

## Chapter 3: Population Dynamics

### Simple Selection

Introductions to evolutionary theory commonly explain the concept of natural selection with examples like this: Walking along a rocky beach, you cast your eye down the shoreline and you notice that pebbles of different sizes, instead of being scattered at random about the beach by the crashing waves, are neatly arranged in bands according to size. This orderly arrangement, though it looks almost like the work of an intelligent mind (or at least a meticulous and industrious mind), is in fact the result of certain simple physical processes. Why each size finds itself stable in each range is of course a rather complicated affair, but why it is that similar pebbles are grouped together is not. Details aside, the basic principle is this: different sizes of pebbles have stable positions at different distances up the beach, and as each pebble lands in a range where it is stable, it tends to settle there. If it lands in a range where it is not stable, it is easily dislodged. Over time, this has an ordering effect, and though individual pebbles may be moved out of their particular “stability zone”, the cumulative effect of the differential tendency of pebbles of different sizes to stay in different locations overwhelms the short term chaos of crashing waves.

The universe, we are informed, is populated by locally stable arrangements and things. Things and arrangements that are locally unstable tend not to be much in evidence. The mechanism at work, natural selection, is so nearly trivial and so completely ubiquitous that it tempts one to accuse evolutionary theory of emptiness, or tautology. And of course, such accusations as well as the various responses are more than familiar by now. Simple selection processes, Dawkins (1986) calls them “sieves”, are of course only the beginning of the story for biological evolution. So, the story usually continues with the introduction of the notions of heredity, replication, and the “cumulative” change that results from the distinctive evolutionary dialectic of variation and differential reproductive success. Before long sexual reproduction and the attendant shuffling of genes which is responsible for most current evolutionary novelty is introduced. And somewhere along the line it turns out that replication is essential to natural selection (or to evolution, or both). If

---

the author anticipates the reader's question, "what about the pebbles on the beach?", the answer is that the pebbles exhibit a sort of simple selection, which is by itself grossly inadequate to account for the wonderfully adapted (and co-adapted) structures we see in modern organisms. Such processes and principles are important for purposes of illustration, but since we don't get very far in understanding biological evolution without the inclusion of heredity and replication, why waste terminology on them? "Natural selection" (and/or "evolution") is surreptitiously redefined in such a way as to include the necessary ingredients for apparently "progressive", cumulative change.

The consequence of this common explanatory development is, as I tried to show in Chapter 1, that when people turn to the question of cultural evolution — of the evolution of science, language, and knowledge in general — they tend to ask "what are the replicators?" If they have read further and learned of the "central dogma" of evolutionary theory — that only changes in the germ plasm and not in the somatic tissue are inherited — they tend to ask "what in cultural evolution corresponds to the genotype and what corresponds to the phenotype?" As I argued earlier, this tends not to get us very far. What we should do, is go back to the pebbles on the beach.

The primary concern of this book is to try to understand knowledge from a biological point of view. But more than that, we are concerned to try to find a way to study knowledge as a biological phenomenon, to "naturalize" epistemology, to turn it into a field of scientific enquiry. What does this have to do with pebbles on the beach? It turns out that the mechanism of information transfer that is essential to understanding the biology of knowledge is already evident in the pebbles case, just as the basic mechanism of evolution is. Of course, before we are done there will be plenty of twists and turns in the story of knowledge, just as there are in the story of the evolution of life.

You cast your eye down the beach, and are struck by the orderliness, the beauty even, of the arrangement. It might have been designed by an intelligent mind (or at least a meticulous and industrious mind). But there is something else, something not just to admire, but to learn. The arrangement or spatial distribution of pebbles tells you something about how things are. It tells you that there is something at work that imposes order on

---

the spatial distribution of pebbles. It also tells you locally, on *this* beach, where the stability zones are for different kinds of pebbles. What the repeated selection of locations has done is left the imprint of myriad local physical processes on the local distribution of pebbles. There is, in short, *information* about the local environment coded into the distribution of pebbles. If we knew more we could, for instance, say something about the average height of the surf, or the recent occurrence of storms.

This is not to imply that it is only selection processes which leave information in their wake. Physical processes in general leave traces, which may (if we know enough) allow us to infer back to their causes. Selection, as the net effect of a number of causal processes does this as a matter of course. In the case of the pebbles on the beach, the kind of information it leaves in its wake is perhaps as trivial as the simple single-step selection processes we are considering. Change the example: I watch the frequency of a dark variant of peppered moth rise while the frequency of the light variant drops. Given that this is the result of selection, what does this tell us? That this is the kind of environment where dark moths do better. Again, this seems pretty trivial. I acquire a new dog, and discover after a year that the potted plants in my garden are all now in plastic pots, whereas originally there was an equal mix of clay and plastic. What does this tell me? That this is now the kind of garden where plastic pots do better. Plastic pots are more stable in the new environment. Again, trivial.

Aside from the triviality of the information transferred, you will now (or should, anyway) be worried about another problem. Aren't I importing the knowledge mechanism into these scenarios? If so, doesn't this beg the whole question of the naturalization of our understanding of knowledge? The answer is yes, though it is difficult to avoid presupposing knowledge of the world when building a theory. But try the following: one finds that in various regions, different styles of clothing predominate. In very northern or southern regions, clothing tends to be heavier, and in tropical and equatorial regions, clothing tends to be lighter. Presumably, Alaskans *could* wear shorts and t-shirts, and Southern Californians *could* wear goose down parkas. But they don't. From this, I conclude that Southern California is the kind of place where shorts and t-shirts do better,

---

and Alaska is the kind of place where goose down parkas do better. This seems like the kind of thing that might be good to know, a kind of non-trivial (or at least potentially useful) information. As for the second problem, forget about me for the moment. The information that Alaska is the kind of place where goose down parkas do well is encoded in the dispositions to acquire and wear clothes of the locals, and would be there (and would be useful) even if no “knower” ever observed the fact. That, I submit, is no trivial thing. And there is no smuggling in of the “knower”, either.

We still need a bit more apparatus if we want to account for something approaching knowledge, since it isn't clear whether the reason that this is the kind of place that favors *this* kind of clothing is due to the clothing being better for people (compensating for harsh environments), or because local trend-setters and determinates of social status make it advantageous. What we need for knowledge (or at least what I am planning to supply in what follows) is that the locals are a basically sensible people capable of “learning” via trial-and-error what kind of clothing is the best and that they *select* styles of clothing not due to fad but due to superior protection from the elements. Trial-and-error learning is easily understood as a cultural- level selection process, and when what is doing the selection is something like the acute discomfort that comes from dressing in a way that threatens to dangerously elevate or lower body temperature, then the distribution of clothing styles encodes valuable information about the local environments, and that it is held in dispositions to dress makes it both exploitable and independent of third-party knowledge mechanisms. Reliable, learnable, exploitable information about the world is a decent candidate for knowledge, at least for the naturalist.

On the other hand, clothing styles are notoriously subject to the dictates of fashion. What this means from our point of view is that clothing styles inhabit a complex ecology. The tendency of humans to decorate themselves and to judge others according to how *they* decorate themselves frequently competes with more utilitarian factors in determining the distribution of dispositions to dress. The discomfort that arises from thermal extremes is only one of a number of selection mechanisms in the cultural ecology of clothing. But the distribution of clothing reflects how the world is, to whatever extent, because the

---

mechanisms that *create* discomfort are reliable, and the fact that that discomfort is correlated with risk of illness or death makes the information valuable. The story behind the reliability of the discomfort mechanisms is of course not a story of cultural evolution, but of genetic evolution. *Our* ancestors did better because they were disposed to dress so as to keep their body temperatures within the “safe” range, and that is why we are sensitive to temperatures in the way we are. We need to understand both the cultural evolution of clothing and the genetic evolution of the discomfort mechanisms in order to understand to what extent *knowledge* about the world is encoded in the clothing distribution. Moreover, the two evolutionary processes interact. The evolution of genetically based learning abilities goes on *while* those abilities are influencing the course of cultural evolution. How well culture does in arriving at and maintaining solutions to problems affects in turn how well people with the aforementioned learning abilities do. Genetically inherited learning abilities affect the cultural “fitness” of cultural items like clothing styles and the local efficacy of cultural items like clothing styles affect the well-being and reproductive success of the learners.

This “fitness feedback loop”, along with the information transferring tendencies of selection are the central components of the model of biological knowledge systems that will be developed in the next few chapters. The picture of multi-level selection processes itself is not a new one. Psychologist Donald T. Campbell, following on the work of Karl Popper, laid out the basic vision in his (1974) “Evolutionary Epistemology.” Campbell’s idea was that selection and variation processes are ubiquitous, occurring on many “levels” in biological and especially human knowledge systems, and it is via the interaction between levels that our beliefs and acquired dispositions get reliably and usefully correlated with the world. The challenge, if one wants to make a science of biological knowledge systems, is to try and make this clear.

### **Modeling Evolution**

There are a number of kinds of models of evolutionary processes. One of the best known is the cellular automaton, of which the most familiar example is the “life” simulator. The cells in a cellular automaton are not, as one might think, the dividing cells of biological growth. Rather, the cells are squares on a grid, much like a chessboard. The system evolves as follows: each cell can be in one of a number of states at a particular time. In the simplest examples, the states are just “on” and “off.” At each time and for each cell, states of the (eight) surrounding cells are recorded, and then the state of the central cell is either changed or left the same according to some rule that dictates the “dynamics” of the automaton. In the “life” simulator, for instance, the rule is “If two adjacent cells are on, then the cell stays in its present state. If three, then the cell is turned on, otherwise it is turned off.” The process is then repeated, and this iterative cycle goes on indefinitely. As you probably know, this particular algorithm results in some very interesting patterns, in particular the selective stabilization of a variety of distinct “life-forms”, and an ecology complete with “predators.”

Another popular approach to modeling evolutionary processes is the genetic algorithm developed by John Holland (1975, see Mitchell 1996 for a good introduction). Genetic algorithms, as the name implies, model the evolution of actual strings of, well, numbers, though the numbers can be used to do all kinds of things. They could, for instance, function to determine the phenotypes of a population of simulated animals. Or, they could be used to set the controls on a complicated machine. Or again, they could be used to determine the values of variables in a set of differential equations. Typically, one starts out with a modest number of such strings (say a hundred). Each string is then “plugged in” to the application of interest — inserted as variables in the system of equations, used to set the controls on the machine, or used to generate a population of simulated animals. The performance of each string is then evaluated according to some test (the “selection mechanism”). Strings that do badly are replaced with variants of the more successful strings. Variations can either be “mutations”, where small random

---

changes are made to successful strings, or they can be “recombinations”, where segments of two successful strings are spliced together to generate a new string. Each member of the new set of strings is then “plugged into” the application, evaluated, and selected. Again, the process is repeated indefinitely.

What cellular automata have in common with genetic algorithms is that both attempt to actually *instantiate* evolutionary processes in the computer. They model the behavior of actual individuals (strings of numbers or cells in a matrix) rather than the average behavior of a large population of such individuals. Cellular automata are particularly good for modeling spatial effects, due to the inherent relationships of proximity that exist between cells on a grid. Genetic algorithms are particularly good for exploring large “solution spaces”, and have come onto their own as tools for finding solutions to systems of differential equations that are intractable under standard techniques. For our purposes however, a older, simpler, and more common kind of evolutionary model will be more appropriate.

### Population Models

Mainstream evolutionary biology has not embraced the replicator as a central concept, but continues to follow in the population genetics tradition deriving largely from the work of R.A. Fisher, Sewall Wright, and William D. Hamilton. Within this tradition, the central object of analysis is the *population*. It is the population, rather than the lineage which evolves, and that evolution consists of shifts in the relative frequency of types within the population. This *frequentist* or population-oriented approach has proved quite fruitful in extensions of evolutionary theory beyond the dynamics of pure genetic transmission, notably in Cavalli-Sforza and Feldman (1981), Boyd and Richerson (1985), as well as the large literature on Evolutionary Game theory following (roughly) from Maynard Smith and Price (1973). It is this mainstream tradition that we will be following in the development of the formal system in this chapter, though some care will have to be taken to avoid the pitfalls of analogical thinking. In particular, the paradigm formalisms of

---

population genetics<sup>1</sup> are geared toward the dynamics of sexual reproduction, which we should expect will not be appropriate for the more general class of evolutionary processes, which includes cultural evolution, and of particular interest, science as an evolutionary process.

Population models are a species of vector models. Vectors, to begin with, are simply lists of numbers. Newtonian physics uses vectors to specify the position and momentum of bodies. The position vector has three components, which give the  $x$ ,  $y$ , and  $z$  components of the body's location. The general form of the vector is written like this:  $\langle x, y, z \rangle$ . The 3-place position vector determines a point in three-dimensional space when the three variables  $x$ ,  $y$ , and  $z$  are given specific values, like this:  $\langle 3.1, 6, 7.8 \rangle$ . The momentum vector also has 3 components, which give the current momentum of body in each of the three spatial directions. Consequently, the full state space of the system has 6 dimensions, three for the spatial location of the body and three for the momentum of the body. The full state of the system, as determined by the values of the two vectors, picks out a single point in that six-dimensional space. The motion of the body is simulated by modifying the position and momentum vectors according to the laws of motion. Vector models thus represent the state of the system by giving specific values to one or more vectors, and they model the *evolution* of systems with sets of equations which modify those values over time, moving the point through the state space of the system.

Population models use vectors in a similar way. A population is a collection of individuals, categorized according to type. The distribution vector  $\vec{p} = \langle p_1, \dots, p_n \rangle$  gives the state of the population in terms of the frequencies of each of the  $n$  types of things in the population. Just as with the position vector in physics, this determines a point in the system's state space, which has  $n$  dimensions — one for each type of thing. Natural selection is simulated by changing the values in the vector according to the relative fitness of each kind of thing in the population. The evolution of the population appears as a trajectory across the population's state space.

---

<sup>1</sup> Specifically, the Hardy-Weinberg and Fisher equations. See Hofbauer and Sigmund 1988.

Normally, the values of the vector  $\vec{p}$  are restricted, due to the fact that  $\vec{p}$  tracks




Figure 3.1

the relative frequencies of types in the population. Frequencies, by definition, range between zero and one. Moreover, they all must add to one. This means that  $\vec{p}$  is restricted to a subspace of its full  $n$ -dimensional space, a region with a distinctive shape. For a population with only three types ( $n = 3$ ), the subspace looks like figure 3.1(A): an equilateral triangle with vertices at  $\langle 0, 0, 1 \rangle$ ,  $\langle 0, 1, 0 \rangle$ , and  $\langle 1, 0, 0 \rangle$ . This triangular region is referred to as the “simplex”. The simplex for  $n = 3$  is a *two* dimensional equilateral triangle. The simplex for  $n = 4$  is a tetrahedron, and so on. The simplex for  $n = 3$  turns out to be a convenient way to represent the dynamic of population with three types. (See figure 3.1(B).) This is sometimes referred to as a “phase portrait.” At the point in the center, all types are equally represented in the population. The vertices are points of “fixation” for each of the three types. On the edges the frequency of one of the types is zero.

As you might imagine, there are many sorts of population models. The ones we will be developing differ in important ways from the standard models of population genetics as well as from the standard models of evolutionary game theory (the most “cultural” of the well-explored evolutionary models to date). However, since we will be working toward models in which genetic and cultural evolution, along with their

interactions, can be simulated in a single model, we need to build our modeling system so that it is general enough to cover both cultural and genetic evolution, learning and inheritance. The trick is to start simple, and provide ways of adding features to the dynamical equations as they become relevant.

What all evolutionary population models have in common is *selection*. Selection consists in shifts in the relative frequency of types in a population due to differential fitness of the types. The basic mathematical representation of fitness and selection is quite simple. For each type  $i$ , there is, along with its frequency  $p_i$ , a fitness which is designated  $w_i$ . It is convenient to collect the fitnesses of the  $n$  types into a second vector  $\vec{w} = \langle w_1, \dots, w_n \rangle$ . Think of the  $w_i$ 's as "growth rates". Things that don't reproduce will have growth rates between zero and one. Things that do can try for higher growth rates. The way that the fitnesses govern the dynamics of the population is quite simple as well. As a first approximation, the new frequency of some type  $i$  is the old frequency times the growth rate. This is written,

$$p_i' = p_i w_i \quad (1)$$

At each generation this multiplication is performed by substituting each of the 1 through  $n$  type indexes for " $i$ " in equation (1), and this gives the new frequency distribution of types in the population. The problem with equation (1), however, is that there is no guarantee that the new sum of the  $p_i$ 's (written  $\sum_i p_i$ ) will add up to one — which is to say, there is no guarantee that the new frequencies will actually be proper frequencies at all. The obvious solution is to divide each new frequency by the new *total*, a process referred to as "normalization." You could write this a number of ways, but the convention is this:

$$p_i' = p_i \frac{w_i}{\sum_j p_j w_j} \quad (2)$$

The quantity in the denominator is just the new population size. But since the old  $p_i$ 's are

---

assumed to be true frequencies, the quantity in the denominator is also the average or *mean fitness* of the population. Mean fitness is written  $\bar{w}$ , so that it is common to express equation (2) like this:

$$p_i' = p_i \frac{w_i}{\bar{w}} \quad (3)$$

Equation (3) just says that the relative frequency of each type will increase if its fitness is higher than average, and will decrease if it is lower than average. The greater the ratio of a type's fitness to the mean, the more that type's frequency will increase. Likewise if the type's fitness is smaller than the mean.

Suppose we have a population with 10 types. Imagine the population to be whatever you like, varieties of hummingbird, clothing styles, or beliefs in various stories of creation. If we choose the initial frequencies of the 10 types at random, and do likewise for the fitnesses of each type, the population might evolve as in figure 3.2.

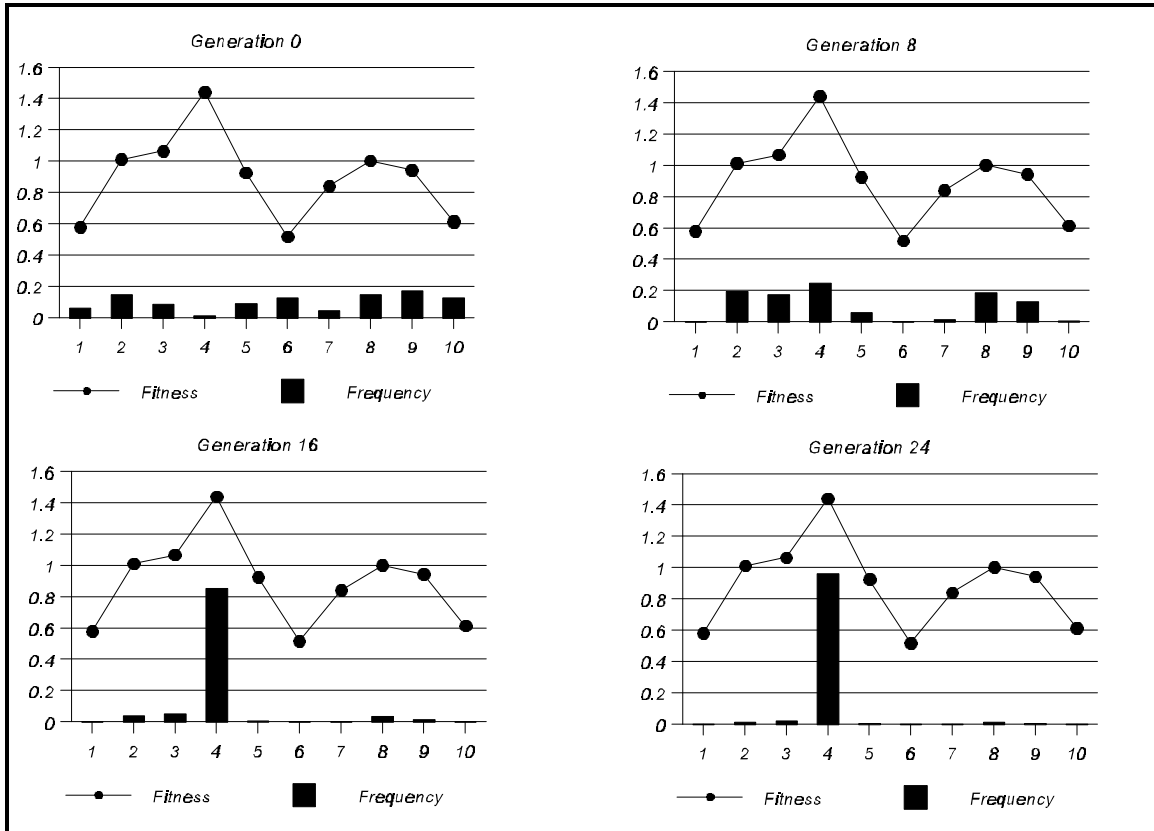


Figure 3.2: Selection under fixed fitnesses.

Given that our distribution vector  $\vec{p}$  inhabits a space of 10 dimensions, and our fitness vector  $\vec{w}$  a space of another 10 dimensions, we need to be a bit clever in order to draw a picture of the state of the population. Figure 3.2 uses a pair of graphs to pick out a point in the 20 dimension state space. The bar graphs give the values of  $\vec{p}$ , and the line graph above gives the values of  $\vec{w}$ . This kind of representation is known as an “adaptive landscape”, and the basic idea is due to Sewall Wright. Over the course of 24 iterations of our simple selection equation (3), we can see that type 4, which has the highest fitness, comes to predominate. Indeed, it predominates despite the fact that it is initially not very well represented. In our example, the fitnesses range between .5 and 1.5. This means that the *absolute number* of individuals of a type can increase by 50%, or decrease by 50% in a single generation. But the relative frequencies may not change at those rates, since it is the relationship to mean fitness that determines the growth in relative frequency, not the

---

absolute growth rate. Indeed, the usual practice in population genetics is to adjust (or “normalize”) the *fitnesses* so that the largest fitness is always equal to one. Here however, partly because absolute growth rates seem more intuitive to me, and also because we will want to use them in the multi-level models we will be building in later chapters, we will continue to use absolute growth rates as fitnesses. The important thing to understand here is that when the population frequencies are being normalized in each generation (as they are in equations (2) and (3)) it is the relative size of the fitnesses that determine the dynamics, not the absolute value. For all the difference it makes, you could use fitnesses that ranged between 1 and 100, or between one million and one billion.

Figure 3.2 provides a useful heuristic for thinking about selection as an information transfer process. It may help to think of the population as a very simple learning device. What it learns is always basically the same, it learns which type of thing in the population has the highest growth rate in the local environment. Alternatively, what it learns is whether *this* is the kind of environment where type 1 has the highest fitness, or whether type 2 has the highest fitness, and so on. When the fitnesses are fixed as they are in our example, and no other sources of frequency change intrude, populations are very reliable learners. Given enough time and a consistent environment, it is inevitable that the population will pick out the most fit type (#4). The population does this picking out by being dominated by the most fit type, which, if you think about it, seems to make more sense than having it go to extinction. Of course, this is all metaphor, and fairly outrageous metaphor at that. The population isn't making use of the information or interpreting it. But the thing to recognize is that the information is of a useful sort, and it is held in a useful way, *and* the gaining of this sort of information is inevitable in favorable situations. Things will get messier and more interesting soon enough.

Recall that we started by complaining that popularizers of evolutionary theory tend to focus on the specifics of biological/genetic evolution, thus encumbering attempts at a theory of cultural evolution with unreasonable prerequisites. Building a general theory up from the mathematics of evolution, on the other hand, leaves us much freer. Selection, from this point of view, does *not* require replication, so there is no need to look for

---

replicators, nor to despair over the prospects of a theory of cultural evolution when we don't find them. What evolution by natural selection requires is that there be a population, whose type frequencies change over time due to fitness differences. Given that one can treat any collection of objects as a population, this reduces to the question of whether differential fitnesses are driving frequency shifts. But since a fitness is, by definition, a type's growth rate, it stands to reason that any group of objects that have fitnesses, where those fitnesses vary, will evolve by natural selection. So what kinds of things have fitnesses? There are a number of distinct notions of fitness, and these have changed considerably since Darwin. Many of these include the very notion of reproductive success (Dawkins 1982). But the notion of fitness at work in the mathematics of population modeling is quite clear: fitness is a numerical expression of a growth rate of types. Our choice of convention — that fitnesses are *absolute* type growth rates — makes things even clearer. And, it turns out that literally every type of thing in every environment has a fitness. Types whose members reproduce obviously have growth rates. But even types whose members do not reproduce have growth rates, it's just that the rates are always one or less. Types of non-reproducing things, strictly speaking, have measures of stability, or "decay rates." But mortality rates are part of reproductive rates, so it would seem rather arbitrary to insist that types which go out of existence at certain rates but do not reproduce have no growth rate, no  $w$ . There is certainly no mathematical reason for such an insistence, and since what interests us is how far we can legitimately go in applying the natural selection *equations*, there is no reason not to include decay rates as a special case of fitness. Consequently, we seem entirely justified in saying that every group of objects, unless every type of object in the group has the *same* fitness in the current environment, will evolve according to natural selection.

This is not to say that all frequency shift in every population is due to natural selection — that fitnesses subsume every cause of such shifts. Evolution, especially in the most interesting cases, proceeds by a kind of dialectic between selection and *variation*. The most commonly discussed sources of variation in biological evolution are genetic mutation and recombination, but importantly, *immigration* is also a non-selective non-

---

fitness related source of frequency shift in populations. Mutation and recombination are *endogenous* sources of frequency shift, immigration is an *exogenous* source of frequency shift. But what is it that groups mortality (or stability) and reproduction as contributors to selection, and mutation and immigration as non-selective sources — contributors to variation? The answer to this question will allow us to distinguish *cultural* forces as contributors to selection or variation according to principle, rather than by appeal to analogy.

The answer to the question of what constitutes selection and what constitutes variation has two versions, one based on kinds of causal processes, and one based on the mathematics. On the causal story what makes something a contributor to fitness and thus factor in selection (assuming fitnesses differ) is that the effect on the type's frequency depends on features of members of the type and the interactions of those features with the local environment. Suppose we have a population with types 1 through  $n$ . Type 1's frequency ( $p_1$ ) decreases due to mortality, some of its members have died. This is an expression of fitness because it was members of type 1 whose features resulted in type 1's frequency shift. Type 2 increases due to reproduction. This is an expression of fitness because it is existing members of type 2 that result in type 2's increase. Type 3's frequency increases due to a massive influx of type 3 from outside the population. This is not a matter of fitness, because the increase in type 3 has nothing to do with features of established members of type 3 interacting with the local environment. Type 3 did not *grow*, but was *added to*. A number of members of type 4 mutate into members of type 5. Consequently, type 5's frequency increases. But again, this is not a matter of type 5's fitness, since members of type 5 didn't do anything to bring about the frequency change. Discussion of what to say about type 4's loss (i.e., is it selection or variation) will be deferred.

On the causal story, what makes a frequency change a matter of fitness/selection or a matter of variation depends on whether pre-existing members of the type are involved in the change, or whether they are “passive observers” of the frequency shift. This way of specifying the difference between selection and variation squares well with virtually all

---

biological uses. It also turns out that, given that variation is simply characterized as sources of frequency shift that are not fitness components, then it follows that selection and variation cover all sources of frequency shift. This is an interesting consequence, in that it indicates that all sources of frequency shift can be accommodated within the conceptual framework provided by evolutionary modeling.

In the example above, I said that after an influx from outside type 3 did not grow, but was added to. The mathematical difference between selection and variation turns on the same distinction. In the causal story, in each case, the affect of fitness components on a type's frequency is proportional to the number of pre-existing individuals of the type. A high reproductive rate won't make much difference in a type's frequency, at least initially, if there are only a few members of the type around. The way this gets expressed mathematically is that the measure of fitness  $w$  is *multiplied* by the old frequency. All selection processes in population models have this feature. Sources of variation, on the other hand, are characterized by something being *added* to a type's frequency. The effect may, of course, be proportional to the frequencies of other types, but not to the type whose frequency is changing. So the easy answer, and as yet I have no reason to be dissatisfied with it, is that growth rates (fitness and selection) multiply frequencies, sources of variation add to them.

The general model of evolutionary processes that is beginning to emerge (and I expect that it may be viewed with alarm, or at least distaste, by some purists) is this: selection and variation processes are indeed ubiquitous, making the tools of evolutionary analysis applicable in a much wider variety of situations than has been imagined. Indeed, any arbitrarily chosen group of objects can be treated as a population. This does not mean that the population will evolve, for it can fail to evolve by not having the relative frequencies of types in the population shift. One would expect this sort of thing to be unusual, however. So we can say that any group of objects is a population, and most populations evolve. What about selection and variation? Every type of things has a fitness, a summation of "multiplicative" growth factors, in a given environment. Not every

population that evolves, evolves by natural selection, however. Only when members of the population have different fitnesses does natural selection occur. Variation can cause the evolution of a population in the absence of selective forces (fitness differences), though one expects that the usual case is a bit of both.

What is potentially disturbing about this treatment of evolutionary theory is that it may seem to trivialize the concepts, by making them apply to everything. I have some sympathy with this worry, at least insofar as it is motivated by the concern with trying to teach people how biological evolution *actually* works, a task which is frequently impeded by too free a use of metaphor. But on the other hand, our aim here is to put together a set of conceptual and mathematical tools which will make tractable the analysis of the interaction between genetic evolution and cultural “evolution”, and it just happens that the existing mathematical models are flexible enough to do this, with minor modifications.

Apologies aside, then, the reason that Campbell was right about the ubiquity of selection and variation processes is that the entrance requirements are fairly lax. In fact, rather than thinking about selection and variation processes as *kinds* of processes, it might be better to think about selection and variation models as *ways* of looking at or analyzing the evolution of populations. The appropriate analogy here is to statistical properties. The question in applying basic statistical concepts is not whether the population has a “mean” or a “median”, but whether looking at the population in those terms is helpful. Likewise, the question of whether to use an evolutionary model to analyze a population is not whether the concepts *apply*. The concepts always apply. If the frequencies are shifting, then it is due to selection, or to variation, or more usually to some combination. It may be that no-one has ever written down the kind of equations that determine changes in frequency in your population, but those equations, when they are written, will either constitute selection, or variation, or some combination. This rather pragmatic attitude toward the application of evolutionary models is made all the more plausible if you consider that, in general, there is no clear answer to what the populations are, where one population ends and the next begins. Moreover, even if the population is given, there is still a rather pragmatic decision to be made as to how many types to divide the population

into, and at some level there may not be a single right answer. This is not to say that as an area of inquiry develops, we don't learn how it is most fruitful pick populations and types. This certainly seems to have happened in biology, with "island biogeography" providing a clear case of rather obvious population divisions.

The remainder of this chapter will be devoted to explaining some of the standard ways in which the simple equations above are made more interesting. I will also add a non-standard formula that will prove especially useful in modeling cultural evolution.

### Mutation

All evolution of populations consists either of selection, or variation, or both. We have already seen the most basic kind of selection model, and making selection more interesting consists mostly in making the fitnesses change in interesting ways. However, we haven't seen any models of variation, or non-selective change, yet. Recall that selection consists in multiplying a type's frequency by its fitness. Variation consists in adding something to the frequencies. The simplest model simulates *uniform random mutation*.

Mutation is a process by which an individual changes its type. This results in two things: the frequency of the type that the individual was a member of is decreased, and the frequency of its new type is increased. In uniform random mutation, every type mutates at the same rate, with an equal chance of mutating into every other type. Let  $\mathbf{m}$  be the rate at which mutation occurs. A value of  $\mathbf{m} = .001$  means that, at each generation or cycle of the dynamics, one tenth of one percent of the individuals spontaneously change type. The frequency change, to begin with, is:<sup>2</sup>

$$p_i' = p_i(1 - \mathbf{m}) \quad (4)$$

Of course, this only covers the *loss* to mutation of each type. How much does each type

---

<sup>2</sup> Notice that the loss to mutation is multiplied by/proportional to the frequency. This will be discussed presently.

gain? If  $\mathbf{m}$  of each type is mutating, then it follows that  $\mathbf{m}$  of the population is mutating. Since this is uniform random mutation, each type has an equal chance of mutating into any other (or back to itself). So the gain to each type from mutation is  $1/\mathbf{n}$  times  $\mathbf{m}$ , or  $\mathbf{m}/\mathbf{n}$ , where  $\mathbf{n}$  is the number of types in the population. Mutation *adds*  $\mathbf{m}/\mathbf{n}$  to each type, making the mutation equation:


$$p_i' = p_i(1 - \mathbf{m}) + \frac{\mathbf{m}}{\mathbf{n}} \quad (5)$$

Finally, we can combine our selection equation (3) with the uniform random mutation equation (5) to yield:

$$p_i' = p_i \left( \frac{w_i}{\bar{w}} - \mathbf{m} \right) + \frac{\mathbf{m}}{\mathbf{n}} \quad (6)$$

though we will find it more convenient in the rest of this section to consider mutation processes in isolation.

The effect of uniform mutation in isolation is, if anything, more predictable than the effect of selection under fixed fitnesses as in equation (3) and Figure 3.2 above. Figure 3.3 shows the difference for a population of three types.



**Figure 3.3:** Uniform Mutation vs. Simple Selection

Just as selection under fixed fitnesses always results in the most fit type taking over the population, uniform mutation alone always results in convergence on the state where every type is equally represented. This is because every type gains the same amount from uniform mutation, but those which constitute more of the population *lose* more. This tends to equal out the frequencies, and when selection and mutation are both occurring, this “flattening out” of the distribution interferes with the tendency of the population under selection to track the environment by becoming dominated by the most fit type.

This makes it sound like variation is an impediment to selection, and it is true that when mutation rates are very high, they can lower the population’s overall growth (mean fitness). On the other hand, sources of variation can do something that selection cannot — namely, increase the frequency of a type when it is zero. Fitness effects, remember, are always proportional to the frequency of a type, so that when that frequency is zero it remains zero, no matter how well members of the type would do if there were any around. Sources of variation, like mutation, being additive in nature (rather than multiplicative) *can* increase frequencies from zero. Look back at equation (5). If  $p_i$  equals zero, then  $p_i$  will equal  $m/n$ , which at least gives it a chance to compete. There are a couple of other ways to put this. One is that selection never introduces the things that it acts on. Other

---

processes, which necessarily fall under the category of variation must do that. Another is that, under selection alone, extinction is permanent. If the frequency of a type ever falls to zero, there is no way that it can increase under selection alone.

Now, why is it that mutation, which is the paradigm of variation processes, involves multiplication as well as addition, translates into selection as well as variation? Mutation, as the term is being used here, is a process by which an individual changes type. As noted above, this necessarily involved the reduction of the number of individuals of the old type, and the increase of the number of individuals of the new type. The increase to the new type is a clear case of a non-selective process, since the increase did not involve any established members of the increased type. It is the decrease that can cause confusion. It seems to me that clarity requires that we say that mutation, or type change, is a composite process which contributes (negatively) to the fitness of the old type and constitutes variation for the new type. Imagine that our population consists of chemical molecules, and that type 1 is less stable in the current environment than type 2. Whether the destabilization of type 1 molecules constitutes “mortality” or mutation is simply a matter of what happens to the remains. If the remains of the type 1 molecule’s destabilization are of type 2, then the process is mutation. If the remains are of a type that are not considered as part of the population, then the process is analogous to mortality. But the cause of the type 1 molecule’s destabilization is the same in either case — the interaction of the destabilized molecule with the local environment. The effect of the destabilization, the reduction of type 1’s frequency, is also the same. So there really seems to be no choice but to characterize mutation as a composite process, at least if mutation is defined, as we have, as a process by which an individual changes type (and this is certainly what the standard mutation models simulate).

This may seem to do some violence to the biological usage of “mutation”, but the tension seems to me to be instructive, rather than debilitating. There are two reasons why the selective aspect of genetic mutation is not ordinarily addressed in biological evolution. The first is that rates of genetic mutation are typically so low that they play relatively little role in the frequency shifts in genetic populations, serving mainly to introduce novelty in

---

minute amounts. Moreover, genetic replication takes place in an environment that is well insulated from the environment to which the species as a whole adapts, so that whatever selective destabilization occurs in genes, it has little to do with the process of phenotypic adaptation that is the primary focus of the study of biological evolution. In short, the frequency reduction in standard cases is negligible.

Second, genetic mutations occur during the process of copying the gene, so that it is not clear exactly what to say about it. Is genetic mutation a process by which a gene changes type, or a process by which some new type is introduced *by* an individual of some other type, in an act of reproduction gone wrong? Since the reduction is negligible in either case, there is no real reason to make a decision on this somewhat baffling ontological question. Nonetheless, one can see that in either case the loss must be, strictly speaking, a contributor to the fitness of the old type. For if it is a matter of the gene changing type, then the local instability of the gene that initiates the mutation is a fitness contributor. And in the case of reproduction gone wrong, this constitutes a reduction in the reproductive rate of the old gene type — again, a fitness contributor. I am not suggesting that this is something biologists should worry about. All model building requires simplification, and treating genetic mutation as a case of pure variation seems appropriate, given that the negative effect on fitness is negligible. We should not expect, however, that in *all* cases mutation rates are so low — indeed, we will see cases in cultural models where virtually all selective destabilization results not in mortality, but in type change.

In the uniform random mutation model (without selection) discussed above, the loss to mutation of each type (equation (4)), while it was of necessity a fitness contributor, had no differential effect on the types' frequencies, since the mutation rate for every type was the same. In general, one might expect that mutation rates will differ between types, and moreover, that types do not mutate into one another with equal likelihood. Types which are more similar are more likely to mutate into one another than types which are quite different.

The usual way to broaden mutation models to accommodate these sorts of differences is to assume that each type mutates into each other at some fixed rate. So to begin with we need a table or *matrix* of mutation rates. Table 1

<i>Mutation</i>	Type 1	Type 2	Type 3
Type 1	.9992	.0002	.0006
Type 2	.001	.9987	.0003
Type 3	.01	.05	.94

**Table 1:** Non-uniform mutation rates

specifies a mutation matrix for a population with three types. The row labeled “Type 1” gives the rate at which type 1 mutates into each type. Notice that type 1 “mutates into itself” at a rate of .9992. What this represents is not the tendency of type 1 individuals to destabilize and then return to being type 1, but just its tendency not to mutate. This was written as  $(1 - \mathbf{m})$  in equations (4) and (5). So type 1's mutation rate is .0008, type 2's is .0013, and type 3's is a whopping 6%. There is, of course, an equation that goes with the matrix. The “self-mutation” rates in the matrix help make it simple. If we call the matrix  $\mathbf{M}$ , then  $M_{i,j}$  will be the contents of  $j$ th cell in the  $i$ th row. We can then write:

$$p_i' = \sum_j p_j M_{j,i} \quad (7)$$

There are other ways to implement mutation matrixes in vector models, including those that combine matrix mutation with selection in a single equation, but this is the simplest. Equation (7) just says that the new frequency of each type ( $i$ ) is the sum of the contributions of each of the various types ( $j$ ) to  $i$ 's frequency, where those contributions consist of the old frequency of each type ( $j$ ) times that rate at which it mutates into  $i$ .

### Frequency Dependent Fitness

Matrices also prove useful in modeling changes in fitness. Though in the simple epistemological models we will be developing in this book we will have little occasion to use fitness matrices, they will become essential later on as elements of competition and “cooperation” between cultural items move to center stage. The basic idea is that

oftentimes, how well a type does depends on how common various types are in the population. Fitnesses depend importantly on the environment (indeed, from the point of view of population models, fitnesses very nearly *constitute* the environment), and for many populations, individuals in the population form an important part of the environment for other individuals. Frequency dependent fitness models focus on the effects of these important internal interactions, assuming for the most part that external determinates of fitness, while they may differ between types, do not change over time.

Frequency dependent fitness models have been used quite productively to investigate the evolution of cooperation. Suppose that a population of individuals inhabits an environment where they frequently have the opportunity to confer a benefit on some other member of the population, at a relatively small cost to themselves. Say the benefit is worth 3 units of fitness to the recipient, but only costs the donor 1 unit of fitness. It stands to reason that everyone will accept the benefit when offered, but only some will chose to confer the benefit. Following the standard convention we will call those who both accept and confer the benefit “cooperators”, and those that accept but don’t confer “defectors”. If we give every member of the

population a “background fitness of 2 (in this case, these are not absolute growth rates, but just convenient numbers), then this determines the payoff matrix in Table 2.

$W$	Cooperate	Defect
Cooperate	4	1
Defect	5	2

**Table 2:** Frequency Dependent Fitnesses

If we let  $W$  be the name of the fitness matrix (also referred to as a “payoff matrix”) then  $W_{i,j}$  is the contents of the  $j$ ’th cell of the  $i$ ’th row of the matrix. The fitness of cooperators will depend on how many cooperators there are in the population — the more cooperators there are, the better cooperators do. We can calculate the fitness of cooperators by multiplying each payoff with the likelihood of running into each type of individual. If we assume that individuals in the population interact at random, then the likelihood of running into an individual of a given type is just the current frequency of the type. In this case, the expected payoff to cooperators is just  $4 \cdot p_{\text{cooperate}} + 1 \cdot p_{\text{defect}}$ . The

mathematical expression for the fitness or expected payoff of an arbitrary type  $i$  is  $w_i = \sum_j p_j W_{i,j}$ . Inserting this calculation of fitness into our selection equation gives us:

$$p_i' = p_i \frac{\sum_j p_j W_{i,j}}{\bar{w}} \quad (8)$$

In this case  $\bar{w}$  is the average of the new fitnesses as calculated via the matrix.

The game our population is playing is known as the “prisoner’s dilemma”, and the news is not at all good for cooperators, nor for the mean fitness of the population as a whole. The problem is that, while cooperators do better when there are more cooperators around, so do defectors. In fact, no matter what the relative frequency of cooperators and defectors, defectors do better. Consequently, unless defection is actually extinct, it will inevitably drive cooperation to extinction. The result is that the population has a mean payoff of 2, rather than of 4 (which it would have if every member of the population cooperated with every other).

As you may know, this result states a very basic problem for the evolution of cooperation. Illustrating the problem mathematically like this, however, has made it possible for us to understand a great deal about how one might get around the problem. Indeed, the evolution of cooperation literature has been quite productive in generating different kinds of solution to the problem of cooperation, which can then function as *empirical hypotheses* for real-life situations where cooperative behavior has stabilized. My hope here is that by approaching the analysis of biological knowledge systems in a similar spirit, a progressive enquiry can be started in that area as well.


As I said, there are a number of interesting ways in which one can modify the model so that cooperation stabilizes. The most well-known is the introduction of “sequential strategies”, in particular, “Tit-for-Tat.” The idea is this: one way in which our basic selection model is unrealistic is that we assume that individuals cannot choose who to interact with nor *how* to interact. Sequential strategies are strategies for sequential

interactions. In this model, we still assume that individuals interact at random, but that instead of just playing the game once with every individual they run into, they play the game a fixed number of times — say five. Now this repetition won't make any difference to the interactions between individuals that always cooperate and always defect. Enter Tit-for-Tat. Tit-for-Tat is just a little bit smarter than Cooperate and Defect. Tit-for-Tat can either cooperate or defect, depending on what the other player did in the previous round. Tit-for-Tat's strategy is to cooperate the first round and then on subsequent rounds do whatever the opponent did on the previous rounds. For a five round sequence, this gives us the payoff matrix in Table 3.

<i>W</i>	Always Cooperate	Tit-for-Tat	Always Defect
Cooperate	20	20	5
Tit-for-Tat	20	20	9
Always Defect	25	13	10

**Table 3**

You can work out for yourself the details of why this is the case, but it turns out that the presence of Tit-for-Tat can cause cooperative behavior to stabilize. The phase portrait for the Tit-for-Tat game is given in figure 3.4.



**Figure 3.4:** The Tit-for-Tat game with 5 iterated encounters.

It turns out that as long as the population starts out with more than 12.5% Tit-for-Tat, the population will evolve to some mix of Tit-for-Tat and Always Cooperate. Defection does well as long as it doesn't dominate the population, but when Always Defect's numbers become large, the fact that TftT does better against TftT than All-D does against TftT becomes the deciding factor.

This is not the place to pursue the analysis of the evolution of cooperation. My main purpose in demonstrating how the Tit-for-Tat game works has been to convey a sense of how population models can facilitate the very precise asking and answering of questions. The models are always very simplified pictures of reality, but that very simplicity can be helpful in isolating mechanisms of interest. That simplicity also makes it easy to duplicate results. If you don't believe that TftT has this effect, see for yourself. Despite the relative informality of my presentation, I have given you everything you need

to duplicate the model. Moreover, if certain simplifications involved bother you, the ability to duplicate *this* model makes it possible for you to extend *this* model. That makes it easy to see whether it was your modifications that had the significant effect, whereas if you built an alternative model from the ground up, this might not be so clear. My feeling is that this ability to duplicate results and incrementally add improvements to models is what lies behind the tangible progress we have made in understanding the evolution of cooperation. Obviously, I would like to see models of biological knowledge systems make the same kind of progress.

### **Selection-Mutation in Fixed Size Populations**

One of the problems we ran into in trying to apply the replicator idea to culture is that cultural entities, despite our tendency to think of them as tangible, are in fact more like patterns of behavior in populations of humans (or other animals), than kinds of things in their own right. We explored the question of how to think about their “transmission” in the last chapter, and it turns out that there is a non-standard population model that will be convenient for modeling processes of phenotypic variability.

Suppose we are having a disagreement about where human beings came from, and you convince me via an ingenious argument involving the probability of beneficial mutations that the Darwinian theory cannot possibly be right, and that the world and everything in it is in fact the result of a special act of creation by an all-powerful Deity. The population of interest here is the population of beliefs about the creation of the world and the life forms in it, and what has just happened is that the frequency of beliefs in Darwinism has just gone down a little bit, and the frequency of beliefs in special creation has just gone up a little bit. If this were not just an isolated incident, but a general pattern, we would have to say that this seems to be the kind of culture where special creation does better than Darwinism, and we could then wonder about why that might be. Obviously, the local culture is somehow discouraging the belief in Darwinism. But is the selection mechanism one that has to do with the ability of Darwinism to help us find our way about in the world, or is it due to factors more internal to our cultural life? Which of the two it is

---

will have something to do with whether we say that I have “learned anything” from our discussion, but the immediate issue here is how to model this process.

Notice that unlike the familiar cases of genetic evolution, what happened was that at the very moment I changed my belief, one “token” of the belief in Darwinism disappeared and another token of the belief in special creation came into existence.<sup>3</sup> It seems natural to say that there are not two events here, but one. What happened is that my belief regarding the creation of life changed type — it went from being the type “belief in Darwinism” to being the type “belief in special creation.” According to the terms we have been developing, what happened was that an individual belief changed type, which is to say, that it mutated. What is decidedly odd about this situation is that, unlike our mutation models above where there was selection on the “source” end to the mutation process and variation on the “destination” end, there seems to be selection on “both” ends here. This is the case because it was the inability of my belief in Darwinism to survive our conversation, combined with the attractiveness of your belief in special creation, that resulted in the mutation process. There is nothing “random” about this either. My belief mutated directly into a more (locally) fit belief just *because* it was locally more fit! Does the oddness of this general notion of mutation never cease?

Yes, I think this is the end of it. Mutation, as a process by which an individual changes type is a composite process, and needn’t involve variation at all. There is always selection at the source end, and there can either be selection or variation at the destination end. That’s it.

For our purposes, however, we won’t need to model mutation with selection at both ends since we aren’t going to model cultural transmission in this book. We have other fish to fry. What we do need to be able to do is model populations that are like our population of beliefs in the following respect: the population has a fixed size. The thing

---

<sup>3</sup> Belief individuation constitutes a challenging theoretical topic which I do not intend to contribute to here. What matters for our purposes is that we can expect to be able to count beliefs in a population if we need to. Sperber (1996) has argued persuasively that this is easier than one might suspect.

about beliefs about creation is that everyone has one, or at least that seems to be the most sensible way to model the population of beliefs. If anyone wants to claim that they don't have a belief we will assign them to the type "none." And if they are of divided opinion, we will give half of their full contribution to the belief-type frequency to one type, and half to the other.

The kind of fixed-size populations we are going to be looking at are ones where every member of the biological population is always performing one of a number of behaviors, and when the behavior token "destabilizes", it changes at random to some other behavior (or back to itself). It turns out that this is a fairly good way of representing navigation in simple organisms like bacteria and (to a lesser extent) bees.

In this kind of system (since there no transmission) the fitness of a behavior is always one or less. Selection proceeds as before by multiplying the fitnesses  $w$  times the frequencies  $p$ . But since the population is of a fixed size, instead of normalizing after selection (dividing the frequencies by the new total) we *redistribute* the surplus generated by selection. Moreover, since the "destinations" of the type change process are random, every type gets added to by the same amount, just their equal share of the total loss to selection (destabilization) in the population. The equation looks like this:

$$p_i' = p_i w_i + \frac{(1 - \sum_j p_j w_j)}{n} \quad (9)$$

Equation (9) just says that the new frequency of each type is the old frequency multiplied by the type's current local fitness (or "decay rate" if you prefer) plus the  $n$ th part of the population's total loss to destabilization. This dynamics has some very interesting properties as we shall see later on, the most important of which is the ability to track whatever is locally determining the fitnesses. It turns out that when the fitnesses are controlled by mechanisms that have been under selection at the genetic level, lots of useful information gets encoded in the distribution of behavior in an exploitable way. But there we're getting a bit ahead of ourselves.



**Sampling Error / Drift**

The forgoing construal of populations assumes implicitly populations which are "effectively infinite," a phrase which, while common in population genetics, seems vague enough to be worrisome, and thus requires some clarification here. If one has a fair coin, then the chance of it coming up heads is 50%, as is the chance of it coming up tails. This does not entail, however, that for any series of flips the proportion of heads will be 50%. Rather, it is consistent with the coin being fair that a series of 100 flips will all come up tails. It just won't happen very often. How often is, of course, the subject matter of statistics. On the other hand, any infinite series of coin flips will (almost certainly) converge toward 50% heads. This is one way of stating the "law of large numbers."

Recall that our basic Fisher-style selection operator takes  $\bar{p}$  to  $\bar{p}'$  as a function of  $\bar{w}$ , such that

$$p_i' = p_i \frac{w_i}{\sum p_i w_i} \quad (10)$$

The new frequency of types is the old frequency, times the fitness of the type, expressed as a percentage of the new population total. The fitness, recall, is just the expected value of increase for the type in the current environment, due to a certain subset of causes (those involving pre-existing tokens of the type). So if the fitness of type  $i$  is 3, then in all likelihood, the absolute numbers of  $i$ 's will triple in each cycle. This does not mean, however, that each individual is guaranteed to have three offspring (to use a biological example). It is consistent with a fitness of 3 for  $i$ 's to have a 90% chance of having no offspring, and a 10% chance of having 30 offspring. Moreover, just as with finite series of coin flips, these likelihoods do not guarantee that in a group of ten  $i$ 's, nine will have no offspring and one will have 30. It is consistent with these probabilities that all ten have thirty offspring, or that all have none. What having a fitness of 3 does entail is that for larger and larger numbers of  $i$ 's, the growth per cycle in the absolute number of  $i$ 's will with almost certainly converge to 3. So it is only in the case of *very* large, or "effectively infinite" populations, that there is any assurance that the population will behave as the

fitnesses, which is to say the environment, dictates.

What we are doing when we assume an effectively infinite population is simply ignoring sampling error for the purposes of the model. One might think that the only real excuse for making such an assumption is that it makes the models more manageable, and this it does. But Sober (1984) offers a more theoretical justification for such assumptions. He characterizes evolutionary theory as a "theory of forces." This means that one undertakes the analysis of a system's dynamics by isolating different kinds of influences on the system's behavior, and modeling their contribution to the dynamics in isolation. Newtonian mechanics is the paradigm here. We begin by modeling the effects of bodies in the absence of influences like gravity and friction, and model the effect of those influences in isolation as well. We then proceed to the combined effects of those "forces" in more complex models. In evolutionary theory, the main "forces" that effect the evolution of populations are selection, mutation, recombination, and so forth. "Drift," or sampling error, is also one of the forces that acts on populations, to be understood first in isolation, and then incorporated into more complex models as it becomes relevant.

We will largely continue to ignore sampling error in the remainder since its effect is mostly to introduce error into the informational dynamics that we are interested in here. The thing to remember is that, the smaller the population, the larger the effects of sampling error, and thus the less predictable the dynamics.

There are two other explanatory features of sampling error that should be noted before moving on. The first is Sewall Wright's "shifting balance theory" of evolution, in which the very unpredictability of the evolution of semi-isolated subpopulations ("demes") serves as an important source of genetic variation for the more inclusive population. Hull(1988), as discussed in Chapter 1 herein, appeals to this theory as a justification for tightly knit subdisciplinary groups in science, though unsuccessfully (or so I argued). Likewise, sampling error may provide "noise" sufficient to break symmetries in the evolutionary process, or to destabilize unstable equilibria (see Skyrms (1994)). The point to keep in mind is that as our models become more complex, sampling error may allow us to account for otherwise mysterious phenomena.

---

The second point is that, according to the way we are constructing our models here, extinction almost always depends on sampling error. For except in the unusual case when a type's fitness becomes zero, our selection dynamics will never eliminate a represented type entirely. Extinction of a type is a matter of eliminating the last token of the type, and as long as we assume that there is no sampling error (i.e. that our population is effectively infinite) we can never be down to the last token of any represented type. Furthermore, the event which results in extinction must be a result of sampling error (failure of the type's frequency to follow the specified dynamics) except, again, in the unusual case where the type's fitness is zero. On the other hand, the dynamics of information gain that we will be exploring do not depend on extinction. But we should keep in mind that anytime we assume that extinction has taken place, sampling error has been at work.

### Conclusion

This chapter began the task of building the tools required to model information transfer in multi-level selection processes. The central concept is that of a population — a typed collection of individuals. For such a collection, type frequencies may or may not change. If the frequencies change (including the introduction of new types), that constitutes evolution in the population. Evolution may be the result of selection, variation or both. Selection is the result of differential fitnesses which are summations of each type's local growth/stability properties (e.g., mortality, reproduction). Variation constitutes the balance of causes of frequency shift, those not due to individuals of the type previously in the population (e.g., emigration, gains from random mutation). These schematic characterizations of evolutionary components are not intended, by themselves, to explain the accumulation of complex adaptations. *That* accumulation is the result of the constraints that specific causal processes place on nature of selection and variation in specific populations. Reproduction, for example, overcomes the natural loss of adaptive variation characteristic of collections of ephemeral individuals (like pots in the garden) via the mechanisms of inheritance. The inaccuracy of those same mechanisms introduce new

variants at low rates which are similar to previous successful individuals, facilitating the gradual exploration of the space of adaptive solutions. Thus, the concepts explored in this chapter provide a framework within which one can say what it is about inheritance mechanisms that are so ideal for accumulating adaptive variation. But those questions are tangential to the point of this book, which is concerned with finding a way of understanding evolution that does not place unnatural restrictions on our characterization of culture, but which allows us to capture the momentary adapting tendency of environmental interaction in its fullest generality.

The concepts of selection and variation that emerge are broad in their application and minimal in their requirements in a way that may make them seem to verge on triviality. But consider: it is not having mass or location, but how much mass or what location an object has, that has explanatory power. Similarly, it is not having a fitness, but how much, and more importantly *why* a type has that fitness, that has explanatory power. Nor is it the simple fact that a population has sources of both selection and variation that explains anything, but why it has the particular sources it does, and why they behave the way they do. The evolutionary concepts presented in this chapter do not and are not intended to explain anything. They are intended, rather, to provide a framework within which explanations can be made. Our explanatory use of them will come later. Moreover, they are for the most part commonplaces of theoretical biology.

I should emphasize one important novelty introduced in this chapter. Mutation, the process according to which an individual changes type, was characterized not as a pure source of variation but as a composite process where selection (strictly speaking) operates on the source type resulting in the destabilization of the token and either random variation or selection increasing the destination type. (The example was changing belief based on an argument.) The latter possibility was forced on us by the need to accommodate populations of fixed size, which we can expect to be common in cultural transmission scenarios. I argued that this should not be taken as a dispute with ordinary biological usage, in which the selection on source types is negligible and in which Weismannian inheritance insures that mutational destinations are random with respect to

fitness.