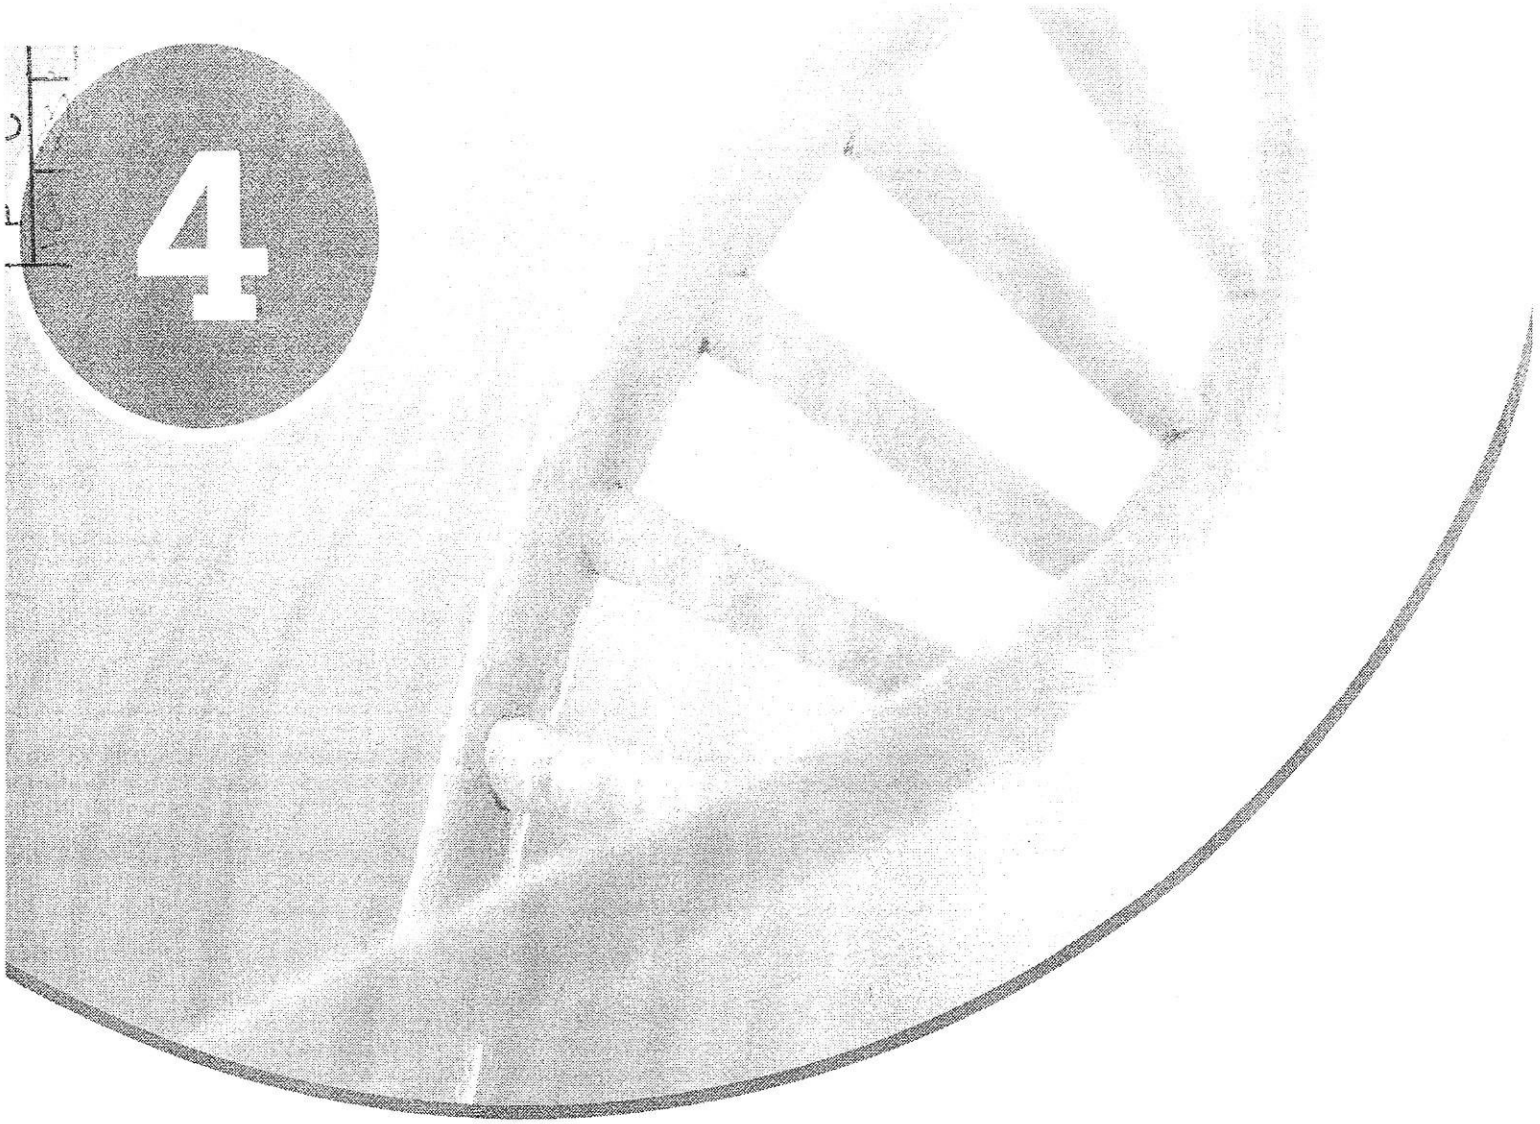

4 Competition

We have now examined two of the pre-conditions for natural selection: the tendency of individuals to be slightly different from one another (Chapter 2, Variation) and the propensity for variation to be passed from parents to offspring (Chapter 3, Heredity). These two principles will only generate natural selection if another condition is met: there must be competition to reproduce. This chapter introduces competition, examines the level at which it is best thought of as operating, and looks at what types of evolutionary outcomes it tends to favour.

4.1 Malthus: checks on reproduction and competition to reproduce

For the principles relevant to competition, we need to look beyond Darwin to an English thinker called Thomas Malthus (Figure 4.1). Malthus published a famous book called *An Essay on the*

Figure 4.1 Thomas Robert Malthus (1766–1834). *With kind permission of The U.S. National Library of Medicine.*

Principle of Population in 1798. Darwin had read the essay and felt it to contain a fact of fundamental importance. However, Darwin's reliance on Malthus led to much misunderstanding of the implications of Darwin's thought. This is because Malthus' book contains many claims, only some of which are relevant to natural selection. In particular, we need to separate *general principles* that Malthus derived about the dynamics of populations from *specific claims* that Malthus made about the future of the human population in Europe. Only the former are necessary for understanding natural selection.

The general principles in Malthus' essay that we need be concerned with are that populations could potentially grow exponentially, but in practice cannot do so, and therefore must be limited by incomplete survival and/or reproduction. This means that there is competition between members of the same population to be in that fraction which manages to survive and reproduce. These are the principles that Darwin realized were key to his theory. Malthus' other claims, for example that the European population was expanding too fast for the available resources, that wars and famines were in prospect because of this, and that the lower classes should be discouraged from reproducing because of the poverty that their increase would generate, are not in any way part of the theory of evolution. Indeed, their factual aspects have not stood the test of time. The European population after Malthus continued to expand, but famine actually became less common, not more so, because agricultural productivity and economic growth went up even faster. Moreover, people in Europe eventually responded to the increased affluence by spontaneously decreasing their family sizes, leading to a stabilization of the population. However, none of this is in any way important for the general theory of natural selection.

Exponential population growth: an example

Let us then examine the key Malthusian principle that is correct—exponential potential population growth. To see what we mean by 'exponential growth' we will take as our example not humans but cats. We will assume the following, although the exact details are not critical: that a female cat can have two litters a year, of six kittens at a time; that the kittens are roughly 50% females; that the kittens mature sexually at 1 year of age; that cats can go on reproducing until they are 5 years old. We start from a single breeding pair.

You can probably see the trend that the population will follow. In year 1, the breeding pair produces 12 kittens. About six of these will be females. Thus, in year 2, there will be six new females plus the original one, giving seven. Seven females can produce 84 kittens, about 42 of which will be female. Thus, in year 3, there will be 42 plus seven, or 49, breeding females, who will produce 588 kittens between them. The number of kittens grows dramatically larger each year. From year 7, some of the older breeding females begin to die off, but this has a negligible effect. By then there are over 100,000 breeding females anyway!

If we graph the size of this hypothetical cat population, we see that it does not just increase, but increases at an ever-increasing rate (Figure 4.2a). Such a pattern is called exponential growth. Under exponential growth, the population would explode very quickly. Figure 4.2b illustrates this by calculating the number of years our cat population—starting from a single pair, remember—would take to reach various milestone sizes. Within a dozen years, the number of cats surpasses the current world human population, and within a century there are more cats than the estimated number of atoms in the observable universe. Of course, this could never actually occur, but it serves to illustrate what would happen if a population increased at anything like its theoretical maximum rate. The nature of exponential functions is such that it does not much matter which species we choose. Cod, which produce several million eggs a year, would reach the milestones a few years earlier and chimpanzees, with one offspring every few years, several decades later. However, this does not detract from the general pattern. Darwin himself calculated that even elephants, the slowest-breeding animals he could find, would ultimately follow a similar pattern, with one breeding pair having 19 million descendants in around 750 years.

Figure 4.2 **(a) Theoretical exponential growth of a cat population starting from a single breeding pair. (b) Time taken in years for such a population to reach various size milestones.**

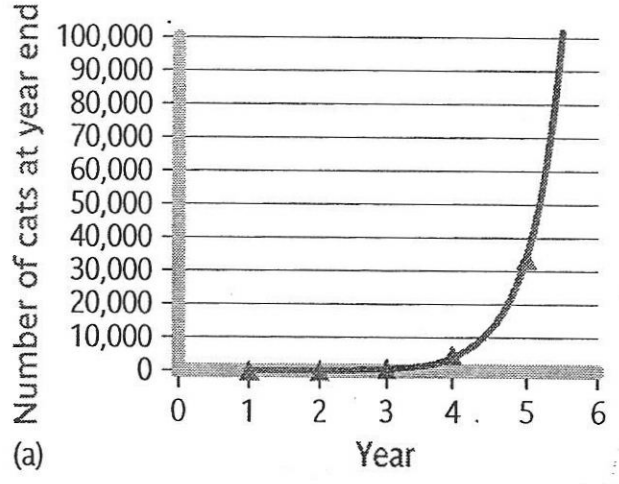

(b)

Milestone	Number of years
One thousand cats	4
One million cats	8
Greater than US population	11
Greater than world population	12
One million times world population	20
Atoms in the observable universe (estimate)	96

Figure 4.3 Effects of adding infant mortality, adult mortality, or both to the growth of the cat population.

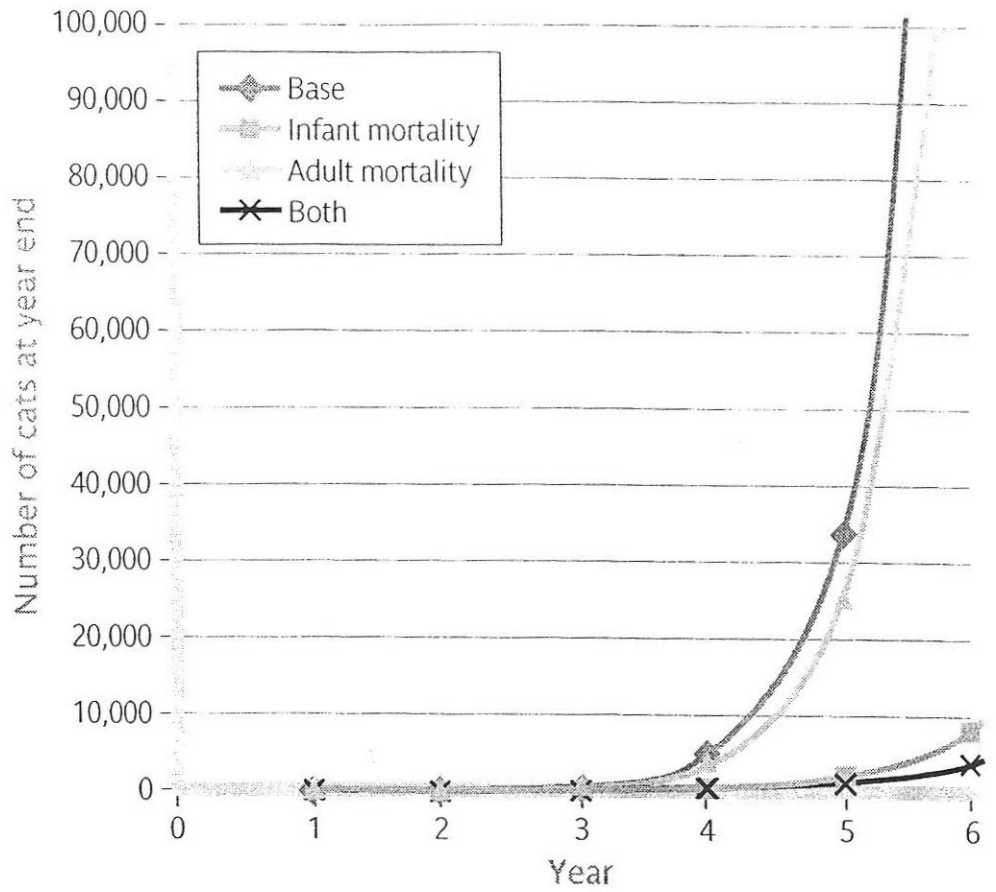

Populations in nature do not, of course, follow such a pattern. If they did, the world would become excessively crowded with organisms in a few years. Populations are sometimes stable, sometimes contract, and sometimes grow, even quite fast, but they never attain anything remotely like the exponential pattern they could theoretically achieve. Why not? You will note that in the cat example, I assumed that all kittens conceived survived into adulthood and also that all adult females survived for the whole length of their breeding lifespan. What happens if we alter either of these assumptions?

Figure 4.3 shows how the population growth is changed by either giving each kitten just a 50% chance of surviving into adulthood (the infant mortality condition), or only a 50% chance of surviving from one breeding season to the next, once it is an adult (the adult mortality condition), or both. You will see that where there is both infant and adult mortality operating, the growth of the population is dramatically flattened off. The curve begins to look a little more realistic.

Mortality is not the only way the effect could be achieved. A similar flattening would follow if, for example, we made some cats unable to find a mate or to conceive. To flatten the curve just requires that for some reason or other not all individuals are reproducing at their potential rate. Malthus understood that such 'checks' on exponential growth must be operative in all biological populations.

4.1.2 There is a differential reproductive success

A biologically realistic model of the cat population requires, then, that not all individuals survive and reproduce as much as they might. Look at it from the point of view of all the kittens conceived. Not every one of them can make it into that much smaller fraction that will in turn have kittens. This means that the population in generation $n + 1$ is not a complete reflection of the population in generation n. Some cats have become ancestors of cats in the next generation and some have not. Another way of saying this is that there is competition amongst the cats in any particular generation to leave descendants in the next. Cats will thus differ in their reproductive success. Reproductive success is the number of viable descendants produced and so is obviously zero for any cat that dies at birth or in kittenhood.

4.1.3 Differences in reproductive success can lead to changes in the population

Natural selection requires that there is differential reproductive success, but differential reproductive success is not sufficient to produce any evolutionary change. To see why, let us introduce a further refinement to the model. Start with 100 cats, 90 of whom are black and 10 yellow. We will simplify things for now and assume that females always produce kittens of the same colour as they are. First consider a scenario where the mortality rate is 50% for the yellow cats and also 50% for the black cats. Although individuals vary in reproductive success, this has nothing to do with their colour and thus the proportions of the two colours in the population remain the same (Figure 4.4, black line). Where reproductive success is unrelated to a phenotypic characteristic, then that phenotypic characteristic does not change over the generations.

Figure 4.4 **Proportion of cats that are yellow over 300 years in the cat model, starting with 90 black and 10 yellow cats, where infant mortality is 50% for both black and yellow cats (black line) and where infant mortality is 50% for black cats and 49% for yellow cats (blue line).**

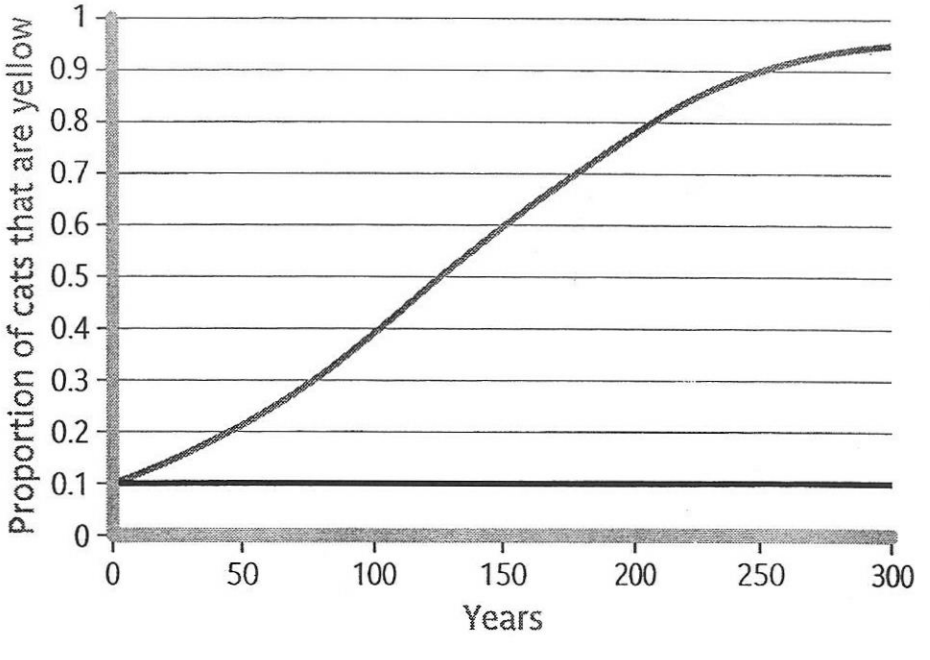

Now imagine the population moves into an open sandy environment. The yellow kittens are now slightly harder to see against the background and so slightly less likely to be killed by predators such as eagles. Eagle predation is only one source of infant mortality and the difference in camouflage is modest, so we will set the infant mortality rate for yellow kittens just slightly lower than for black ones—49% versus 50%. This means that yellow females in generation n, through greater survival to breeding, have a just slightly elevated chance of leaving offspring in the next generation, offspring who in turn will have a slightly elevated chance of leaving their own offspring, and so forth.

Observe what happens. The proportion of yellow cats increases generation on generation, until after about 300 years, almost all cats are yellow (Figure 4.4, blue line). Natural selection has driven a process of adaptation to the sandy environment. What we mean by this is that someone coming along at the end of the 300 years would be struck that these cats look well designed for the sandy environment in which they live. We know, of course, that there has been no designer, simply that the regime of differential mortality has caused a gradual change in the composition of the population, from mostly black cats to all yellow ones.

Note a few key features of the adaptive process. The sandy environment did not cause the yellow mutant to come about. Mutation is unrelated to the demands of the environment and in this example we just assumed that yellow was already around in the range of colour variation of the population. No individual cat changed in any way during the adaptive process. Black cats always had black offspring and yellow cats always had yellow offspring. However, whereas at the beginning we had a type of cat that was basically black, at the end we had a type of cat that was basically yellow. Note also that it took some time for adaptation to occur. Three hundred years is longer than the life of any one cat. On the other hand, it is a blink of an eye in biological timescales. Given the tiny survival advantage we gave the yellow cats (mortality of 49% rather than 50% in the first year of life, the same life history thereafter), it is extremely impressive how fast the adaptive process can occur. If we had made the competitive advantage of yellow cats bigger, then adaptation would have been even faster.

4.2 Natural selection at the genotypic level

In section 4.1, we examined the effect that natural selection was having at the level of the phenotype of the cats. This is a reasonable enough approximation to start off with. However, although success in competition is related to phenotypic characteristics, what natural selection is actually doing is changing the frequencies of the underlying alleles and so that is the process we should really model. This changes things somewhat, since the mapping between genotype and phenotype in a sexually reproducing population is not one-to-one (as we learned in Chapter 3). In this section, we examine how natural selection changes frequencies of alleles. We use the same cat example, but because we are now considering genotypes rather than phenotypes, we have to set up the model in a slightly different way. We track generations rather than years. Since there were two litters of kittens per year, we can think of each year in the previous model being approximately equivalent to two new generations in the current one.

4.2.1 Increase in frequency of an advantageous dominant allele

Let us assume that the colour of the cats is controlled by a single gene and that the allele A, which is dominant, gives the yellow colour, whereas *a*, recessive, gives black. The frequency of allele A in the population is equal to *p* and we will start where $p = 0.1$. This means that 10% of all copies of the gene in the population are A and thus that the other 90% are *a*. Where mortality is the same for yellow cats as for black cats, the allele frequencies do not change over the generations. This just follows from the Hardy–Weinberg principle (Chapter 3), given that genotypes AA and *Aa* (yellow cats) and *aa* (black cats) all have the same expected reproductive success.

Now let us return to the case where infant mortality is 49% for yellow cats and 50% for black cats. What we need to calculate is the fitness of each of the two alleles. The fitness of an allele is the number of copies in the next generation that a copy in this generation leaves. Where there are two alleles in competition with one another, their relative fitness determine how their proportions change over time.

The fitness of the dominant allele A is completely determined by the reproductive success of yellow cats, since all the cats in which A appears are yellow. The recessive allele *a*, by contrast, appears in yellow cats with genotype *Aa* and also in black cats with genotype *aa*. Thus, to calculate the fitness of *a* we need to do a weighted average of the reproductive success of the black cats it appears in and the reproductive success of the yellow cats it appears in.

Figure 4.5a graphs the way the allele frequencies change over the generations, given the slight survival advantage of yellow cats. It is the same general shape as in Figure 4.4, but you will note that the increase of allele A is a little slower than when we just considered phenotypes, particularly once A is common. This is because the allele *a*, although it is the allele 'for' black coats, spends quite a lot of its time in yellow cats (i.e. heterozygotes of genotype *Aa*), particularly when allele A is common. It is only exposed to its selective disadvantage, as it were, on the occasions when it is homozygous. The rest of the time it 'hides from selection' in the bodies of yellow cats. This means that it takes longer for natural selection to weed it out than it would otherwise.

Figure 4.5 **The change in frequencies of cat coat colour alleles over the generations, where yellow coats have a survival advantage. (a) The black line is the frequency of a dominant allele causing yellow coats and the blue line is the recessive allele associated with black. (b) The black line is the frequency of a recessive allele causing yellow coats and the blue line is the dominant allele associated with black.**

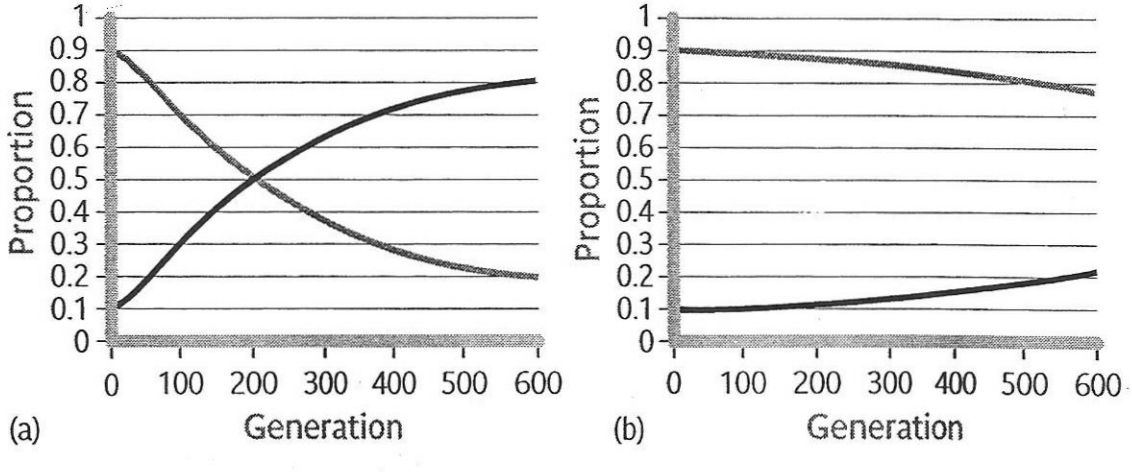

4.2.2 Increase in frequency of an advantageous recessive allele

Figure 4.5b considers the case where the advantageous yellow coat is produced by the recessive allele. The increase in frequency is much slower than the dominant case (although it would eventually become common if we ran the model for longer). This is also due to an imperfect association between alleles and phenotypes. Since the 'yellow' allele is recessive, it mainly appears in heterozygous black-coated cats, especially when it is rare. It can only show its advantage to selection, as it were, on those occasions where it happens to be homozygous and thus causes the cat's coat to be yellow.

4.2.3 Follow the alleles

The difference between Figure 4.5a and 4.5b illustrates a fundamental point that will recur in future sections. Natural selection changes allele frequencies, increasing the frequency of those alleles with high fitness and decreasing the frequencies of, and ultimately eliminating, alternative alleles with low fitness. The fitness of the allele is the weighted average of the relative reproductive success of all the different phenotypes it appears in. We will see in Chapter 5 that this means alleles with harmful effects in some individuals can persist, as long as they are causing reproductive success benefits in other individuals who also carry them. It is the *alleles* with the highest fitness that natural selection will preserve, not necessarily the *organisms* with the highest fitness.

This means that when we are trying to predict evolutionary outcomes, we should always calculate the relative fitness of *alleles*, not organisms, species, or any other unit. In Figure 4.5, the advantage *to the individual* of having a yellow coat is exactly the same in (a) and (b). Just by thinking about individuals, we would have predicted the same result in the two cases. However, when we do our accounting at the level of the alleles involved, we make the correct prediction that yellow coats will increase much more slowly when the allele for them is recessive than when it is dominant.

4.2.4 Competition revisited

We have just seen that the allele should always be the unit of accounting in thinking about evolution. Adaptive evolutionary change—for example cats in a sandy environment becoming yellow—is a consequence of one allele defeating another in competition within the gene pool of a population. This is the part of Darwinian theory that is most often misunderstood. People tend to think that the competition that drives adaptation is competition between different species, or between different populations, or between different individuals. All of these types of competition do exist in nature, but the relevant question for the evolutionist should always be: why would the alleles for [the phenomenon of interest] have out-competed alternative alleles in the same gene pool?

4.2.5 Natural selection and polygenic characteristics

The cat example considered so far is a single-gene characteristic. When we consider alleles that contribute to a polygenic characteristic such as height, things become a little more complex. If many genes affect height, an allele X whose phenotypic effect is to increase height by some amount will appear in all kinds of phenotypes, including some phenotypes that are shorter

than average. (If you need to review why this is, look back at Chapter 2, section 2.4.) However, the phenotypes in which X appears will on average be taller than those in which it does not and thus, if height is advantageous in a particular environment, X will gradually increase in frequency. Since, where tallness is advantageous, *all* alleles whose average effect is to increase height increase in frequency, the average height of the population increases over time. Thus, for polygenic characteristics too, differential reproductive success can lead to adaptive change in phenotype driven by changes in the underlying allele frequencies.

A general principle for thinking about change in polygenic traits as a result of natural selection is that the response to selection is the product of the selective pressure and the heritability. That is, evolutionary change gets faster with increasing selective advantage of whatever trait is being selected and also gets faster the more heritable the trait is (and, of course, the trait must show *some* level of heritability for there to be any evolutionary change at all).

4.3 Group selection

The power of natural selection is that, through a blind and goalless process of differential fitness of alleles, it can gradually create structures that appear well designed for their environment. This leads on to a difficult question, however: well designed for whom? For example, a tendency to attack neighbours might be a good design for an individual animal, but very bad for its neighbours. For whom should we expect adaptations to be optimized: the organism, the group, the population, or the species?

This issue is known in biology as the levels of selection debate and much has been written about the various approaches to it. For our current purposes, two points are important. First, it is generally best to do the accounting of advantages and disadvantages in terms of competing alleles, as we did in section 4.2, rather than any higher unit. The individual organism can often be used as an alternative, but this is best seen as a kind of approximation, for reasons we have partly touched on and to which we return below. Second, behaviours that promote the collective advantage of groups of organisms at the expense of individuals within the group will not usually evolve. This section shows why this is the case.

The idea that behaviours might exist because they benefit the group (e.g. the colony, the herd, the tribe, the population, or whatever) rather than the individual organism is known as group selection and it was widespread in biology until the late 1960s. The following example is from Wynne-Edwards (1962).

4.3.1 Wynne-Edwards and reproductive restraint in birds

In years when food is scarce, birds of various species lay a smaller number of eggs than they do in years when food is abundant. This is called reproductive restraint. Wynne-Edwards' interpretation of reproductive restraint was the following. If too many chicks hatch for the resources available, then there will be competition for food and all the chicks will be stressed. This can mean that the entire population dies out. Populations thus survive better if the birds within them restrain their reproduction, and reproductive restraint is thus an adaptation that has evolved because it increases the likelihood of the population surviving through bad years.

This argument is wrong, but to see precisely why it is wrong, we need to set up a model. Let us define a type of bird called an 'altruist' who, when times are hard for the population, limits

Figure 4.6 (a) The number of populations surviving over time containing all selfish individuals (dark blue line) or some altruistic individuals (light blue line). (b) The proportion of altruistic individuals over time within a population that starts out mostly altruistic.

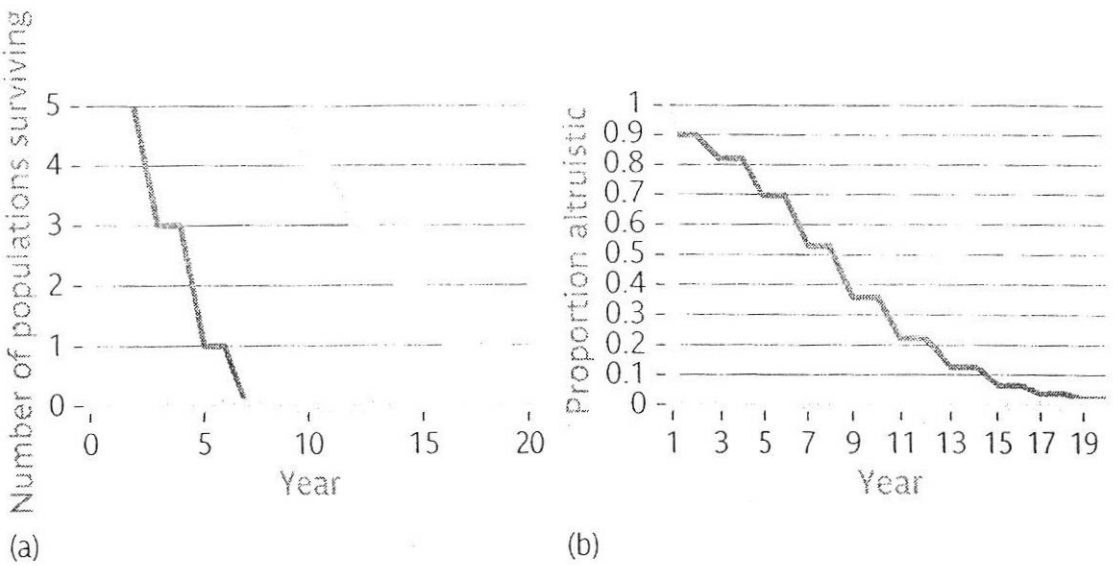

the number of eggs it lays. Populations which contain altruists will suffer less stress and will thus be less likely to go extinct in bad years than populations made up of 'selfish' individuals, who always have as many chicks as possible. For Wynne-Edwards' argument to work, 'altruism' as defined has got to be able to out-compete 'selfishness' as defined.

Let us set up ten populations of ten breeding female birds each. (For simplicity, we do not worry about the males in this scenario and we assume females only live for one breeding season. These assumptions are not important for the conclusion.) Populations A, B, C, D, and E are all selfish. Populations F, G, H, I, and J all contain one selfish and nine altruistic females. In good years, all birds lay two eggs and all chicks survive. In bad years, selfish females lay two eggs, whilst altruistic females only lay one egg. In addition, we stipulate that, in a bad year, if less than half the population practises reproductive restraint, there is a 50% chance that the entire population dies out. We will alternate good and bad years.

Figure 4.6a shows that populations containing altruistic individuals (populations F–J) are indeed more likely to survive than populations A–E. This is as Wynne-Edwards argued. However, now consider what is happening *inside* any one of the populations F–J. The selfish individuals within these populations increase at a faster rate than their altruistic fellows. This is because they are laying two eggs every year, whereas the altruists are laying only one in the bad years. Thus, over time, populations F–J contain fewer and fewer altruists, until there are no altruists at all within them (Figure 4.6b). Thus, even though populations containing altruists do better than populations without them, altruists go extinct because they are disadvantaged as individuals relative to their competitors *within* their groups.

One might respond by saying that this outcome came about because, in my model, there were some selfish individuals in the largely altruistic populations from the start. A fairer test of group selection would be to start populations F–J with *only* altruists in them. Setting up the model this way might delay the inevitable, but it would not change the result. As long as even very occasionally a selfish individual migrates into a population composed of altruists, or alternatively a

genetic mutation occurs which turns an altruistic individual selfish, then selfishness will always end up invading.

4.3.2 The evolutionarily stable strategy

Selfishness always out-competes altruistic behaviour in the above example because altruism is not an evolutionarily stable strategy (ESS). An ESS is a behavioural policy that, once common in a population, cannot be out-competed by any alternative behavioural policy. In our model, selfishness is an ESS, whereas altruism is not. Natural selection will always find the ESS in the end.

Follow the alleles, again

We could have come to the same conclusion about Wynne-Edwards' explanation for reproductive restraint in a simpler way by doing the fitness accounting at the allelic level rather than that of the bird. Assume simply that an allele causing selfishness is in competition with an allele causing altruism in the above scenario. We can calculate the fitness of these two alleles, in the three contexts in which they are each found: in good years, in populations that go extinct in bad years, and in populations that survive in bad years (Table 4.1). The fitness of the two alleles is the same in good years. They are also the same (i.e. zero) in bad years in populations that go extinct. However, in bad years in populations that survive, the selfish allele has twice the fitness of the altruistic one. Altruists are very slightly less likely to be found in populations that do go extinct (because being an altruist increases the number of altruists in your group by one), but this effect is not strong enough to offset the large fitness differential within the group. Thus, summed across the three types of situation, the average fitness of the selfish allele is higher than that of the altruistic allele. This means that it will eventually go to fixation and the altruistic allele will go extinct.

Implications

Let us summarize the implications of the results of this section. They mean that behaviours will not usually evolve if they benefit some larger group at a cost to the individuals performing them (although see section 4.5 for special cases). This is simply because the alleles underlying the behaviours will always have lower overall fitness than competitor alleles which do not take this cost. This is even true if the behaviours that evolve are damaging to the collective interest and drive the whole population extinct in the long term. Natural selection operates on immediate allelic advantage and cannot 'see through' this to the long-term viability or sustainability of a species or population.

Table 4.1 **Fitness of the 'altruistic' and 'selfish' alleles in the bird example. The fitness is identical in two out of three possible situations, and the selfish allele has higher fitness in the third, so the selfish allele has higher fitness overall and hence must spread.**

	In good years	In bad years when population goes extinct	In bad years where population survives
'Altruistic' allele	2	0	1
'Selfish' allele	2	0	2

This conclusion places important constraints on the types of evolutionary explanations we should consider. First, evolution does not produce outcomes that are 'good for the species'. There are many examples in nature of things that are clearly bad for the species yet still evolve. For example, amongst mammals, a major cause of death amongst young individuals is infanticide by members of the same species. This is clearly bad for the species and yet it has still evolved because alleles for committing infanticide can out-compete alleles for not committing infanticide. Wherever you encounter a 'good for the species to have members that do X'-type argument, you need to reframe it in 'good for individuals to do X because . . .' terms, or, even better, in 'the fitness of an allele causing X would be higher than competitor alleles because . . .' terms.

Second, we need to be cautious about arguments that invoke the good of a group or population rather than individuals. For example, if I wanted to explain the existence of celibate priestly castes in many human societies, I might be tempted to speculate that it is good for societies to have a group of individuals within them who have no partisan interests and are devoted to the broader social good. This speculation could well be true, but it is not sufficient to explain how celibacy evolved. Rephrase it in allelic terms. An allele that caused its bearers to become celibate would always be out-competed by competitor alleles that caused their bearers to reproduce, so the fact that celibacy would be good for the group is insufficient to make it evolutionarily stable.

4.4 Kin selection

The previous section has shown that adaptations that are better for groups than they are for individuals are likely to be out-competed and disappear. Thus, our general expectation should be that the organisms that we see today will be designed to behave in such a way as to promote their individual interests rather than anyone else's. A quick glance at animal behaviour provides plenty of evidence that this view is correct; individual animals feed themselves first, run away from predators, and may attack or even kill other group members when it seems to be in their interests to do so.

However, there is a striking group of exceptions to this picture, of which parenthood is the central one. Mothers gestate, suckle, and protect their offspring, and do not seem to receive anything in return. Intuitively, this is to do with the fact that a mother's offspring are genetically related to her. If you will, they represent the future of her genome and thus when she invests in them, she is investing in the fitness of her genotype (albeit, a different copy of her genotype than the one which happens to be in her body). This section formalizes this intuition and considers the extent to which it would be adaptive for individuals to invest in the copies of their genome that are inside bodies other than their own.

The part of evolutionary theory that deals with this issue is called the theory of kin selection and its development is largely credited to the English biologist William Hamilton (Figure 4.7). Kin selection theory does not just deal with parents and offspring. It can also be applied to sibling, nephew, and any other family relationships. The central component of the theory is one we have already met, the coefficient of relatedness (section 3.3.1). Recall that this coefficient represents the size of the expected increment of allelic similarity between two relatives above and beyond the similarity to be found between two randomly selected members of the population.

Figure 4 / William Hamilton (1936 – 2000), father of the theory of kin selection.
© *James King-Holmes/Science Photo Library.*

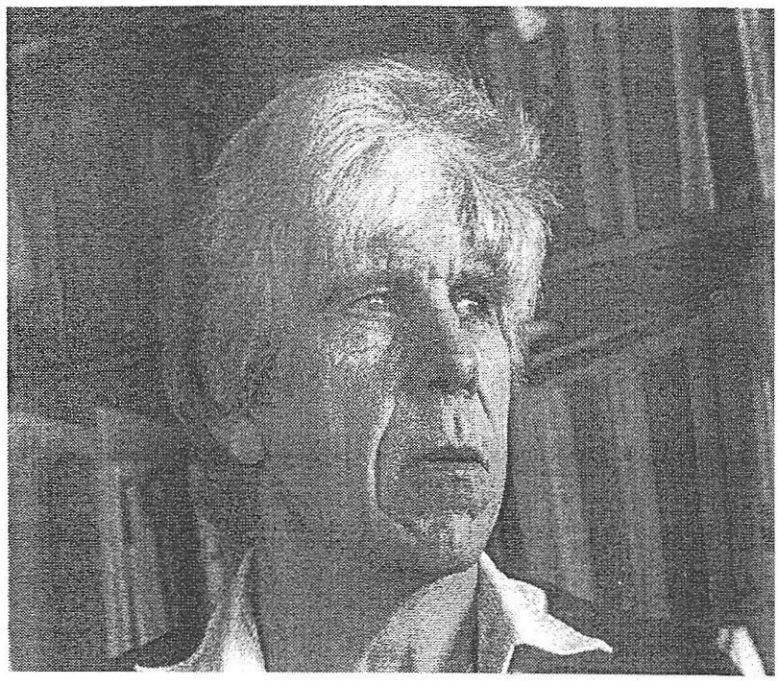

4.4.1 Hamilton's rule and the concept of inclusive fitness

Let us consider a hypothetical example of a bird species, in which a mutant allele a_1 arises. The phenotypic effect of this mutant is to make younger sisters forego their reproduction and instead aid the reproduction of an older sister (e.g. by guarding the nest or bringing food for the sister's chicks). We assume that when this allele is found in an individual with no older sister, that individual reproduces as normal. The competitor of a_1 is a_2, which causes the younger sister to reproduce for herself. Under what conditions would a_1 out-compete a_2?

At first glance, the answer might seem to be 'never', because the effect of a_1 is to reduce its bearer's individual reproductive success. However, assume that, by helping, the younger sisters with a_1 increase the reproductive success of their older sisters. Because they are related, those older sisters are disproportionately likely *also* to be carrying a_1. More precisely, in around half of all cases (because the coefficient of relatedness is ½), the sister will have an allele which is identical by descent and therefore bound also to be a_1. By helping a sister, bearers of a_1 are disproportionately helping other copies of a_1 to reproduce and therefore disproportionately not helping copies of a_2.

However, siblings are not genetically identical. The older sister could be carrying a_2 and thus aid given to a sister will sometimes benefit the competitors of one's own alleles. Thus, there seem to be both benefits and hazards to investing in kin rather than oneself. William Hamilton showed that a simple inequality described the circumstances under which such investment could be adaptive. This inequality is known as Hamilton's rule and it states that a kin-directed behaviour can be favoured by selection whenever:

$$c < rb$$

Box 4.1 Deriving Hamilton's rule

Where does Hamilton's rule come from? It is mathematically complex to derive it formally, but one simple way of showing where it comes from is to calculate the relative fitness of an allele for foregoing reproduction to help an older sibling. As before, allele a_1 causes younger sisters to forego reproduction, at cost c, to help their older sisters, who thereby get benefit b. It causes singletons and older sisters to reproduce normally. This allele is in competition with an alternative allele a_2, which causes younger sisters to reproduce for themselves. We can tabulate the expected fitness of alleles a_1 and a_2 in the three contexts in which they each occur: in the bodies of singletons, in the bodies of older siblings, and in the bodies of younger siblings (Table 4.2). In the table, w represents an average level of fitness for the population. The allele a_2 always has fitness w. The allele a_1 has fitness w when in an offspring with no siblings, $w + \frac{1}{2}b$ when it is in an older sister, since the younger sister may share the allele and thus give aid to the value of b, and $w - c$ when it is in a younger sister, since it causes the younger sister to forego reproduction to the value of c.

The sum of the top row of the table is $3w + \frac{1}{2}b - c$, and the sum of the second row is $3w$. The *difference* in fitness between a_1 and a_2 across all contexts is the difference between the top row and the second, namely:

$$3w + \tfrac{1}{2}b - c - 3w$$

This is equal to:

$$\tfrac{1}{2}b - c$$

When will a_1 have higher fitness than a_2? Exactly when the difference in fitness between them is greater than zero; that is, when:

$$\tfrac{1}{2}b - c > 0$$

or in other words, when:

$$c < \tfrac{1}{2}b$$

Thus, an allele for helping siblings has higher fitness than its competitors when $c < \frac{1}{2}b$, exactly as Hamilton's rule says. Similar reasoning could be followed for other coefficients of relatedness.

Table 4.2 Fitness of an allele, a_1, whose effect is to make younger sisters forego reproduction to the value of c but provide aid to their older sisters to the value of b, compared with a competitor allele a_2 that has no such effect. w is the average level of fitness in the population. The situation has been simplified by assuming a_1 is dominant and rare. a_1 will out-compete a_2 exactly where $c < \frac{1}{2}b$.

Allele	Fitness in offspring with no sibling	Fitness in older sister	Fitness in younger sister
a_1	w	$w + \frac{1}{2}b$	$w - c$
a_2	w	w	w

where c is the reduction in the actor's reproductive success, *b* is the increase in the recipient's reproductive success, and *r* is the coefficient of relatedness. For the sibling case described above, since $r = 1/2$, the behaviour could spread whenever the average increase in the older sister's reproductive success caused by helping was more than twice the reproductive success foregone by the actor to help. For a nephew, where *r* is only 1/4, the benefit to the recipient would have to be at least four extra chicks for every one forgone by the actor, in order for the behaviour to evolve.

Hamilton's rule means that if we want to calculate the reproductive success of an individual animal, we should not restrict ourselves to the animal's number of personal descendants, but also add in any extra reproduction by relatives that results from the individual's behaviour, adjusted by the coefficient of relatedness. Reproductive success so calculated is called **inclusive fitness**.

4.4.2 Applications of kin selection

Hamilton's rule has very wide applicability. Most obviously, it can be used to understand why and to what extent adults invest in their children and grandchildren (as we shall see in Chapter 8). There are also many phenomena in nature similar to the sibling example described above, where individuals under certain conditions invest in the offspring of their siblings or the later offspring of their parents, rather than reproducing for themselves. These behaviours are known as **alloparenting** and although there can be other benefits of alloparenting behaviour, kin selection appears to play a strong role (Griffin & West 2003).

Kin selection within the individual

Kin selection also has some less obvious applications. For example, most of the cells in your body have no chance of ever reproducing. Muscle cells, red blood cells, and cells in the immune system are all doomed to die out when you die and yet they go on working all your life to keep your body functional and intact. On the other hand, sperm and egg cells do nothing but sit around waiting to be used in reproduction. Thus, in one sense, cells in all other tissues forego their own reproduction to allow the gametes a chance to reproduce. How could such reproductive restraint be evolutionarily stable? The answer of course is kin selection. The cells in your body are genetically identical (coefficient of relatedness is 1) and thus investing in reproduction of gametes is, as far as a cell in your brain or blood is concerned, just as good as reproducing for itself.

It follows from this that genetic mutations arising in some cells of the body but not others are potentially disruptive to the body's functioning. This is because once such a mutation has occurred, there are allelic differences between cells (i.e. the coefficient of relatedness is no longer exactly 1) and there will be competition between different cell lineages. Selection within the organism then favours cells that proliferate differentially at the expense of other body tissues. For this reason, cells in which genetic mutations occur, for example during mitosis, are detected and destroyed by the immune system. However, sometimes such cells evade the immune system and then they become cancers—mutation-containing cell populations that favour their own expansion at the expense of the proper functioning and ultimate interests of the whole body. Cancers are related but non-identical organisms to their hosts, with their own fitness interests. Lest this seem a fanciful way of looking at it, there is at least one independent organism that started out as a group of cancerous cells within a larger one. Canine transmissible venereal sarcoma is a venereal disease of dogs. The cells of this parasite are descended from cancerous tissues of a particular dog, tissues that accumulated mutations allowing them to spread by mitosis, first through the body of the original host and later by genital contact with other hosts too.

Figure 4.8 Proportion of times that black-tailed prairie dogs (*Cynomys ludovicianus*) give alarm calls when approached by a predator, as a function of the kin present in the group. *Data from Hoogland (1983). Photo © Adam Mandoki/istock.com.*

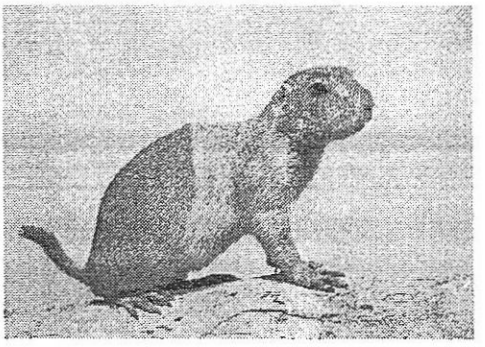

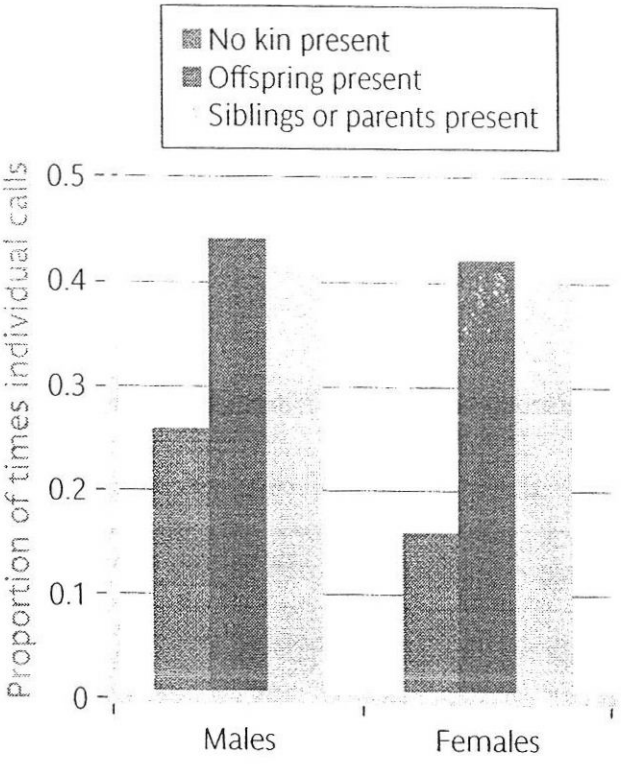

Alarm calling in prairie dogs

Kin selection is the key to understanding many behaviours that seemed to earlier generations of biologists to exist for the good of the group. For example, in the black-tailed prairie dog, *Cynomys ludovicianus*, adults sometimes give alarm calls when they detect a predator. These calls alert other individuals to the presence of the predator, but will tend to draw attention to the individual giving the call. A classic group selection account would stress how groups containing callers might survive better than groups with no callers, but this account cannot be correct, as we saw in section 4.3. Instead, calling adults might disproportionately benefit their kin. If this was the case, then the behaviour could evolve by kin selection as long as Hamilton's rule is satisfied.

Hoogland (1983) approached prairie dog colonies with a model predator (a stuffed badger) and noted how often individuals gave alarm calls, having first established the kinship relationships between animals in the colony. The results show that calling is much more common when the caller has kin in the colony than when it does not (Figure 4.8). These kin can be offspring, but calling is also common when the kin in the colony are siblings or parents.

4.4.3 Conditions on Hamilton's rule

There are several constraints on the applicability of Hamilton's rule that should be mentioned. The first is that a behaviour will not evolve, even if $c < rb$, if c is too large in absolute terms. For example, an allele that made all its bearers sterile could not spread, even if they thereby helped

their siblings have 100 extra offspring. This is because all the siblings they helped to reproduce would necessarily *not* possess the allele, otherwise they would be sterile too. Thus, Hamilton's rule is only applicable to behaviours with a weaker selective disadvantage, such as foregoing *some* reproduction under *some* circumstances. The second point is that the coefficient of relatedness alone is not sufficient to predict which behaviours will evolve. For example, parents nearing the end of their lifespan might well do more to protect their offspring than those offspring would do to protect the parents. The coefficient of relatedness is the same in the two cases, but the parents can have no more offspring, whereas the offspring have all of their reproduction ahead of them. We can incorporate such asymmetries into Hamilton's rule by assuming that b and c are not fixed for particular behaviours, but variable depending on the future prospects of the recipient and actor, respectively.

Finally, and unlike the case of cells within the same body, the coefficient of relatedness of one animal to another is always less than the coefficient of relatedness of that animal to itself. Thus, although the theory of kin selection predicts that there will be widespread investment in kin found in nature, it is perfectly compatible with there being behaviour directed *against* kin too. For example, in the very same prairie dogs that give calls to help kin, a major source of death in the young is being killed by the mother's sister when she has pups of her own (Hoogland 1985). Such females are reducing their inclusive fitness by killing their nieces and nephews. However, they are more closely related to their own offspring than they are to their sister's and so, providing that the benefit to their own pups is great enough, the behaviour can be adaptive. Behaviours that harm kin are particularly likely to evolve where there is local competition between relatives for finite resources. A clear example is found in certain bird species, such as kittiwakes, where nestlings may be killed by their siblings. Parental provision to the nest is limited and the benefit of receiving more of it is clearly greater under certain conditions than the inclusive fitness benefits through the lost siblings.

4.5 Advanced topics: evolutionary transitions, levels of selection, and intra-genomic conflict

We end this chapter by reviewing some more advanced topics in evolutionary theory, which build upon the principles set out so far. You can progress to Chapter 5 without covering these topics, but they are areas of great current research interest and since they go to the very heart of how evolutionary competition works, understanding them is useful for deepening your understanding of evolutionary theory more generally.

The argument so far in this chapter has been that behaviours that benefit groups at the expense of individuals will not be evolutionarily stable against competitor behaviours that benefit just individuals. This is usually true. However, there are examples in nature of collectives of unrelated entities who all work for the common good. In such groups, many of the adaptations that arise appear designed for the good of the collective, not the component elements.

The most compelling example of such a collective is the animal. Calling an animal a collective seems paradoxical, but the body of an animal such as yourself is created by around 25,000 genes working together and very few of these genes are closely related by kinship. Indeed, your

mitochondria are descended from free-living bacteria that became incorporated within the cells of a very distant ancestor. They even have their own genome. Thus, your genome can be seen as a group of many unrelated elements, including some that originate from a different species, and yet they work together for the common good of making an integrated, functional body.

This leads us to re-examine the question of the level of selection. It seems plausible that genes, by working together on a common body, can do better at replicating than any of them could by trying to replicate alone and that this is the key to the evolution of complex organisms. However, it is not immediately obvious why this works for a group of genes uniting for the common good of the body and fails to work, for example, for a group of birds regulating their reproduction for the common good of the population. Would not the same issues of evolutionary stability that we saw in section 4.3 arise?

4.5.1 The Price equation

A useful tool for thinking about this problem (and all the others considered in this chapter) lies in a framework called the Price equation, developed by theorist George Price in the 1970s (for a non-mathematical introduction, see Okasha 2006). The Price equation is complex, but its message can be simplified for our purposes as follows. The evolutionary change we should expect in a characteristic depends on that characteristic's covariance with the fitness of the alleles coding for it. Covariance is basically like correlation; it is positive when increasing the characteristic increases fitness, negative when decreasing the characteristic decreases fitness, and zero when there is no association between the characteristic and fitness. Thus, characteristics that covary positively with fitness will increase over evolutionary time, those that covary negatively will decrease, and those with zero covariance with fitness will not change. Behaviours that benefit a collective are no exception to the general picture. They will evolve as long as the overall covariance between the fitness of the allele and the functioning of the collective is positive.

This condition explains why group selection did not work in the reproductive restraint example of section 4.3. There was a positive covariance between the amount of altruism in a group and the average fitness of that group's members, due to differential group survival. However, this was very weak, since selfish individuals that happened to find themselves in groups with altruists benefited as much as the altruists did. There was also a strong negative covariance between altruism and fitness *within* the group because altruistic individuals reproduced less than their selfish group-mates. This negative covariance more than offset the smaller positive one, leading to an overall negative relationship between altruism and fitness, and hence altruism's demise.

4.5.2 Suppression of within-group competition

Now consider genes cooperating to make a body. As our bodies are constituted, the only way for genes to replicate themselves is for the whole genome to get replicated (i.e. for the phenotype to survive and reproduce). If it does so, they all benefit to the same extent. If it fails to do so, they all die out. Thus, the covariance between phenotypic success and allelic fitness is positive and strong. This allows the complex genome to emerge as a functional collective; all of the individual genes have the same interest in making it work. In general, wherever there is a functional collective in nature, there are mechanisms in place that abolish any differences in reproductive success *within* the collective, so that any element's fitness is determined only by the functioning of the whole (Frank 2003).

In complex genomes, such mechanisms are in place. During meiosis, genes are bound together on chromosomes and either the whole chromosome goes forward to the gamete or none of it does. This abolishes any competition between neighbouring genes to, for example, replicate themselves faster than their neighbours. They are either all going forward together or all failing together. The way meiosis works makes it a lottery which of the two copies of a particular chromosome ends up in a particular gamete. It is not, for example, dependent on the chromosome's size. A fair lottery abolishes competition because no characteristic of the alleles on the chromosome makes any difference to the result. Because of mechanisms such as these, the fate of any one of your genes is highly correlated with the fate of all of the others and so the Price equation predicts that they will cooperate for the common good.

Earlier in the chapter I mentioned that, although the best way to keep track of fitness in evolutionary models is at the level of the allele, we can use the reproductive success of the individual as an approximation for many purposes. The Price equation allows us to understand why this works: because the covariance between the reproductive success of an organism and the fitness of all its alleles is substantial. We can also refine our conclusion concerning the level of selection. Rather than adaptations for the collective good never being evolutionarily stable, they will persist exactly where mechanisms have evolved that ensure high overall covariance between allelic fitness, the functioning of the collective, by suppressing competition between elements within the collective.

To visualize this more concretely, imagine two different reality television shows. In both shows, a group of people are taken to a remote location. They are set various challenges. In one show, the individual who completes more challenges than all the others gets £1 million. The Price equation predicts that people in this scenario will try to do each other down because there is a negative covariance between a participant's winnings and the success of the other group members in the challenges. In the second show, the participants each get £1 million if all of them complete all of the challenges. Here, you would expect the participants to help each other because there is a positive covariance between one individual's winnings and the performance of all the others. The former type of game is the one that free-living organisms are engaged in, which is why they look after themselves, not other group members. The latter type of game is the type that the genes in your genome are playing, which is why they work together to make a functioning body.

4.5.3 Evolutionary transitions

There are a number of points in the history of life where several previously free-living elements come together and start to operate as a collective. Each of these points heralds a major evolutionary transition, where a new form of biological organization arises (Figure 4.9; Maynard Smith & Szathmáry 1995). One transition is from independent molecular replicators to groups of genes. Another is the emergence of more complex genomes with chromosomal organization. Still another is the eukaryotic cell, which, you will recall, is formed by more than one species and then the multicellular organism from a single-celled ancestor. A further transition that a few species have undergone is the transition to eusociality. Eusociality is the situation where a whole colony of individuals work together to further the reproduction of one or just a few of their number. Eusociality has evolved at least 15 times in nature and is best known in, although not restricted to, the ants, bees, and wasps. Eusocial colonies have impressive functional organization, with different castes fulfilling different roles and a number of features existing that look well designed for the colony, not for the individuals within it. In fact, the eusocial colony of a

Figure 4.9 The history of life shows a succession of transitions at which elements form collectives which then begin to evolve apparent adaptations of their own. The lower-level elements continue to exist independently alongside the new entities. *Photo bottom left: © merrymoonmary/istock.com; photo bottom middle: © Graham Cripps/NHMPL.*

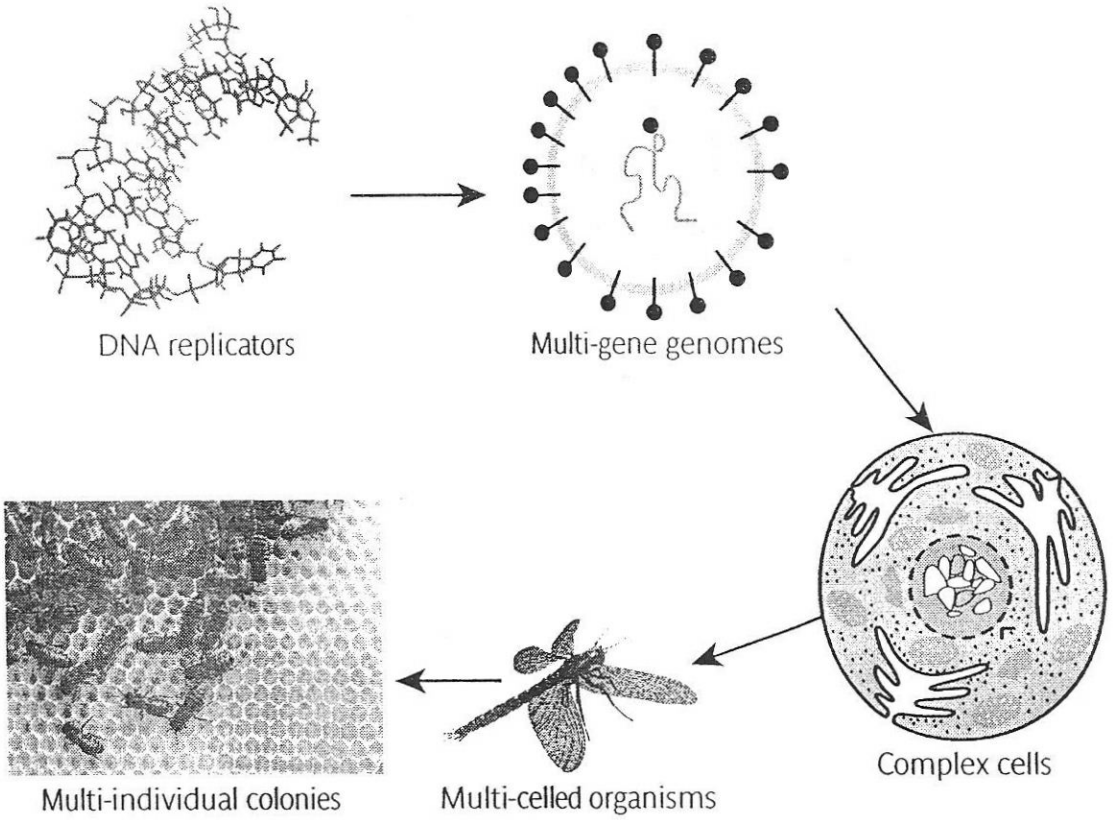

bee or wasp can be viewed as a kind of organism rather than a group of individuals, just as you can be viewed as an organism rather than a coalition of genes.

The general pattern with the major transitions appears to be that the threshold from loose aggregation of independent entities to functional collective is very hard to cross because collective cooperation is usually disrupted by competition within the group. On those rare occasions that the threshold is crossed, there is a point of no return—it is very unusual for part of a multicelled organism to return to separate existence, as canine transmissible venereal sarcoma has done—and the resulting life forms can quickly radiate and prosper, forming a whole new branch of the tree of life.

4.5.4 Intra-genomic conflict

Even when the point of transition to collective functioning is crossed, there can still be some internal competition simmering away. Within the complex genome, for example, it has become clear that the mechanisms suppressing competition between genes are not always perfect and genes in the same individual can differ in fitness. Such competition and the effects it produces are known as intra-genomic conflict.

Transposable elements

Intra-genomic conflict arises whenever genes can favour their own interests above that of the whole. We have seen one example from transposable elements (Chapter 2). These are genetic sequences that have the ability to make extra copies of themselves at meiosis. Thus, over evolutionary time, they can proliferate within a complex genome—recall that humans have over 1 million copies of the *Alu* element. This serves *Alu*'s own fitness interests, not the interests of the organism. Indeed, the rest of the genome probably incurs a small cost of the extra *Alu* material and is thus under selection to shut *Alu* activity down. The evolution of the genome will thus be a dynamic of cat and mouse between selfish elements like *Alu* and suppressive adaptations elsewhere in the genome.

Segregation distorters

Genetic variants that distort fair segregation provide another case. Under normal meiosis in a diploid organism, the chance of a particular allele going forward to the gamete is 50%. In mice, there is a genetic variant called the t haplotype. When one copy of the t allele is present in males, 90%, rather than 50%, of the viable sperm produced carry t. The t variant seems to achieve this by disabling most non-t sperm. It thus gives itself an advantage relative to all its competitors. The t variant is found in mice all over the world and has persisted for many thousands of generations. It has only remained at the low frequencies that it has—around 5% in many populations—because the homozygote is lethal in males and has negative effects in females too. It clearly persists because of its own ability to distort segregation rather than any beneficial effect at the level of the organism.

Cytoplasmic male sterility

Another example is the tug-of-war between the mitochondrial portion of the genome and the rest. Mitochondrial genomes are inherited down the female line only. This occurs because the larger, female gamete provides the cellular environment of the zygote, whereas the male gamete provides basically only nuclear DNA. This causes a conflict of interest between the mitochondrial genome and that of the cell nucleus. The former will only succeed in female offspring, whereas the latter succeed in offspring of either sex. Thus, any mutations within the mitochondrial genome that increase the proportion of females produced will have a selective advantage over their competitors.

Plants are usually hermaphrodite, producing both pollen and ovules. In many types of plant, there are mitochondrial genetic variants that shut down pollen production and make individuals only female. This phenotypic effect is called cytoplasmic male sterility (CMS). Genetic variants causing CMS can become extremely common and when they occur, there is selection on variants in the nuclear genome to counteract their effects and restore pollen production. Over evolutionary time, there is a dynamic arms race between mitochondrial CMS mutants and counter-mutants in the nuclear genome. Many plant species are found with a mixture of female-only individuals and hermaphrodite individuals as a result. Examples such as these remind us that organisms are not quite perfectly unified; their origins as groups of elements with distinct interests surface every now and then.

Summary

1. Within all biological populations, there are differentials in reproductive success and thus competition to reproduce.
2. Some phenotypes lead to greater success than others in reproductive competition in particular environments and it is these differences that drive the process of adaptation.
3. Although natural selection produces changes in phenotypes, it does this by changing the frequencies of the underlying genotypes. The best way to model evolutionary outcomes is to track the fitness of competing alleles, rather than competing individuals or competing populations.
4. Any behaviour or characteristic that, once common in the population, cannot be displaced by any competitor behaviour or characteristic is said to be an evolutionarily stable strategy (ESS).
5. In free-living organisms, innovations that enhance the interests of the group or species at the expense of the individual are unlikely to be evolutionarily stable. Adaptations in nature are generally well designed for the individuals' genes, rather than for groups or species.
6. Behaviours that benefit relatives at the expense of the actor can evolve, if Hamilton's rule is satisfied, because of allelic relatedness between kin.
7. Selection can favour adaptations for the good of higher-level collectives, as long as mechanisms have evolved that abolish differentials in reproductive success amongst their constituent elements, so that there is a positive overall covariance between the collective's functioning and the reproductive success of the elements within it. Complex organisms, which are coalitions of different genes, represent prime examples of such collectives.
8. The covariance between fitness of different genes within an organism is not always total, and this leads to intra-genomic conflict and patterns that benefit some genes at the expense of the whole organism.

Questions to consider

1. In the cat model in section 4.1, adding mortality slows the exponential growth of the population (see Figure 4.3). You will note, however, that the population is still going to explode exponentially eventually. Adding a certain probability of mortality has merely retarded the increase by a few years. Such explosions do not tend to occur in real biological populations. Why do you think this is? How would you model the pattern more realistically?
2. Section 4.3 argued that Wynne-Edwards' group selection explanation for birds having smaller clutches in bad years must be wrong. However, birds do actually do this. How might it be explained without recourse to group selection ideas?
3. Identical (MZ) twins have identical genomes (coefficient of relatedness of 1). Thus, the theory of kin selection predicts that they will treat the offspring of the other twin exactly as they treat

their own. There is no reason for them to favour their own interests over those of the other twin. In fact, their behaviour is more like that of ordinary but close siblings. Why might the behaviour not follow the theory in this instance?

4. Shakespeare's plays frequently centre around brothers banishing or killing their brothers. For example, in *As You Like It* and *The Tempest*, one brother banishes the other and takes over the dukedom, and in *Hamlet* and *Richard III*, brothers kill brothers to become king. Are these aspects of Shakespeare's stories biologically implausible?

5. Given the section on intra-genomic conflict, how do the interests of the Y chromosome differ from those of the X? What kinds of alleles on the Y might be favoured, and what response from the X would be favoured?

Taking it further

The models of evolution presented in this chapter are all simplified for expository purposes. For more mathematically rigorous approaches, the reader is referred to McElreath & Boyd (2007). The classic—and brilliantly written—statement of the need to consider the relative fitness of alternative alleles to understand evolution is Dawkins' *The Selfish Gene* (Dawkins 2006b). The debate about the correct level at which to conceive that selection is acting rumbles on—very different views are given, for example by West *et al.* (2006) and Wilson & Wilson (2007). The problem stems from the fact that frameworks using the language of group selection are formally equivalent to frameworks using the language of (inclusive) individual fitness. It thus becomes somewhat semantic which language is preferred.

The major transitions in evolution are reviewed by Maynard Smith & Szathmáry (1995), and the special role of repression of within-collective competition is modelled by Frank (2003). For a recent review on eusociality, see Wilson & Hölldobler (2005). Okasha (2006) shows how the Price equation and refinements of it can be used as an overall language for thinking about levels of selection and the evolutionary transitions. On intra-genomic conflict and selfish genetic elements, Burt & Trivers (2006) is the authoritative but vast source; shorter introductions are provided by Hurst *et al.* (1996) and Hatcher (2000).