of the dynamics: the distribution $\mu(s)$ will not be increased uniformly (different types of mutations are affected differently by knocking out a particular mutation repair system); the relative likelihood of two-hit mutations increases; the depletion of some of the beneficial mutations speeds up; and deleterious mutations play a larger role.

The quantitative agreement of the experiments on yeast evolution is satisfying — but we must remember that the experiments were explicitly designed to test the simplest situation beyond acquisition of a single beneficial mutation. And one might argue that the results are biologically boring: Why should one care about many small changes in a mildly stressful non-interacting environment not that different from environments the organisms have experienced before? Such an attitude may be reasonable as far as understanding *current* biological function. But for understanding evolution, it is totally unreasonable: How can one even begin to understand the dynamics of interesting evolutionary processes without understanding the simplest? And if nothing really makes sense except in light of evolution ...

We now turn to other — and surely more interesting — aspects of evolutionary dynamics. But even to frame good questions, the insights from the simplest evolutionary processes is essential.

4. Recombination and sex

Sex, in the general sense of combining some of the DNA from two organisms, surely plays a crucial role in evolution. Almost all successful groups of multicellular organisms reproduce sexually — some always, others only occasionally. The long-term benefits of sex must thus outweigh the shorter term costs, including maintaining the complicated mechanisms for sexual reproduction and the "wasting" of reproductive effort by females when they produce males rather than reproducing parthenogenetically: sex requires producing more offspring to pass on their genetic material. [16]

Bacteria, which reproduce asexually by division, nevertheless have various mechanisms for picking up DNA from other bacteria (and viruses). A well known benefit — to the bacteria! — is acquisition of antibiotic resistance by acquiring a functional group of genes — an "operon" — from another bacterium that has evolved mechanisms for dealing with similar chemicals. [8]

Many potential benefits of sex have been discussed in the literature for both single-cell and multi-cellular organisms. [16] But which are the dominant benefits in which circumstances is controversial and little understood.

A concrete benefit of sex can occur in large populations with many potentially beneficial mutations available: the situation we have been analyzing. As we have seen, with purely asexual reproduction most mutations are wasted in large populations: only those in the already-fittest individuals tend to matter and the rate of evolution increases only logarithmically with the population size. With sex, the evolution rate could be much faster. An extreme model is instructive: in each generation assume all the genes (or even parts of genes) recombine randomly so

that an individual genome includes one of each gene but with the specific allele chosen independently from the pool of variants of that gene. If the effects of mutations are roughly additive, then — at least naively — each gene will evolve separately. The rate of change of the fitness would then be the sum of the rate of changes of the fitnesses of each gene and the overall speed of evolution would continue to be mutation limited. Thus v would grow linearly with the population size even when a large numbers of different mutations are simultaneously present in the population. This suggests an enormous advantage of sexual reproduction in large populations. But what happens in more realistic models of recombination?

Various types of recombination can be studied in the simple staircase model: a large supply of beneficial mutations each with the same fitness advantage and their combined effects additive. The diversity within the population now plays a crucial role: it is not sufficient to know how many individuals there are with a given number of mutations. As they can recombine to be in the same individual, one needs to know the distributions of the *different* specific mutations among the subpopulations. There have been various efforts to analyze some forms of recombination in this model. [17–19] But it is still far from understood. For example: with a fixed recombination rate, for very large populations is the evolution speed proportional to N, to a power of $\ln N$, or something in between? Indeed, is the linear dependence on N correct even in the extreme model?

The applicability of the additive approximation for beneficial mutations is much more questionable when there is recombination. In the asexual case, the crucial feature is that the *distribution* of potentially beneficial mutations does not depend much on the past history: i.e., on which mutations have already been acquired. Thus while interactions between mutations are important, they are primarily so in a certain statistical sense: how the interactions affect the evolution of the distribution of available beneficial mutations. But as soon as there is recombination, the interaction between the specific mutations that have accumulated in two different individuals is crucial. In general, there are likely to be incompatibilities. For example, if in lineage 1 there is a mutation B_1 that was conditionally beneficial on an earlier beneficial mutation A_1 (i.e. $s_{B_1} < 0$, $s_{A_1} > 0$, and $s_{A_1B_1} > s_{A_1}$ and in lineage 2 mutation D_2 is similarly conditionally beneficial on C_2, then it is more likely than not that, e.g., the recombinations A_1D_2 and B_1C_2 are less beneficial or deleterious. Thus, in this case, sex breaks up beneficial combinations — one of its negative effects. Indeed, it is just such an effect that can be a source of speciation in separated populations: these can accumulate different mutations which are incompatible so that the populations can no longer productively mate. But for long term evolution, what matters most (as we have seen for large populations of asexuals) are the anomalously fit individuals. Thus the rare matings that produce individuals far fitter than average — for example

by combining different beneficial mutations even when their effects are simply additive — can be the most important.

In organisms such as yeast which can reproduce either sexually or asexually, sex can provide a valuable probe of asexual evolution, particularly the distributions of beneficial (and approximately neutral) mutations and interactions between these.

To even begin to understand the effects of sex requires far more knowledge of how multiple genetic changes together determine the phenome of organisms: again, this depends crucially on the biological architecture and past evolutionary history. We comment briefly on such issues — whose addressing requires going well beyond phenomenological theory — at the end of these lectures.

5. Deleterious intermediaries and combinatoric possibilities

So far, we have considered the evolutionary effects only of mutations that increase the fitness in the current environment. This corresponds to the conventional picture of asexual evolution by a series of uphill steps. But at least in large microbial populations, two-hit mutations that involve an intermediate downhill step can occur on reasonable time scales. Even if the intermediary is lethal so that the two mutations must happen the same generation, this can occur. For example, if the roughly 10^{15} bacteria in a human body divide every few days, even with point mutations rates as low as 10^{-9} per cell division, in the lifetime of a single human host a large fraction of the possible *simultaneous* two-point mutations are likely to have occurred in the most common species of the human's bacterial ecology. But this is a drastic underestimate of the rates of double mutations.

Consider a beneficial double mutation which increases the fitness by s, but with the intermediary deleterious with loss of fitness δ. We will refer to these as two-hit beneficial mutations even when the two mutations occur in different generations. If the mutation rate for the second mutation is μ, then we need to estimate the probability that a first mutation to the deleterious intermediary gives rise to a second mutation that establishes. Since the fate of each first mutation is independent, this depends only on the probability, ϵ, of a single intermediary mutant individual giving rise to an established favorable double mutant: ϵ is thus the *establishment probability* for the double mutant which plays the role of the establishment probability, $\epsilon = s$, for a single beneficial mutation.

First consider the neutral-intermediary case, $\delta = 0$. The lineage from a typical individual will die out in a few generations, so the probability that a mutation occurs and establishes from one of this lineage is of order μs. But with probability of order $1/\tau$, the lineage will survive for more than τ generations, and if it does, so, its population size, $n(\tau)$, will become of order τ. The probability that such

a lineage gives rise to an established mutant is $\mu s \int_0^\tau n(t)dt \sim \mu s \tau^2$ until this becomes of order one: lineages that survive longer than $\tau \sim 1/\sqrt{\mu s}$ are very likely to do so. Thus the probability that a single first mutant gives rise to an established beneficial double mutant is

$$\epsilon \sim \sqrt{\mu s} \tag{5.1}$$

This establishment probability is dominated by the rare lucky intermediary that lasts for an anomalously long time.

If the intermediary is weakly deleterious, then (as discussed earlier) it lineage is effectively neutral for times up to of order $1/\delta$. Thus if $\delta < \sqrt{\mu s}$ the neutral result applies. In one specific context, this condition on δ thus provides a concrete answer to the question: How neutral does a mutation need to be to be "neutral"?

The non-neutral regime obtains when

$$\delta > \sqrt{\mu s}. \tag{5.2}$$

The longest a deleterious lineage is likely to last is of order $1/\delta$, in which case it will reach a population size of order $1/\delta$ and the probability that such a lineage gives rise to an established second mutant is $\sim \mu s/\delta^2 \ll 1$. Since this happens with probability of order δ, we conclude that

$$\epsilon \sim \frac{\mu s}{\delta}. \tag{5.3}$$

[In this regime, the process is loosely analogous to quantum mechanical tunneling through an intermediate state with energy higher by δ.] Note that for $\delta \sim \sqrt{\mu s}$ this becomes the neutral result as it should.

The establishment rate for a two-hit mutation in a population of size N with first mutation rate μ_A, second mutation rate μ_B, a least-deleterious intermediary mutant, A, with $s_A = -\delta_A$, and selective advantage of the double mutant, s_{AB}, is thus

$$N\mu_A\epsilon_{AB} \sim N\mu_A \min\left[\frac{\mu_B s_{AB}}{\delta_A}, \sqrt{\mu_B s_{AB}}\right]. \tag{5.4}$$

This result should obtain as long as the population size is large enough that the maximum intermediary population size needed for the above argument is $\ll N$: this is not a stringent condition, at worst requiring $N \gg 1/\sqrt{\mu_B s_{AB}}$.

For three-hit mutations, ABC, with two deleterious intermediaries, A and AB, the above argument can be iterated using the establishment probability ϵ_{BC} for the a BC double-hit mutation from an A individual (in place of s_B in the above), to obtain the establishment probability ϵ_{ABC}. For the case with both intermediaries almost neutral — now requiring A to be extremely close to neutral

and the population sufficiently large — the overall establishment rate of the triple mutant is

$$N\mu_A \epsilon_{ABC} \sim N\mu_A \sqrt{\mu_B}[\mu_C s_{ABC}]^{\frac{1}{4}} \tag{5.5}$$

while if the intermediaries are more deleterious:

$$N\mu_A \epsilon_{ABC} \sim N\frac{\mu_A \mu_B \mu_C s_{ABC}}{\delta_A \delta_{AB}}. \tag{5.6}$$

The above results for two-hit mutations were recently obtained from exact calculations [20], but without the heuristic arguments that aid their understanding and are needed for more complicated situations.

For a K-hit beneficial mutant with deleterious intermediaries of typical strength δ that is not tiny, and mutation rates of order μ, the establishment rate is similarly

$$\frac{1}{\langle \tau_{\text{est}} \rangle} \sim Ns\mu \left[\frac{\mu}{\delta}\right]^{K-1}. \tag{5.7}$$

Note that this is just μs times the average steady-state number in the population, n_{K-1}, of the multiple mutant that gives rise to the beneficial mutant.

The problem with the deterministic approximation for this process (discussed earlier) is now apparent. Because of the exponential growth as $n_K(t) \approx e^{s(t-\tau_{\text{est}})}/s$ of the fitter final mutant once it is established at time τ_{est}, averaging $n_K(t)$ over the distribution of τ_{est} is dominated by extremely rare anomalously fast establishments for which $\tau_{\text{est}} \ll \langle \tau_{\text{est}} \rangle$. At such an establishment time which dominates the average, the actual n_K is almost always zero, but in the extremely rare cases when it is non-zero, this quickly gives rise to a population that dominates the average and rises to be much larger than N while the average is still small and the chance of establishment still tiny.

We have found that he rate for any particular multi-hit mutation is very low, even with weakly deleterious intermediaries. But this brings us to a crucial question: how many potentially beneficial K-hit mutations are there? For point mutations, in a genome of length G there are of order G^K K-tuples. But surely most of these are unlikely to be beneficial. Nevertheless, it is reasonable to expect that the number of potentially beneficial K-hit processes grows exponentially in K, say as Q^K with Q a large number. This means that that the *total* rate for K-hit processes is proportional to

$$Ns\left(\frac{Qm}{\delta}\right)^K \tag{5.8}$$

a product of a large and a small number to a power. Whether the product Qm/δ is a large or a small number is one measure of whether the mutation rate should

be considered "large" or "small" in this environment. Loosely speaking, this determines whether the exploration of genome space in an evolving population is local or far-reaching: surely a crucial question.

Even for single-point beneficial mutations in microbes, how many there are in a typical environment with broad stresses is not known — although the experiments on yeast discussed above and other experiments on *E. coli* give some indications. About multi-hit possibilities, nothing is known. Phenomenological analysis is useful for raising such questions and considering their potential consequences. But, once again, one cannot even hope to answer them without far more understanding of the biology and of evolutionary histories.

6. Beyond the simplest questions

In these lectures we have focused exclusively on understanding evolutionary dynamics when the mapping between the genome and phenome is given. And even that only in constant environments in which the fitness is a single quantity which depends on specific aspects of the phenome and is hence a function of only the organism's own genome. We have only considered the simplest interactions between organisms: competition for total resources which implicitly also assumes no spatial structure of the populations. We have seen that even with these gross simplifications the evolutionary dynamics can be subtle. And, as soon as there is sexual recombination, very little is understood.

We end by discussing briefly three general directions in which far more theory is needed. For all but the last, laboratory experiments and close interactions between theory and experiments are essential.

6.1. Space, time and ecology

Most theoretical studies of population dynamics focus on phenotypes with genetic variability assumed. This is already an enormously rich subject even if the physical environment is constant. But once mutational processes are included, very little is understood.

Temporal variations of environments are surely crucial for long term evolution: the many different "tasks" organisms face — as individuals during their lifetimes and as populations on longer time scales — mean that their "fitness" is a poor concept. One has to, at a minimum, consider multiple aspects of fitness: most simply, fitnesses in different contexts. As far as one type of organism is concerned, time-dependence of the physical environment and of the biological environment are similar (although the stresses they cause may be very different). And once genetic changes are considered, little is known: even with only two distinct environments, the interplay between the time scales for the environmental

variations and the genetic changes on evolutionary dynamics is only beginning to be explored. Of course, organisms feed back and change their own environments. Once there are several types of organisms, this gives rise to ecologies.

Simple ecological interactions between organisms which depend on the species or strains involved can give rise to stable coexistence, to oscillations, and to chaos (studies of these phenomena by Robert May and others played an important role in the developments of understanding of chaotic dynamics). The interplay between these effects and genetic changes — i.e. the evolution of simple ecologies — has been little explored and is a ripe for both experiments — some underway — and theory.

Another essential complication is spatial variation and mobility of populations. Even in the simplest models with organisms that are phenotypically identical but differ by some neutral mutations (most commonly, mutations in protein coding regions that do not change the amino acids because of the redundancies in the three-nucleotide to amino-acid genetic code), the spatial dynamics of populations is interesting and subtle. In recent years, such neutral genetic differences have been used to track human migrations. But even some of the simplest questions in the simplest models with stochastic spatial motion are not yet answered. Again, once genetic changes occur, additional complications arise. And of course, phenotypic variability brings in the full richness of evolution.

Even within the simplest model of acquisition of beneficial mutations that we have discussed, the interplay between sexual recombination and spatial variation is essential to understand. When mating is within separate fractions of the population on short time scales but there is mixing of the populations on longer time scales, the evolution can be very different than in fully mixed populations. And, of course with interactions between mutations...

6.2. Biological architecture

To even begin to address questions about how interactions between genetic changes combine to give phenotypic changes, one needs to understand many aspects of biological organization of cells and organisms.

There are two extreme caricatures of the architecture and functioning of a cell. One is that each protein (and its regulation) has a well defined function or functions and that these are grouped into modules which are themselves linked together to perform higher level functions, etc. With this as a paradigm, molecular and cell biology (aided recently by genomics) have established many connections between genes (and gene regulation) and phenotypic traits. But these successes have given rise to a bias towards thinking of evolutionary processes in too restrictive terms.

The opposite extreme is a holistic network of multiple interactions between proteins and other components. If this were the structure, changes of almost any part would affect much of the rest — and therefore multiple aspects of cell behavior. This scenario is closer to the macroscopic paradigm of evolution via selection on quantitative traits that are affected by many genes. In this paradigm, most genetic changes have side-effects — pleiotropy. And microscopically many genes have multiple functions at present, different functions in the past, and, potentially, new functions in the future.

These caricatures represent very different views of genotype to phenotype mappings. How does the extent to which each of these is true affect evolutionary potential — the *evolvability*? And how does it affect evolutionary dynamics? Conversely, how does evolutionary history affect the extent to which these caricatures represent reality? Even rough answers to these are needed to address some of the questions we have raised earlier about interactions between mutations — most crucially impacts of changes on each other and thus the combinatoric possibilities.

An advantage of modular architecture is that changes in one module are less likely to result in deleterious effects on others and thus incompatibilities between different beneficial mutations. But the network picture suggests enormously more combinatorial possibilities of changes, even if few of these are beneficial. Laboratory evolution experiments together with genetic methods and resequencing to track down changes are just beginning to start addressing these questions: for now in a small number of specific contexts, but in the future potentially in a wide enough spectrum of evolutionary contexts that general lessons can be learned. Better theoretical understanding of the evolutionary dynamics is needed both to design and to interpret such experiments. And such experiments will enable more useful general modeling and analysis.

6.3. Abstract models

To develop understanding — qualitative and quantitative — of broad issues in evolutionary dynamics, abstract modeling is also needed. Scenarios can be studied in toy models that are crude caricatures of a few potentially important features. And general issues of dependence of key quantitative parameters — e.g. population and genome sizes — can be analyzed.

As discussed in the introduction, to many — from Darwin on — the biggest puzzle in evolution is the evolution of complex functions. But to even begin to think about how this occurs so — seemingly to many — fast, one needs some quantification of degree of complexity and of how fast is "fast".

Independent of their motivations for doing so, some advocates of "intelligent design" as an "alternative" to evolution have tried to introduce notions and ques-

tions that should be taken far more seriously than they have been. In particular, some have focused on the apparent "irreducible complexity" of certain biological functions — loosely, how many components and interactions these need to function at all. Whatever the legitimate criticisms of the examples they have chosen (and the lack of understanding of biology these might represent) the general issue cannot be waved away. This is not an issue of marginal improvements of existing functions, but of the evolution of "new" — whatever that means — functions. It is surely true that even many of the simple functions that cells perform could not be evolved by a series of purely beneficial mutations or other genetic changes *in their current biological context.* Perhaps they could occur by a route that also changes many other functions of the cell. But to invoke such an explanation relies on murky assumptions about the biological architecture and dependence of fitness on multiple functions. And this is a major part of what one is trying to understand. In the real world, the evolution of complex functions surely relies crucially on past evolution of other functions. This issue is thus at the heart of evolvability.

The state of understanding of the difficulty of evolving any even moderately complex processes is so poor that almost any progress on abstract models that include some of the essential aspects of evolution would be valuable.

A long term goal is to formulate and address questions about how the difficulty of evolving functions depends on their complexity, and how this depends on the basic biological architecture — crudely, modular or holistic network — on recombination, and on quantitative parameters. For a class of functions that have increasing complexity loosely parametrized by some H, how does the difficulty of evolving these grow with H? To make sense of this question, one first needs well defined classes of biological-like functions, such as pattern recognition, for which knowledge from computer science should be invaluable. For these, one then wants a definition of H — such as the minimum number of components needed, although this may not be a good measure. And one has to consider definite classes of architectures and types of mutational or recombinational processes. An obvious measure of difficulty, D, is the total number of "cell divisions", NT, needed. However, as discussed in the introduction, it is not even clear that this is the most relevant combination of population parameters: indeed, the analysis of the staircase model and ensuing discussion suggests it is not. If the difficulty, D, was proportional to a power of H, highly complex functions could evolve readily. But the naive expectation is that D grows exponentially with H. This would be expected if all H components were needed for increased fitness and the probability of these arising were the product of H small factors (as in the estimate of rates of K-hit processes). If this were the case, the largest evolvable H would be modest even with enormous NT. But high evolvability might mean that how D depends on H is intermediate between these behaviors.

This could only arise from the nature of the architecture and concomitant maps between genome and phenome that are the *least unlikely* to have evolved in an *ensemble* of environmental histories.

The crucial issue here, as already arose in the simple processes that we considered in these lectures, is extrapolation. For example: for many purposes, once the difficulty, D, of a problem grows faster than a power of its size — here H — whether D grows as e^{CH} or as, say $e^{cH^{\frac{1}{3}}}$ does not much matter. But with $NT \sim e^{100}$, it matters a great deal. As emphasized earlier, this means that simulations are of little use in the absence of a theoretical framework: extrapolation over a wide range of parameters is not possible.

Progress in developing abstract models and in framing quantitative questions would be major steps forward conceptually and certainly advance our quantitative understanding.

7. The state of the field

If nothing else, I hope these lectures have made the case that there is a huge amount to be done to even begin to understand evolutionary dynamics. I would thus like to end with a comment on the state of the field. As this author, this school, and much of the audience are physicists, it will take the form of an analogy with the field of physics. It is loose, but I think instructive.

But first a quote from Richard Lenski from a review of the late great evolutionary biologist Ernst Mayr's last book: [21]

"*Mayr argues that the precise mathematics that underlie physics are not applicable to biology, in which determinism, typological thinking, and reductionism have limited utility. ... [H]e builds on this point by splitting biology into two distinct domains, functional and historical. While functional biology may fit within a framework similar to that of physics, Mayr argues that the historical domain of biology–in essence, evolution–requires a different framework. My [Lenski's] own view is that evolutionary history, reflects dynamical processes (e.g., mutation and natural selection) that can be described mathematically and tested experimentally (as indeed they often are), although evolving biological systems are more complicated than what physicists study. [Arguing for a] distinction between the functional and historical domains of biological understanding may reflect [a] limited interest in evolutionary dynamics per se.*"

If the basic laws of evolution are analogous to the laws of quantum mechanics, then the simplest evolutionary process (well described by population genetics) is like the hydrogen atom. The most complicated evolutionary processes directly observed are like simple molecules. And the statistical dynamics of multiple

neutral mutations is analogous to ideal gasses. At the opposite extreme is most of evolutionary theory. This is much like geology: the constituents are known, many patterns are observed — some qualifying as laws — with varying degrees of understanding, and historical scenarios are well developed and can be predictive.

But there are many levels between simple chemistry and geology. Most (although by no means all) of these are understood, each level in terms of lower levels: largely from condensed matter physics and geophysics. And these enable extrapolation over a huge range of length and time scales.

In contrast, for evolutionary dynamics understanding of most of the intermediate levels — or even what these are — is very limited. And the lack of quantitative understanding masks, I believe, severe limitations of the qualitative understanding. Genomic data is, perhaps, starting to provide the more complicated chemistry. And the simple models of acquisition of multiple beneficial mutations are perhaps like ideal periodic solids or one-dimensional Ising models. But surely a major effort combining experiments, sequencing data, observations, and theory is needed.

One can hope that in the near future far more interest will be sparked in evolutionary dynamics, *per se*.

Evolution is — in contrast to popular perceptions in the United States — a fact. The evidence, reinforced by understanding of the basic laws, is overwhelming. But it will take far better understanding on multiple levels for evolution to become a fully fledged theory.

8. Acknowledgments

My understanding of evolution — albeit very limited — owes much to colleagues too numerous to mention. Thanks are due to all these. Almost all the original work summarized in these notes was carried out in collaboration with Michael Desai to whom I am most grateful. He and Olivia White graciously provided many useful comments on these notes. This work was supported in part by the National Science Foundation, Merck Corporation, and the National Institutes of Health.

References

[1] There is an extensive literature on most of the topics discussed here. Below are only a few general and historical references, together with some particularly pertinent recent papers.

[2] D.E. Nilsson and S. Pelger, Proc. Roy. Soc. B 256, 53 (1994).

[3] R. Dawkins, Nature 368, 690 (1994).

[4] E. Mayr, *What Makes Biology Unique? Considerations on the Autonomy of a Scientific Discipline* (Cambridge, New York 2004).

[5] A. Novick and L. Szilard, Proc. Natl. Acad. Sci. USA 36, 708 (1950).

[6] S.E. Luria and M. Delbruck, Genetics 28, 491 (1943).

[7] S.F. Elena and R.E. Lenski, Nature. Revs. Genetics 4, 457 (2003); T.F. Cooperl, D.E. Rozen and R.E. Lenski, Proc. Natl. Acad. Sci. USA 100, 1072 (2003).

[8] W.F. Bodmer, in *Prokaryotic and Eukaryotic Cells*, edited by H.P. Charles and J.G. Knight, (Cambridge University, Cambridge 1970); M.G. Lorenz and W. Wackernagel, Microbio. Review 58, 563 (1994).

[9] W.B. Whitman, D.C. Coleman and W.J. Wiebe, Proc. Natl. Acad. Sci. USA 95, 6578 (1998).

[10] J.H. Gillespie, *Population Genetics: A Concise Guide*. (Johns Hopkins, Baltimore 1998).

[11] W.J. Ewens *Mathematical Population Genetics: I. Theoretical Introduction* (Springer, New York 2004).

[12] M.M. Desai and D.S. Fisher, submitted to Genetics.

[13] I.M. Rouzine, J. Wakeley, and J.M. Coffin, Proc. Natl. Acad. Sci. USA 100, 587 (2003).

[14] P. Gerrish and R.E. Lenski, Genetica 102/103, 127 (1998).

[15] M.M. Desai, D.S. Fisher, and A.W. Murray, submitted to Curr. Bio.

[16] J. Maynard-Smith, J. Theor. Biol. 30, 319 (1971).

[17] Y. Kim, and H.A. Orr, Genetics 171, 1377 (2005).

[18] N.H. Barton and S.P. Otto, Genetics 169, 2353 (2005) and references therein.

[19] I.M. Rouzine and J.M. Coffin, Genetics 170, 7 (2005).

[20] Y. Iwasa, F. Michor, and M.A. Nowak, Genetics 166, 1571 (2004).

[21] R.E. Lenski, Bioscience 55, 3568 (2005).