1

INTRODUCING EVOLUTIONARY THINKING

Stephen C. Stearns

Human pathogens have evolved and are evolving drug resistance; emerging diseases have evolved and are evolving to infect us; we have evolved and are evolving in many ways, including symptoms, physiological responses, and behaviour patterns. However, medical research and medical education have paid relatively little attention to evolutionary biology. This book explains why medical doctors might want to consider evolutionary thinking as a standard part of their tool kit (also see Nesse and Williams 1994, one motivation for this book).

Evolutionary biology is a rich collection of well-developed alternative approaches to the interpretation of biological diversity and organismal design. It is not just selectionist thinking about adaptations, although it is certainly that. It is also the study of what genealogies and phylogenies can tell us about relationship and history. (The methods used to reconstruct history make use of the observation that much of the variation in DNA and proteins is neutral or close to neutral—that part of evolutionary biology is certainly not selectionist.) It is the study of conflicts: conflicts between hosts and parasites, between parents and offspring, between genes with different transmission patterns. Participants in conflicts often come off worse than they would have without the conflict. In such cases, evolutionary biology is as much the study of maladaptation as it is of adaptation. Evolutionary biology is also the study of genetic and phenotypic dynamics, regardless of whether they lead to adaptation or not. Sometimes they do not, and sometimes they cannot. No ideological monolith, it is rich in alternatives that can be played off against each other to provide a self-critical, well-tested, and increasingly reliable interpretation of the natural world, of which humans are a part.

This book describes how evolutionary thinking gives insight into human health and disease; this chapter summarizes evolutionary principles. We begin by describing how evolutionary biologists think, then give reasons why medical doctors might be interested in what evolutionary biologists have to say.

HOW EVOLUTIONARY BIOLOGISTS THINK

Evolutionary biologists want to understand how the variation in reproductive success that causes selection arises, how the genetic variation that enables a response to selection originates and is maintained, and how that response is constrained by geography, time, inheritance, conflicts, development, and history. Their insights apply to all organisms, including humans. Evolution contains many special areas of research in which different questions are asked and different approaches are used. There are many ways of thinking about evolution. Here are some of the more important ones. The views of each of these specialties are often useful alternatives to dogmas that develop in another.

Population geneticists think about microevolution, which occurs within populations over relatively short periods of time, about genes, about the different forms that one gene can take—its alleles—changing frequency or being held by various mechanisms in stable intermediate frequencies. Their central problem is what maintains genetic variation. Among the candidate explanations are natural selec-

tion, gene flow, and the drift of neutral genes. Population geneticists tend not to worry about the design of phenotypes. This area is well represented in medical research and education. For example, population genetic models of HIV make it clear that treating patients with several drugs at once is better than treating them with the same drugs sequentially, for such treatment prevents or delays the emergence of resistant mutants.

Evolutionary ecologists and anthropologists think about the design of phenotypes for survival and reproduction, particularly about traits like age and size at maturity, number and size of offspring, lifespan and ageing (these are all life history traits and this specialty is referred to as life history theory), various strategies of favouring offspring of one sex over the other (this is part of sex allocation theory), and the consequences of competition for mates and of choosing mates based on particular criteria (this is the field of sexual selection). They tend not to concern themselves with genetic details. This area has been largely missing from medical thought; it is represented here in Chapters 7, 8, 9, 10, and 22.

Molecular evolutionists think about history recorded in DNA sequences. Some of them examine parts of the genome that are not translated into proteins and that have no influence on the phenotype and processes occurring in non-coding DNA that have little to do with adaptive, selection-driven change. They tend not to worry about adaptive change, either in gene frequencies or in genotypes, except as distractions that need to be noted and controlled. This area is well represented in medical genetics and in Chapters 3, 15, 16, and 17.

Systematists think in terms of evolutionary trees, give great weight to history, and concentrate on variation among species. For them, the major evolutionary problem is to infer relationships among species so that they can reconstruct the structure of life on the planet, not to understand why gene frequencies change or how phenotypes are designed for reproductive success. This area may not appear directly relevant to medical research, but it has greatly aided the interpretation of sequence data from viruses and bacteria (Chapters 15 and 16). The methods of systematics were used, for example, to trace the persons responsible for the transmission of HIV in Florida and Sweden. It has become a part of forensic medicine.

Paleontologists think in deep time and concentrate on large-scale trends and major events, such as adaptive radiations, mass extinctions, and increases in body size within large groups over long periods of time. They often do not notice processes occurring within periods of less than a hundred thousand years, but they do see the big picture with particular clarity. This area, not directly relevant to medical treatment, provides useful cultural background. For example, it tells us how old, in evolutionary terms, the different parts of our bodies are: our hands are very old (hundreds of millions of years), our chin is quite young (less than a million years).

That was a brief description of evolutionary biology in terms of research specialties. The field can also be divided by analytical approach. Evolutionary change in a population of organisms of both sexes, with each individual containing thousands of genes affecting many traits with impact on survival and reproduction, is too complex to analyse in detail. There are four major ways of simplifying it. Each leads to a different way of thinking about evolution with its own advocates, its own school, and its own focus.

One approach concentrates on changes in gene frequencies, focuses on genetic mechanisms, and simplifies the interactions of phenotypes with their environment. This is the genetic dynamics approach that finds wide application in population and quantitative genetics and in sexual selection theory. It seeks to answer the question: how will gene frequencies change? A second approach seeks to explain the design of whole organisms at evolutionary equilibrium by analysing the interactions of the phenotypic traits contributing to reproductive success and by simplifying their genetics. This is the optimization approach used in life history evolution and behavioural ecology. It seeks to answer the question: what is the state of the phenotype at evolutionary equilibrium? The third approach is used whenever selection is frequency dependent, when the fitness of one thing depends on its relative frequency in the population, when success depends on how an opponent responds. This is the game theory approach, which has had great success in explaining the evolution of behaviour, sex ratios, and evolutionary conflict. It also seeks to answer the question: what is the state of the phenotype at evolutionary equilibrium? However, it allows for more complex interactions than does optimization theory. A fourth approach is used when we want to understand the history of a population by studying the patterns in its genes and comparing them with those found in related populations. This phylogenetic approach uses the information in DNA sequences and the logic of modern methods implemented on computers to reconstruct with increasing reliability the history of populations and species, human and non-human. It makes the assumption that the DNA variation used is neutral, close to neutral, or under equivalent selection pressures in all lineages. It seeks to answer the questions: what is the history of this set of populations or species, and how can we reconstruct that history by inferring their evolutionary relationships?

All four approaches are legitimate simplifications of a complex process. None of them retain all important features of that process. Therefore one should remember, when adopting one of them, to check the consistency of assumptions, interpretations, and predictions with those that would have been made had one adopted another approach. (This sounds easy in theory but is surprisingly hard in practice, for both technical and psychological reasons.)

EVOLUTIONARY EVIDENCE: SPECIAL PROBLEMS IN HUMANS

Evolutionary biologists are accustomed to dealing with kinds of inferential evidence that are not commonly used in medical research, including phylogenetic reconstructions of trees of relationship, comparisons across species and higher taxa, and descriptions of unmanipulated field populations. They use these kinds of evidence because they study processes that, in many cases, cannot be subjected to experimental analysis. Their work is often made more rigorous by quantitative mathematical models that make testable predictions and strengthen the interaction of ideas with evidence. Procedures of inference are now in principle reliable; their strengths and weaknesses are well understood. You will encounter the results in many of the chapters of this book.

Experiments are also done in many branches of evolutionary biology. In a particularly strong form of experiment, experimental evolution, the assumptions and predictions of evolutionary models are tested by creating conditions under which certain traits are expected to evolve, then seeing whether they actually do evolve. The expectations are shaped by a theoretical analysis that can be quite sophisticated. This approach requires organisms with short generation times that can be cultured in large numbers, and until now it has only been applied to bacteria, single-celled algae and protozoa, and fruit flies. The results (e.g. Elena *et al.* 1996) promise to strengthen the empirical basis of evolutionary thought.

Humans have long generation times and are not ideal material for evolutionary studies. Selection pressures strong enough to change gene frequencies measurably in a single generation are quite unusual, and with human generation times of 25–30 years, direct observations of gene frequency changes driven by selection are particularly difficult to obtain in our species. Nevertheless, there is enough inferential evidence to lend credibility to many evolutionary principles in humans.

Some of the best evolutionary evidence pertaining to medical research and practice has been obtained on infectious diseases, particularly the evolution of antibiotic resistance and the evolution of virulence. In both cases, evidence from human medical research is also widely discussed among evolutionary biologists because it is some of the best data available on problems they find of central interest, including the potential rate of evolution and the possibility of quite local adaptation.

REASONS FOR MEDICAL DOCTORS TO THINK ABOUT EVOLUTION

1. Each human individual has had a slightly different evolutionary history, and each has a different genetic make-up. This leads to differences in the way that different human individuals react to drugs and to diseases, differences that can result in life or death (Chapters 3, 4, 5, and 6).

2. Micro-organisms and cancer cells rapidly evolve resistance to drugs. This has important implications for drug design and treatment (Chapters 11 and 13).

3. The vaccination of a human population against a disease exerts a strong selection pressure on that disease; it will show an evolutionary response. Evolutionary analysis of vaccine design and application helps to reduce the chances of unpleasant surprise (Chapter 12).

4. Evolutionary theory tells us why virulence evolves to a certain level and no further and what measures could be taken to reduce it. Changes in our lifestyle, in treatment, and in public health measures could all cause virulence to evolve, for better or for worse (Chapters 14 to 18).

5. Why are so many sperm needed for fertilization? Why are so many eggs ovulated but discarded? Why do both the placenta and the ovary make apparently excessive quantities of reproductive hormones during pregnancy? And why are some fetal proteins derived only from the father's genes while others are derived only from mother's genes (genomic imprinting)? The answers may come from the evolutionary analysis of genetic conflicts (Chapter 7).

6. Human sexual behaviour, reproduction, and the assurance of parenthood are affected by evolutionary forces, often with consequences for the welfare of sons versus daughters. Some of the reasons for the neglect and abuse of children are evolutionary. Understanding why such things occur should help us to prevent them (Chapters 8, 9, and 10).

7. Symptoms may be adaptations. They may also simply be by-products, the reactions of organisms to novel environments. In either case, the best treatment requires understanding of why symptoms evolved (Chapter 2).

8. The problems of ageing result from selection operating on the whole human life cycle, from conception to birth to maturity to death. Because evolution operates on reproductive success, selection pressures drop with age and disappear in postreproductive individuals. Because up to a point more fitness can be gained by investing in reproduction than in maintenance that would improve survival, most organisms must evolve senescence. By understanding why we age, we can understand the consequences of treating the symptoms of ageing and of attempting to prolong life (Chapters 19, 20, 21, and 24).

NATURAL SELECTION

We distinguish between selection, which depends on variation in reproductive success, and the reaction to it, a genetic change in the population, which depends on inheritance. Organisms vary in how many offspring they have, whether they have them earlier or later in life, and in how many of their offspring survive to reproduce. Natural selection is simply this variation in reproductive success; it is necessary for adaptive evolution to occur. A second condition is also necessary. Only if some of the variation in reproductive success is inherited will a response to natural selection take place, for only then will the offspring from the better-performing parents inherit some of the capacity for better reproductive performance.

Natural selection is deceptively simple. Although it is hard to believe that such a mechanism could design an eye or a brain, it is incredibly powerful and rapidly produces highly improbable states. Consider the 31 letters in the sentence, THERE IS GRANDEUR IN THIS VIEW OF LIFE, as a sequence of 31 genes each with 26 alleles. If we assembled such sequences at random, we would have to sort through 26^{31} combinations of letters to find this one—as though enough monkeys typing long enough did eventually produce *Hamlet*, except in this case we are asking them to produce the first sentence of the last paragraph in *The origin of species*.

Note, however, that combinations of genes that are good in one context are often good in others, and inheritance of naturally selected variants will preserve them. Once part of the message is found, that part need not be lost. In our example, strong selection retains the correct letter whenever it occurs. If we start with any random sequence of 31 letters and retain all the letters that happen to be correct, then repeat the process by generating new letters at random for the ones that are not yet correct, we get to the right sequence in about 100 trials. This is 30 orders of magnitude faster than a random search. Natural selection is so efficient that in this case its performance was 15 orders of magnitude faster than the difference between the blink of an eye and the origin of the dinosaurs.

This apparently artificial example is readily translated into an experiment on the evolution of RNA in test tubes (Schuster 1993; Schuster *et al.* 1994). One can extract, from a virus that infects a bacterium, an enzyme that replicates RNA. Given some RNA and a supply of the four nucleotides from which RNA is made, this enzyme rapidly produces a large population of RNA molecules in solution in a test tube. At the start of the experiment, the lengths of the molecules, and the sequences of the nucleotides, are close to random. By transferring a drop of the solution into a new test tube every 30 minutes, one selects those RNA molecules that are present at highest frequency, for they are most likely to be transferred in the drop. Replication is quite good but not exact, and in about 1 in 10000 cases the wrong nucleotide gets substituted. Thus the growing population of RNA molecules becomes variable, and some variants are replicated faster than others. Two types of molecules have an advantage in this situation: small ones, and those with characteristic structures enhancing affinity to the replicating enzyme.

After many transfers—more than 100—a rather large, complex molecule dominates the population—which molecule in particular depends on the details of the experiment. One such dominant molecule, with a structure that greatly enhances its affinity to the replicating enzyme, is 218 nucleotides long, and it, among others, emerges repeatedly from such experiments. The chances of such a molecule appearing even once, randomly, are 1 in 4^{218} or 1 in 10^{131}. Since the experiment is set up in such a way that there are only about 10^{16} molecules in a test tube just before transfer, the procedure screens about 10^{16} molecules every half hour, and if the search were random, the chances of finding this molecule would be 1 in $10^{131}/10^{16}$ or 10^{115} per half hour of screening. This amounts to about 10^{112} years to find such a large, complex molecule using a random search. In contrast, selection produces the molecule that is best at getting itself replicated in about 2 days. It works because each step leads to a molecule that is better than the previous one. The improvements are preserved, and they accumulate.

Recall this example when people claim evolution cannot work because mutations occur at random. Although mutations do occur at random with respect to fitness, natural selection filters and preserves them remarkably efficiently. *Natural selection extracts order from randomness.*

The last example was pitched at the molecular level and helps to explain the origin of life; the origin could have been extremely rapid. Now consider an example involving whole organisms. If you were given the task of breeding race horses to run faster, you would select the fastest parents and breed them with each other. This method is effective and has been used for thousands of years to produce desired characteristics in domestic plants and animals. It makes an essential point: when some parents with certain characteristics have more offspring than others, those characteristics will spread if inherited. Artificial selection differs from natural selection in that the trait selected is determined by some human preference. *In natural selection the trait selected is always reproductive success.*

The design of organisms for reproductive success can only be changed by changing the

stored genetic instructions; genetic instructions are changed by natural selection within a population when organisms with different genetic instructions vary in their reproductive success. Thus evolution occurs both in information and in matter. Organisms function both as *replicators*, replicating the information they have stored in their genes, and as *interactors*, interacting with their environments and with each other to survive, reproduce, and get their genes into the next generation (Williams 1992).

ADAPTATION

A response to selection occurs whenever heritable variation in reproductive success improves reproductive performance. If this improvement continues for enough generations, a *process* called adaptation, it results in a *state* also called an adaptation. The state of adaptation is normally recognized in a particular trait, for example: the opposable thumb is an adaptation for grasping objects with the hand. Claims of adaptation can be controversial, for we usually have not seen the process that produced it, and there are alternative explanations. For starters, we define adaptation as a state that suggests to us that it evolved because it improved survival or reproduction or both.

Adaptations can be incredibly precise and complex. If our ears were any more sensitive, they would detect the random noise of Brownian motion in the atmosphere. Our eyes can detect a match struck at a distance of 15 kilometres on a dark night. Our intermediate metabolism would make the engineers who design oil refineries shrink with envy.

Natural selection operates whenever there is variation in reproductive success, and there is virtually always some variation in reproductive success. Therefore natural selection has always acted and is currently acting in virtually all populations, including our own. *A trait only experiences selection pressure if variation in that trait is correlated with variation in reproductive success, and it only responds to selection if some of that variation is heritable.* When both conditions are fulfilled, the *process* of adaptation begins. Whether it will ever result in the *state* we call adaptation depends on whether or not other factors are present that can constrain the response to selection. Since natural selection has acted and is acting on all traits that contribute to survival and reproduction, if such a trait is not well adapted, then something must be constraining its evolution. Several constraints are particularly important: gene flow, sufficient time, trade-offs, and historical accidents.

CONSTRAINTS ON ADAPTATION

Gene flow

Genes 'flow' from one place to another when organisms born in one place reproduce in another place, introducing their genes into the local gene pool. When natural selection favours different things in different places, movement of organisms transports genes that have been selected in one place to other places where they are not appropriate. Gene flow, like mutation, can introduce new genetic variants into local populations, and it can also produce local maladaptations. For example, the gene for sickle cell anaemia is adaptive where malaria is prevalent but maladaptive in other environments.

Despite gene flow, local adaptations do evolve when selection is strong, and selection is often strong. A classic example (Antonovics 1971) is heavy metal tolerance in plants. Plants on mine tailings grow on toxic soil and rapidly evolve adaptations to deal with it. Along a transect across a zinc mine tailing and into an uncontaminated pasture, the index of tolerance changed from 75 to 5 per cent in less than 10 m; the plants that were zinc tolerant flowered later, were smaller, and were more tolerant of inbreeding than those that were not zinc tolerant. Plants growing within 10 cm of a galvanized fence had significantly higher zinc tolerance than those just 20 cm from the fence. Zinc tolerance does not spread throughout the population because it is costly: when zinc is not present, it does not pay to be zinc tolerant.

Thus strong selection can produce very local adaptation despite gene flow; species often consist of genetically quite different populations each displaying different adaptations. For gene flow to prevent local differentiation, selection must be weak and the mean distance that genes move in each generation must be large. In modern humans this is often the case for many traits.

Sufficient time

With or without gene flow, it takes time for a population to adapt to an environmental change. Consider the absorption of milk sugar, lactose, by human adults (Simoons 1978; Durham 1991). Like all other mammals, human children come equipped with the biochemical machinery needed to digest milk. Most children lose that ability at the age at which they were normally weaned prehistorically: 4 years. A minority of humans retain the ability to digest fresh milk into adulthood, including the populations of northern Europe and some in western India and sub-Saharan Africa. The ancestral condition was the inability to digest fresh milk after the age of 4 years, and the new, recently evolved condition is the ability to do just that.

How long would it take that ability to evolve? The origin of dairying can be traced to between 6000 and 9000 years ago, and the ability to digest fresh milk after the age of 4 years has a simple genetic basis: it behaves as a single dominant autosomal gene. This is important, for dominant genes increase in frequency under selection much more rapidly than do recessive genes. Imagine a human population of about 10 000 people in which dairy milk production has begun. A mutation occurred that allowed people to utilize fresh milk after they were 4 years old. It had an advantage, for people who drink milk but cannot absorb lactose suffer from flatulence, intestinal cramps, diarrhoea, nausea, and vomiting, which reduces their reproductive performance. Lactose absorbers benefited from an additional high-quality food source, especially when other food was scarce and especially for nursing mothers and growing children. Selection for such a mutation could have been very strong during serious famines.

Suppose that the ability to absorb lactose conferred a selective advantage of 5 per cent, so that for every 95 surviving and reproducing children of non-absorber parents, the same number of absorber parents produced 100. Initially the gene was rare, and simply because it was rare, it could only increase slowly, for very few people carried it, enjoyed its advantages, and produced a few more surviving children than did those who did not carry it. As the gene increased in frequency, and more people carried it, it began to spread more rapidly through the dairying culture. However, when it became common, its rate of spread decreased, for then most people carried it, and there were very few who suffered from the disadvantage of not having it. How long did it take to increase from a single new mutation to a frequency of 90 per cent? The answer is about 350 to 400 generations or 7000 to 8000 years. This answer would change if we assumed a weaker selective advantage. If the estimated age of the milk-drinking habit is accurate, then it must have conferred substantial benefits if we are to explain its current frequency in northern Europe. Even for a gene under strong selection—and a 5 per cent advantage is strong selection—time is a constraint.

Trade-offs

A trade-off exists when a change in one trait that increases reproductive success causes changes in other traits that decrease reproductive success. For example, some organisms have to pay for improved reproduction with decreased chances for survival. This trade-off exists in the fruit fly *Drosophila melanogaster*, where flies that mate and reproduce have shorter lifespans than virgins. We do not know whether this trade-off exists in humans—it might—for we cannot do the kinds of experiments on humans that we can on fruit flies. However, the evolutionary design of human biology is certainly subject to other trade-offs. For example, the immune system protects us against pathogens, but it does so at the cost of

autoimmune diseases. Evidence from birds also suggests that one of the major costs of reproduction is a reduction in immunocompetence and increased susceptibility to pathogens and parasites.

Whenever one analyses some feature of organismal design in terms of the costs and benefits of changes in traits, trade-offs are involved. They place limits on how much fitness can be improved by changing traits, for most traits are connected through genetics, development, and physiology and cannot be changed by selection independently of one another.

Historical constraints

Organisms are not soft clay out of which natural selection can sculpt arbitrary forms. Natural selection can only modify the variation currently present in the population, variation that is often strongly constrained by history, development, physiology, and the laws of physics and chemistry. Natural selection cannot anticipate future problems, nor can it redesign existing mechanisms and structures from the ground up. We illustrate this principle with two examples.

The first concerns the vertebrate eye, often cited for its astonishing precision and complexity. It contains, however, a basic flaw (Goldsmith 1990). The nerves and blood vessels of vertebrate eyes lie between the photosensitive cells and the light source, a design that no engineer would recommend, for it obscures the passage of photons into the photosensitive cells. Long ago, vertebrate ancestors had simple, cup-shaped eyes that were used only to detect shades of light and dark, not to resolve fine images. These simple eyes developed as an out-pocketing of the brain, and the position of their tissue layers determined where the nerves and blood vessels lay in relation to the photosensitive cells. If the layers did not have the correct position, relative to one another, then the mechanisms that induce differentiation would not function correctly, for they rely on an inducing substance produced in one layer that diffuses into the neighbouring layer. Once such a developmental sequence evolved, it could not be changed without seriously damaging optical performance in the intermediate forms that would have to be passed through on the way to a more 'rationally designed' eye.

The second example concerns the length and location of the tubes connecting the testicles to the penis in mammals (Williams 1992). In the adult cold-blooded ancestors of mammals, and in present day mammalian embryos, the testicles are in the body cavity, near the kidneys, like the ovaries in the adult female. Because mammalian sperm (for some unknown reason) develop better at temperatures lower than those found in the body core, there was selection to move the testicles out of the high-temperature body core into the lower-temperature periphery and eventually into the scrotum (in some species they only drop into the scrotum during the breeding season). This evolutionary progression in adults is replayed in the developmental progression of the testes from the embryo to adult, and as they move from the body cavity towards the scrotum, they wrap the vas deferens around the ureters like a person watering the lawn who gets the hose caught on a tree. If it were not for the constraints of history and development, a much shorter vas deferens would have evolved that did the job just as well or even better.

THE NEUTRALIST–SELECTIONIST PROBLEM

Natural selection has been responsible for some of the variation in human DNA and proteins, but some of that variation has been caused by random drift, where selective advantages were small or lacking. Two mechanisms introduce randomness into evolution: mutations and meiosis. The randomness of mutations with respect to fitness was discussed above.

The randomness of meiosis consists of the 50 per cent chance that each copy of a chromosome has of getting into a particular gamete. Since only some gametes succeed in forming a zygote, developing, and reproducing, the random effects of meiosis are particularly important in small populations. This can be seen by the limiting case of a population of two individuals, one male and one female, that produces just two offspring. Consider a gene sitting on a chromosome in the female. It has just two chances to get into the next generation—one chance represented by each offspring—and each chance is determined by the flip of a fair coin. Thus even if a gene is strongly selected, there is a $0.5 \times 0.5 = 0.25$ probability that it will be lost from such a small population. As population size increases, so do the number of chances that each gene has of making its way into the next generation (many organisms will carry the same genes), and the effects of drift diminish. However, even in large populations, drift is the only force acting on portions of the genome that do not experience any selection pressure; thus drift is not an exclusive characteristic of small populations. However, it is only in small populations that drift is important enough to overcome the effects of strong selection.

In contrast to the randomness introduced by meiosis, the randomness introduced by mutations increases with population size, but not nearly fast enough to compensate for the declining effects of meiotic drift.

Current evidence is not sufficient to decide what proportion of genetic variation is caused by selection and by random processes. In humans, the amount of genetic variation to be explained is large: the human genome contains between 50 000 and 100 000 structural genes coding for proteins, of which about 30 per cent may be polymorphic. In many proteins only a few amino acids are critical to their function; substitutions at many other positions may be selectively neutral or close to it. On the other hand, the fact that no selective function is known for most human polymorphisms does not mean that selection has been absent, for modern civilization has changed our nutrition, our diet, and has eliminated or reduced many pathogens that were selective agents in the past. Not all functional differences among enzymes must have influenced fitness in the past, but without such an assumption it is hard to explain, for example, why the rare variants of polymorphic enzymes generally show lower activities. 'The neutral hypothesis, when applied to the study of human polymorphisms, might even have a counterproductive effect if it discourages the search for sources of natural selection' (Vogel and Motulsky 1996).

EVOLUTIONARY CONFLICTS

An evolutionary conflict occurs when genes have different patterns of transmission but interact, directly or indirectly, in the organism or organisms that carry them. A simple conflict occurs between the genes of hosts and the genes of parasites, especially parasites that can be transmitted horizontally, from host to host, as opposed to vertically, from host parent to host offspring. The host evolves mechanisms to reduce the damage inflicted by the parasite, and the parasite evolves adaptations to extract resources from the host and to improve the chances that its descendants will be transmitted to infect new hosts. The result is a coevolutionary arms race in which both species introduce measures and countermeasures in an open-ended escalation. Both host resistance and parasite virulence evolve, both are often costly, the costs of resistance and virulence ensure that both evolve to some intermediate level, rather than escalating open-endedly, and the result is a state of reduced adaptation in both host and parasite—reduced relative to the state we would have observed if the conflict had not been present.

Conflicts can occur between species, between relatives, and between genes within individuals. For example, mitochondria, which are inherited through the female line, and Y chromosomes, which are inherited through the male line in mammals, have different transmission patterns than autosomal nuclear genes and are therefore potentially in conflict with them.

Natural selection does not always produce adaptations. When evolutionary conflicts are not resolved, all parties suffer.

However, sometimes the conflicts are resolved. One case is suggested by the answer to the question: why are there no parthenogenetic mammals? In mammals, development requires one egg-derived and one sperm-

derived nucleus. Sperm- and egg-derived nuclei are marked differently by DNA methylation (genomic imprinting). Early development requires the expression of some genes derived from the father and some genes derived from the mother, which can be recognized by their sex-specific imprinting. If all the genes came from the mother, some genes would not be turned on at the right time, and development would fail—a strong constraint on parthenogenesis. Why did it evolve? It appears to be the result of a conflict between nuclear genes and mitochondrial genes that the nuclear genes have won. Mitochondrial genes have zero fitness if they occur in males, whose sperm transmit no mitochondria, and mutant mitochondrial genes that induce parthenogenesis increase in frequency. The strong constraint on parthenogenesis in mammals suggests that genomic imprinting may protect nuclear genes from subversive feminization by rogue mitochondria.

PRINCIPLES OF INFORMATION TRANSMISSION

All evolutionary change is based on genetic change; here is a short summary of the basic principles of genetic information transmission.

Sex versus asex

Information is transmitted in fundamentally different ways in asexual and sexual organisms. In asexual organisms (to simplify a bit, for there are many types of asexual reproduction), the entire genome is transmitted as a unit, and the only genetic differences between parents and offspring are the result of mutations. In sexual organisms (to simplify a bit, for there are many types of sexual reproduction), offspring are a 50:50 mixture of genes derived from each parent, and the primary consequence of sexual reproduction is to produce genetically diverse offspring. A sexually reproducing population generates much more genetic variation in each generation than does an asexually reproducing population, where the influx of new variation is limited by the mutation rate. Because the response to selection depends on the amount of genetic variation available, sexual populations may evolve more rapidly than asexual populations. They do so by combining from different parents favourable mutations that can be selected and disadvantageous mutations that can be eliminated by natural selection. Especially in small asexual populations, the favourable mutations must accumulate gradually, one after the other, and the opportunities for eliminating disadvantageous mutations are limited.

This may make it sound like human hosts, which are sexual, would have the upper hand over their principle pathogens, viruses and bacteria, which are asexual. For three reasons, this conclusion is too optimistic. First, some pathogens are actually sexual, for they frequently recombine genetically (e.g. *Neisseria gonorrhoeae*, which appears to be panmictic—see Chapter 16). Second, they have enormous populations, which greatly increases the probability that an advantageous mutation will arise in a short period of time. Third, they have short generation times, especially at 37°C inside the human body. Recombination, enormous numbers, and short generation times could combine to give pathogens an evolutionary advantage over their hosts.

There are also reasons why the pathogens do not always win. First, most hosts are sexual, and sexual reproduction plays an important role as a defensive measure that creates a moving genetic target (Jaenicke 1978). Second, all hosts have some kind of defence system, and vertebrate hosts like humans have very sophisticated defences in the immune system, which functions according to evolutionary principles on a generation time as short as that of pathogens. Third, finding a new host is an extremely risky business for pathogens; most die in the attempt. On balance, both hosts and pathogens have a fighting chance; neither can dominate for very long in evolutionary terms; but sometimes a pathogen can dominate long enough to drive a host population to extinction.

Meiosis conserves gene frequencies

The information stored in the genes is copied precisely both as a DNA sequence, where replication and repair are very accurate, and in the

population, where gene frequencies do not change from generation to generation if there is no selection, mutation, gene flow, or variation due to random sampling in small populations ('genetic drift'). The fairness of meiosis, the large size of most populations, and the accurate replication of genes as material units are what make information transmission in sexual populations a fundamentally conservative affair.

Why gene frequencies do not change from generation to generation is a basic principle, the Hardy–Weinberg law, that can be found in any genetics text. Here the important point is not why they do not change, but *that* they do not change. With Mendelian inheritance, the genetic variation necessary for a response to selection is preserved in populations, not destroyed. If gene frequencies changed very much, very often for reasons that had nothing to do with natural selection, then systematic change in response to selection would be impossible, beneficial changes could not accumulate, and adaptive evolution could not occur. The background noise would be too great for the signal to emerge (as happens in genetic drift). Thus the conservatism of Mendelian inheritance is the stable foundation of all adaptive evolutionary change in sexual organisms.

That is not only an important principle, it is also rather peculiar. Why did a complicated mechanism like meiosis that can be so fair to all the genes, distributing them with precisely equal chances to the gametes, ever evolve? Current evolutionary thinking suggests that the fairness of meiosis resulted as a defensive strategy of the nuclear genes to counteract the distorting effects of rogue genes that have the effect of overrepresenting themselves in the gametes at the cost of other genes—meiotic drivers (Hurst 1992).

SELECTION DESIGNS PHENOTYPES FOR REPRODUCTIVE SUCCESS

Fitness is relative reproductive success

The basic insight of population genetics, that the evolutionary process can be reduced to the analysis of the factors that increase or decrease the number of copies of a gene in a population from one generation to the next, has great simplicity and power. It is a good starting point. However, focus on gene frequency change is not sufficient to explain phenotypic evolution. If we want to understand why organisms are designed in some particular way for reproduction and survival, then we must analyse the organism as an interactor representing the genes in all activities contributing to survival and reproduction—as an organism with an ecology, with food to find, predators and parasites to avoid or combat, and mates to convince: as a phenotype with a certain lifetime reproductive success, with a certain fitness.

Natural selection has several components

The analysis of reproductive success begins with the factors determining the number of offspring produced by a single individual over its lifetime. This is the most general component of reproductive success, *individual fitness.* Selection on offspring number per lifetime is called *individual selection.* In growing populations, offspring produced earlier in life contribute more to fitness than offspring produced later in life, for they produce grandchildren earlier, and the effects of shorter generation time accumulate multiplicatively. Thus individual fitness is not just lifetime reproductive success; it often depends on the timing as well as on the amount of reproduction. This type of selection is often sufficient to account for the states of many traits which appear to be shaped to increase the number of offspring per individual per lifetime, and it is sufficient to account for ageing and senescence.

In sexually reproducing organisms, individual selection contains an important component associated with mating success, with interacting with a partner of the opposite sex to produce offspring. This component of natural selection is called *sexual selection.* Traits under sexual selection are subject to evolutionary changes that improve mating success but may reduce survival. For example, the male peacock's tail improves his reproductive success by making him attractive to females but

reduces his chances for survival by making it harder for him to fly. Sexual selection involves the two sexes in a complex interaction with fascinating properties. Females have evolved preferences for certain kinds of males and by mating with them transfer the preferred traits to their sons and their preferences to their daughters. The results can be surprising. Mothers may pass on their preference for mates that take risks to their daughters, while their sons inherit the risk-taking trait itself. This would explain why car insurance rates are higher for 20–25-year-old males than they are for females (Daly and Wilson 1985).

Organisms living in groups of related individuals experience a third kind of selection, one that has resulted in deep insights. *If what matters to evolution is the relative number of copies of genes that exist in the population in the next generation, then it does not matter through whose reproductive activities those genes were replicated—directly, by an individual, or indirectly, by its relatives.* Thus if an individual can influence the reproductive success of its kin, it should do so if the benefits—the increase in number of genes in the next generation through the reproductive activities of relatives—exceed the costs—the reduction in the number of genes in the next generation it gets through its own reproductive activities (Hamilton 1964). This is called *kin selection*, and it has helped us understand the evolution of apparently self-sacrificial, co-operative, altruistic, and other kin-related behaviour. Its empirical success has convinced most evolutionary biologists that their focus on genes is probably correct (Williams 1966; Dawkins 1976).

The gene-centred point of view also explains why senescence is a property of the soma, not of the germ line. If the gene-centred view is correct, then evolution 'cares' about the germ line—the genes—whereas medical doctors treat the soma, which is, from the point of view of evolution, disposable.

Traits do not evolve for the good of the species

Formerly one often heard that some adaptation had evolved for the good of the species, helping it to avoid extinction. As a general explanation this is fundamentally wrong, and that statement represents the broad consensus of the evolutionary community. Traits evolve because they improve the reproductive success of individuals and their kin, and if the species to which those individuals belong happen to survive longer because of those changes, their longer survival is a by-product of the essential process and not the reason for it.

This insight was achieved through a fascinating episode whose main result can be summarized in a single phrase: *selfish mutants invade.* If a trait did evolve that benefited the species at the cost of the individual, some mutant that selfishly exploited the more altruistic individuals would invade and take over the population. It could do so because selection on individuals is much stronger than selection on species. Individuals have much shorter generation times than species, and in the time that it takes for new species to form and go extinct, a process spanning many thousands of individual generations, hundreds of millions of the individuals that form those species have lived and died. For that reason, selection has much greater opportunity to sort among individuals than it does to sort among species, and species selection simply cannot shape adaptations (Maynard Smith 1964; Williams 1966).

BIOLOGICAL CAUSATION: PROXIMATE AND ULTIMATE

Biologists want to understand all the features of living organisms. One natural approach is to study the immediate causes. How does respiration work? What determines the sex of an organism? What causes senescence? These are a few examples of questions about immediate causes answered by physiology, genetics, biochemistry, development, and related fields. Here the aim is to identify the factors that cause the trait or process during the lifetime of a single organism through the study of *mechanism* or *proximate causation*. Much of biology is devoted to it.

Evolutionary biologists ask different questions and investigate different kinds of causes. Why does respiration occur in the mitochondria and not in the cell nucleus? Why do most species have approximately equal numbers of males and females? Why do many animals senesce, but many plants and fungi hardly at all? Why are so many sperm needed to traverse the female reproductive tract, and why do millions of ova develop when less than a thousand are ovulated in the course of a female's lifetime? These are some questions asked by evolutionary biologists. They are also questions about causation, but on a time scale of many generations and at the level of populations and species rather than individuals. This is the study of *evolutionary* or *ultimate causation*. Whereas in mechanistic analysis the causes can be described as biochemical and physical processes, in evolutionary analysis one often describes the causes as how natural selection, evolutionary conflicts, historical contingencies, or chance events shaped the trait under study.

All traits have both types of causes; therefore a complete biological explanation demands the analysis of both. It would be a strategic error to isolate the two kinds of analysis from each other. We should be able to see the world both ways—from the bottom up, from molecules to populations, and from the top down, from selection to molecules.

SUMMARY

Some key points:

1. In natural selection the trait selected is always reproductive success.
2. Natural selection has great power to shape precise adaptations; it can rapidly produce highly improbable states.
3. Natural selection does not always lead to adaptation. In situations of evolutionary conflict, it is often the case that all parties suffer.
4. Traits do not evolve for the good of the species.
5. Natural selection cannot anticipate future problems, for evolution proceeds by tinkering with what is currently available.
6. It does not matter through whose reproductive activities genes are replicated—directly, by an individual, or indirectly, by its relatives.
7. Because selfish mutants invade, arguments that traits evolved for the good of the species are usually invalid.

Evolution combines with physics and chemistry to explain all biological phenomena, and it is the only part of biology containing basic principles not implicit in physics and chemistry. It has three major principles—natural selection, inheritance, and history—and one fundamental property: selection acts on organisms, but the response to selection occurs in stored information. The first principle, natural selection, is a great law of science, the only mechanism known that can maintain and sometimes increase the complexity of organisms, extracting order from randomness, producing systems organized to overcome, locally, the dissipative effects described by the second law of thermodynamics. Variation among organisms in reproductive success produces natural selection; this happens in physical and chemical material. Populations of organisms respond to selection when some of that variation is genetically based; this happens in stored information. The result is a genetically based change in the phenotypic design of offspring from the more reproductively successful parents. Genetic changes also occur at random as mutations and as the consequence of meiotic sampling and they persist in parts of organisms that are not under selection and in small populations.

Those processes have produced the complexity of all living creatures.

ACKNOWLEDGEMENTS

Comments by Dieter Ebert, Ed LeGrand, Randy Nesse, Beverly Strassmann, Jacob Koella, Arno Motulsky, Andrew Read, and an anonymous reviewer improved a draft.