

Chapter 3

Natural Selection, Kin Selection and Group Selection

ALAN GRAFEN

3.1 INTRODUCTION

'Animals maximize their inclusive fitness.' 'Animals do not sacrifice their own fitness for the good of their group.' Such statements can be read in many of the chapters of this book, as well as in many recent interesting papers in behavioural ecology and animal behaviour. But they are not obviously true—there are counterexamples in theory to both of them—and my main purposes in this chapter are to trace the logic underlying them, to identify the justification that has been found for them, and to display the connections between these usable principles of behavioural ecology and the more fundamental principle of natural selection.

This chapter will mainly justify what many people already believe: but this is an important exercise. The central concept of inclusive fitness is routinely misdefined and occasionally misused. Some confusion still surrounds the problem of group selection. Only by concentrating on the logical grounds of our orthodoxy can we clarify and defend it.

The rest of this chapter is divided into three parts. The first describes how population genetics is fundamental to behavioural ecology, but also how we may be able to avoid dealing with its complications in everyday practice. The second traces the logic underlying inclusive fitness. The last discusses what natural selection tells us about the evolution of animals in groups.

3.2 POPULATION GENETICS UNDERLIES BEHAVIOURAL ECOLOGY

The starting point for much behavioural ecology is that animals are maximizers of one sort or another—efficient predators or foragers, or elusive prey. The only ground for believing this is that natural selection made them so. If not now then at some time in the past (Dawkins 1982a, pp. 20–24), there existed heritable variation in hunting and foraging techniques and in ploys to escape predators. Changes in allele frequencies have made animals good at what they do (see also section 4.1).

The behavioural ecologist, though, does not usually know the genetics underlying the character he studies. While he would be interested to know this genetic system, it is not of primary importance to him. His main aim is to uncover the selective forces that shape the character. The behavioural ecologist has to hope in his ignorance that his method will work almost regardless of which particular genetic system underlies the character (Lloyd 1977). This hope raises two questions. First, is it justified? Secondly, is the assumption so powerful and plausible that a whole research strategy should be based on it?

3.2.1 The phenotypic gambit

Let us start with a brief caricature, with examples, of an important method in behavioural ecology. It has two elements.

(i) *A strategy set.* This is a list or set of (perhaps all) possible states of the character of interest. Here are three examples of strategy sets. McGregor, Krebs and Perrins (1981) studied the song of male great tits, and in particular their repertory size. The strategy set they used was simply every different repertory size they observed: integers from one to five. Brockmann, Grafen and Dawkins (1979) studied the nesting of great golden digger wasps. These wasps sometimes acquire their nest by digging one, and sometimes by entering an already existing one. Brockmann and co-workers were interested in the relative frequency of these two ways of acquiring a nest, and so the strategy set was simply all possible proportions of digging rather than entering: numbers from zero to one. In the hawk-dove game devised by Maynard Smith and Price (1973), the strategy set consists of the two strategies, hawk and dove.

(ii) *A rule for determining the success of a strategy.* The success of a strategy is the number of offspring left by an animal adopting it, or alternatively its inclusive fitness (see section 3.3.1). The rule for determining success may involve the frequency with which strategies are adopted in the population. We may observe the operation of the rule, as McGregor, Krebs and Perrins did. They counted how many offspring every male fathered in his lifetime, and averaged within all males sharing the same repertory size. If it is necessary to know how the successes of strategies change when their frequencies change, then we may model the rule. Brockmann, Grafen and Dawkins did this, and used data to estimate parameters in the rule. When the purpose is to investigate theoretically the consequences of a particular form of frequency dependence, a rule exhibiting this form is simply assumed: in the hawk-dove game the rule is represented in the payoff matrix.

The phenotypic gambit is to examine the evolutionary basis of a character as if the very simplest genetic system controlled it: as if there were a haploid locus at which each distinct strategy was represented by a distinct allele, as if the payoff rule gave the number of

offspring for each allele, and as if enough mutation occurred to allow each strategy the chance to invade.

The gambit implies that all strategies that occur in the population are equally successful, and that they are at least as successful as any non-occurring strategy would be if it arose in small numbers. The application of the gambit to a given strategy set and payoff rule is a powerful test of the joint hypothesis that the strategy set and payoff rule have been correctly identified, and that the gambit is true.

In their first model, Brockmann, Grafen and Dawkins rejected this joint hypothesis when two existing strategies turned out not to be equally successful. They adopted a new strategy set in their second model. See Dawkins (1980, 1982a) for a full discussion of what conditions an act must satisfy to be a 'strategy'.

The joint hypothesis might be false because the genetic system underlying the character does not produce the same phenotypic effects as the very simplest genetic system, the one assumed in the gambit. The mere fact that the prediction of equal success is rejected does not reveal which element in the joint hypothesis is false. The research strategy implied by the phenotypic gambit is to treat such rejection as evidence that the payoff function or strategy set is wrong, and not that the gambit is wrong. Maynard Smith (1978a) discusses this more fully.

3.2.2 Is the gambit true?

Taken literally, the gambit is usually false: few species studied by behavioural ecologists are haploid. But will the genetic system that does underlie the character produce the same phenotypic effects as the genetic system the gambit assumes?

Two points are important here. First, an example is known in which the gambit would be extremely misleading. In some human populations affected by malaria, there are three distinct phenotypes corresponding to the three possible genotypes at a diploid locus with two alleles (Allison 1954). One type almost invariably dies from sickle-cell anaemia before reproducing. The other two types differ in their resistance to malaria. The coexistence of these three phenotypes with markedly different fitnesses would be very puzzling to a behavioural ecologist applying the phenotypic gambit. The mechanics of Mendelian segregation prevent the whole population from sharing the optimal phenotype, because it is produced by the heterozygous genotype. Here, as undoubtedly elsewhere, it is essential to know the underlying genetics in order to understand the distribution of phenotypes observed in the population.

The second point is that such cases are probably rare. Only certain features of genetic systems, such as overdominance in the sickle-cell case, can sustain dramatic differences in fitness, and these features are not known to be common. Maynard Smith (1982a) has analysed how

well different genetic systems support the simplification represented by the gambit, and he concludes that by and large they do so very well. The sorts of character studied by behavioural ecologists are likely to be controlled by many loci, and this reduces the scope for the maintenance of large fitness differences.

Genetic systems are themselves subject to evolution. In its simplest form, this is the creation of a new allele by mutation, but more substantial changes could occur. In the sickle-cell case, a (functional) gene duplication of the locus would allow one locus to fixate for each allele. Every individual in the population could then have the 'intermediate' genotype that confers malarial protection without sickle-cell anaemia. The existence of fitness differences between genotypes creates selection for evolution of the genetic system itself.

The behavioural ecologist hopes that most genetic systems support the gambit, and that those that do not are rare or transient. If the discrepancies produced by different genetic systems are smaller than the accuracy of data, then field workers can safely ignore them. We know that this might not be so, and we should be anxious to find out whether this hope is justified. The dependence of behavioural ecology on population genetics is such that the soundness of our methods depends on arguments concerning population genetics, but our method is designed to avoid doing genetics.

We have seen that the gambit cannot certainly be made with safety. It is a leap of faith. But should we then refuse to use it in our research? To answer this, suppose that we did refuse. What would behavioural ecology be like? It would be very different. Detailed studies in which the precise nature of a character is examined as an adaptation would have to be accompanied by a study in which the genetic mechanism underlying the character was uncovered so precisely that an explicit genetic model could be constructed. The motto would be: no decimal points without genetics. The range of characters that could be studied would be drastically reduced. Genetically simple and well studied characters are rarely of evolutionary interest. They are usually straightforwardly disadvantageous mutants maintained by judicious artificial selection in strains which have spent tens of generations in the laboratory. If we had to work out the genetics of every character chosen for its evolutionary interest, the size of the study would become very large. In some cases it would be impossible to complete the study within the lifetime of a scientist.

Another serious point is that if the gambit is generally true, then the genetics discovered would be almost an irrelevant complication in understanding the evolutionary significance of the character. The gambit makes truly phenotypic explanations possible, and the effort expended in discovering the genetics would be wasted. Better to allocate that effort to studying in an evolutionary way characters of evolutionary interest, and in a genetic way characters of genetic interest.

These are the reasons why the gambit is so attractive—they

should not be mistaken for reasons why the gambit is true. Nevertheless, these advantages seem to me to justify continuing to employ the gambit, always provided we remember that we may be wrong. We should also recognize the urgency of the need to provide a proper justification for employing this convenient simplifying assumption.

3.3 INCLUSIVE FITNESS AND HAMILTON'S RULE

Textbooks on behavioural ecology or animal behaviour usually have a section on kin selection or inclusive fitness in which the reader is advised that what animals really maximize is inclusive fitness. They then either fail to define inclusive fitness, or define it wrongly (Grafen 1982); this section is intended to set out what inclusive fitness is. Even in their eulogies the textbooks are not usually right, so this section will also explain why inclusive fitness and Hamilton's rule are extremely useful additions to our theoretical armoury, although by no means a replacement for 'number of offspring' as a measure of reproductive success. It also deals briefly with what is currently known about the scope of their applicability.

3.3.1 What is inclusive fitness exactly?

Inclusive fitness (Hamilton 1964) is a device that simplifies the calculation of conditions for the spread of certain alleles. These alleles have an effect, through their bearer's phenotype, on how many offspring other animals in the population produce. We can see the simplification by comparing the analyses of a very simple model of sib altruism by standard population genetics and by inclusive fitness. Maynard Smith (1982b) carries out a similar exercise.

Suppose males in a species disperse little, so that every breeding male has exactly one brother of the same age breeding nearby. A single locus controls how a male behaves towards his sib, and the population is at fixation for an allele a at that locus. We consider an allele A that alters the behaviour of its (homozygous and heterozygous) bearers so that each bearer has c fewer offspring, and the bearer's sib has b more offspring.

Will A spread when rare? When A is rare, the homozygote AA is extremely rare and can be neglected. The number of offspring a mated pair produces depends only on the male's genotype and that of the male's brother. How many offspring will the pair have on average if the male is aa ? If his brother is also aa , then the pair produces the standard one offspring. If his brother is Aa then on average the pair produces $(1 + b)$ offspring. The overall average for an aa male therefore depends on the chance that his brother is Aa . Let the overall proportion of Aa be p . Then an aa male has an Aa brother with chance $p/2$. The average number of offspring of an aa male is then $(1 + bp/2)$.

The chance that an *Aa* male has an *Aa* brother is $(1 + p)/2$, and the *Aa* male loses c offspring through the effect of the *A* allele, so the average number of offspring of an *Aa* male is $1 - c + (1 + p)b/2$.

I digress to explain how the chances $p/2$ and $(1 + p)/2$ are calculated. In the absence of any information, an animal would calculate the chance that his sib was *Aa* to be p . If the animal is itself *aa*, then it knows that half of the available four parental alleles are not *A*; and so the chance of his sib containing *A* is halved, i.e. it is $p/2$. If the animal itself is *Aa* then a more complex calculation is required. If the chance that an animal is *Aa* is p , then the chance that any allele is *A* must be $p/2$. One of the parents of an *Aa* individual is *A*? and the other is *a*?, where ? means '*A* with chance $p/2$ and *a* with chance $1 - p/2$ '. The sib therefore has a chance $1/2 + p/4$ of receiving *A* from the *A*? parent, and a chance $p/4$ of receiving *A* from the *a*? parent. The total chance of receiving *A* is therefore $(1 + p)/2$. Charnov (1977) uses this method of calculation. It is approximate because it depends on p being small, and it assumes Hardy–Weinberg equilibrium. Now let us return to the calculation of the condition for *A* to spread.

The fraction of *Aa* males in the next generation goes up if *Aa* males have more offspring than *aa* males. What must be true of our variables p , c and b for this to be so? Well,

$$1 - c + b(1 + p)/2 > 1 + bp/2$$

reduces to:

$$b/2 - c > 0.$$

We have just derived the condition for *A* to spread by calculating simply the number of offspring produced by *Aa* and *aa* males. How is the advantage of altruism shown using this approach? Through the extra probability that the brother of an *Aa* male is *Aa* and so increases the male's own number of offspring.

Inclusive fitness arises from a different accounting procedure (Abugov & Michod 1981), in which instead of counting the effect of everybody's actions on one individual's number of offspring, we calculate the effect of one individual's actions on everybody's numbers of offspring. The count is weighted by the relatedness (for a precise definition see section 3.3.4 below). Inclusive fitness was invented and defined (mathematically) by Hamilton (1964). His paper is devoted to proving that the alternative accounting procedure that underlies inclusive fitness gives the same answer as the standard and logically prior procedure, an example of which we have just worked through. Hamilton described inclusive fitness as:

'the animal's production of adult offspring . . . stripped of all components . . . due to the individual's social environment, leaving the fitness he would express if not exposed to any of the harms or benefits of that environment, . . . and augmented by certain fractions of the quantities of the harm and benefit the indi-

Table 3.1. Illustration of the two different accounting procedures implied by two measures of reproductive success: number of offspring and inclusive fitness. The 'advantage to *Aa*' is the same in both systems. The genotype that has more offspring will also have higher inclusive fitness. Number of offspring counts the effects of everybody's acts on an individual; inclusive fitness counts the (weighted) effects of the acts of one individual on everybody.

Measure of RS Genotype	Number of offspring		Inclusive fitness	
	<i>aa</i>	<i>Aa</i>	<i>aa</i>	<i>Aa</i>
Basic nonsocial fitness	1	1	1	1
Cost of act		<i>c</i>		<i>c</i>
Benefit of act	$pb/2$	$(1 + p)b/2$		$b/2$
Total	$1 + pb/2$	$1 - c + (1 + p)b/2$	1	$1 - c + b/2$
Advantage to <i>Aa</i>		$b/2 - c$		$b/2 - c$

vidual himself causes to the fitnesses of his neighbours. The fractions in question are simply the coefficients of relationship . . . ' (Hamilton 1964).

Applying this to our example, the *aa* males have an inclusive fitness of one, because the extra *b* they sometimes receive (i.e. when their brother is *Aa*) is disregarded as a 'help from the social environment'. The *Aa* males have $(1 + b/2 - c)$ because their relatedness to their brother is $1/2$, which we use to devalue *b*. The condition for *Aa* males to have a higher inclusive fitness than *aa* males is then that

$$1 + b/2 - c > 1, \quad \text{or} \quad b/2 - c > 0,$$

which of course is the same answer as before. We have just demonstrated a (simple and) special case of Hamilton's 1964 result. Table 3.1 illustrates the two accounting procedures. The two methods give the same answer by different means.

3.3.2 How not to measure inclusive fitness

Them things that you're liable
To read in the Bible
They ain't necessarily so.

Porgy and Bess

In the example of the last section we saw what inclusive fitness is. Many textbooks give one of two erroneous definitions, and studies have calculated inclusive fitness from data using one of these definitions. It is instructive to examine these errors.

Erroneous Definition 1 (from Barash 1980, p. 212): 'the sum of individual fitness (reproductive success) and the reproductive success of an individual's relatives, with each devalued in proportion as it is more distantly related.'

Erroneous Definition 2 (from Wilson 1975, p. 586): 'The sum

of an animal's own fitness plus all its influence on fitness in its relatives . . .'. I assume it is intended to weight the influences by relatedness.

The original definition of inclusive fitness, as given in the quotation from Hamilton (1964) in section 3.3.1 above, has a component from 'self' and a component from others. ED1 counts all relatives' offspring, whereas inclusive fitness counts only those the relatives had because of the actions of 'self'. ED1 and ED2 count all the offspring of 'self', whereas inclusive fitness does not include those offspring gained through the actions of others (the 'harms or benefits of that environment' in the above quotation from Hamilton).

The reason such erroneous definitions persist is that in most cases they are not applied. In a general discussion the definition itself is never called on with any precision and so the error is in a sense silent, even unimportant. But as soon as data are used to calculate inclusive fitness, the precise definition obviously does matter.

Measures of reproductive success must have one essential property. If bearers of one allele have a higher (lower) reproductive success than non-bearers, then the allele must increase (decrease) in frequency. ED1 and ED2 lack this property. Number of offspring and inclusive fitness do have this property. See Grafen (1982) for a fuller discussion.

In the face of the obvious difficulties of calculating the differences in number of offspring that helping causes, an alternative to inclusive fitness is to use number of offspring as a measure of reproductive success. If these data are available, it will be much simpler to calculate.

When it is desirable to use the inclusive fitness approach in analysing field data, it is better to aim at using Hamilton's rule than to calculate inclusive fitness itself. The next section discusses how to do that. The inclusive fitness approach allows us to separate the success of an allele into components of 'own offspring' and 'relatives' offspring'. Also, it may sometimes be the only approach we can use with certain data. See the discussion of Noonan's example in the next section.

3.3.3 Hamilton's rule and how to use it

Hamilton's rule is that animals are selected to perform actions for which $rb - c > 0$, where r stands for relatedness. Inclusive fitness in Hamilton's 1964 paper was just a tool used in the construction of the rule, and the only reason we have dealt with it at length is that it is surrounded by so much confusion. Hamilton's rule is more important and more illuminating than inclusive fitness; it is also easier to apply data to it because the form of the rule encourages us to use the correct logic of differences. The rule has been derived recently by Charlesworth (1980), using a simple population genetics approach, and by Hamilton (1975) and Seger (1981) using the selection mathematics of Price (1970, 1972).

The first application of Hamilton's rule to data complete with decimal points was by Brown (1975). We will come later to his example of helping at the nest in birds.

The very first step in applying Hamilton's rule is to choose the decision we are interested in—being as explicit as possible about the alternative course of action. To calculate b and c we must think through all the consequences on lifetime number of offspring that follow from doing one thing rather than another. The simple difference in number of offspring will also include the extra b 's contributed by those relatives, and therefore does not give a proper estimate of c .

A difference of c in the animal's lifetime number of offspring results from choosing to do Y rather than X . Any consequences that would follow from doing Y not X should be taken into account—decreased longevity, retribution and so on. It may seem at first sight that a simple way to estimate this from data is to take the difference in lifetime number of offspring between animals that do X and animals that do Y . However, this seemingly reasonable procedure may give the wrong answer. The reason is that animals that do X will have relatives who tend to do X , and animals that do Y will have relatives that tend to do Y . The principle is that an animal that helps n times, and is helped m times, should have $mb - nc$ offspring as a result.

The same caution applies to measuring b .

The value of r has been assessed in a number of ways. Bertram (1976) modelled the structure of his lion prides to arrive at relatednesses; Brown (1975) used simple ancestry; and Metcalf and Whitt (1977b) used electrophoresis. See section 3.3.4 below for a further discussion of r .

Finally, before proceeding to examples, Hamilton's rule in the form $rb - c > 0$ has definite advantages over the more popular form $b/c > 1/r$ for the purposes of statistical testing. For one thing the first form is correct whatever the signs of r , b and c . Also, if r is known from *a priori* grounds, then the mean and variance of the difference $rb - c$ are calculable simply from the means and variances of b and c . The ratio b/c on the other hand has mean and variance that depend in a more complicated way on the distributions of b and c .

Example 1: Brown's analysis of helping at the nest in the superb blue wren

Juveniles in many bird species sometimes stay behind at their parents' nest and help rear their siblings instead of leaving to try to raise offspring of their own. Using data on the superb blue wren (*Malurus cyaneus*) from Rowley (1955), Brown (1975) knew how many young a nest produced in a year according to whether or not there was a helper present that year. He estimated the benefit to the parents as the difference that the helper made, and estimated the cost to the

juvenile as the average number of offspring produced by an unhelped pair.

We now examine the assumptions made in assessing b and c in this way. It is possible that parents survive better if they are helped by their young, and so produce more offspring themselves in later years. It is also possible that helpers do not help, but that able parents have many offspring each year, a fraction of whom stay behind. This would produce a correlation between number of 'helpers' and number of young raised, but b would be correctly assessed as zero. The general assumption lying behind the measure of b used is that the number of young produced in a year depends only on the number of helpers in that year. It is possible that experience of breeding in the first year is better or worse preparation for breeding in the second year than helping, or that there is differential mortality in the two groups. It is probably true that helping rather than breeding alone affects the chance of being helped in the second year. These would upset the measure of b , as it is the lifetime number of offspring that matters. Here the general assumption is that the rest of the juvenile's life is unaffected by its decision in the first year.

The point to notice about these assumptions is that although crucial they are interesting and biological and perfectly amenable to investigation, as is borne out by many later studies by Brown and his associates. They discover (among other things) with increasingly powerful methods whether or not helpers in the gray-crowned babbler really help (Brown *et al.* 1978; Brown *et al.* 1982).

In the example of the superb blue wren, the basic data were that pairs with helpers produced 2.83 offspring on average and those without produced only 1.50. The benefit of staying was therefore 1.33 to the parents. For females, the cost of staying was 1.50, as it was for males who could find mates; while for males who could not find mates, the cost was zero. The value of r was one, because the choice is between creating siblings and creating offspring, which are equally related. (Alternatively, we can say that the helper helps both his parents increase their number of offspring, and the sum of his relatedness to his parents is one.) The conclusions were that females should not help because

$$rb - c = 1.33 - 1.50 = -0.17 < 0,$$

and that males who could find mates should not help by the same calculation. Males who could not find a mate should help because

$$rb - c = 1.33 - 0 = 1.33 > 0.$$

Brown (1975) also discusses the case of the Florida scrub jay, using data from Woolfenden (1975), in which females should help. Emlen (1978) discusses at more length the application of Hamilton's rule to helping in birds.

Example 2: Noonan's study of joint nesting in Polistes fuscatus

Noonan (1981) studied the founding of nests by females of the social brown paper wasp *Polistes fuscatus*. In this species, nests may be founded by from one to about ten females; one of these females becomes the queen, who does most of the egg-laying. These females are almost always closely related. One of the questions Noonan asked was whether a female who joined a group of relatives as a worker did better than a female who decided to become a solitary queen. Specifically, did the fact that she was helping close relatives swing the balance in favour of social cooperation? We will not follow Noonan's own analysis.

To apply Hamilton's rule, it is vital to be precise about the wasp's decision. Suppose a female is confronted with a nest of $N - 1$ relatives and she knows that she is the last to decide whether or not to join it. If she does not join it we will assume, in two separate applications of the rule, first that she leaves and becomes a single foundress and second that she leaves and dies.

To estimate the benefit to her relatives, we must calculate how many young her relatives would produce as a nest of $N - 1$ females without her, and then in a nest of N females with her as an extra worker. The difference is the benefit she confers on her relatives. The cost to herself is the difference between the eggs of her own she lays as a worker in a nest of N females, and the young she rears in her alternative role. This is the number of young reared by a workerless queen in the first application, and the zero young reared by dying in the second. Noonan's paper contains all the necessary information to carry out these calculations, and the results are shown in Table 3.2.

The main conclusions are as follows.

1. In terms of her own number of offspring, a female is much better off as a solitary queen than as a worker, and much better off as a worker than dead.
2. In terms of her contribution to her nest-mates' number of offspring, her effect depends strongly on how many workers there already are at the nest. If she would be the only worker, then her effect is strongly positive; if she would not, then she has a negative effect on their reproduction. She lays more eggs as a worker than the extra she provides for the nest as a whole.
3. If she would be the only worker, then her relatedness to the queen may well favour her joining. The condition is that $r > 0.48$.
4. If she would not be the only worker, then her relatedness to the members of the nest will act against her joining. She would do better to join strangers and parasitize them. The value of r can even be high enough to swing the balance in favour of dying rather than joining close relatives in a group.

The implications of these conclusions do not matter to us here, but

Table 3.2. A kin selection analysis of the decision of a female *Polistes fuscatus* whether to join a nest of relatives and become a worker in a nest with a total of N females. Two alternatives are considered: becoming a solitary queen, and dying. Data are from Table 2-2 in Noonan 1981. The conditions for joining are calculated using ' $rb - c > 0$ means join'. We assume that r is never negative. Nests of sizes 1, 2, 3, and 4 were in fact common (Noonan 1981).

	N		
	2	3	4
Number of eggs of a solitary queen	18.25	18.25	18.25
Number of eggs of a worker in nest of size N	4.6	4.8	3.6
Cost of joining instead of becoming a solitary queen	13.65	13.45	14.65
Cost of joining rather than dying	-4.6	-4.8	-3.6
Number of eggs by the rest of the nest if joined (making N)	46.4	40.5	39.1
Number of eggs by the nest if not joined (making $N - 1$)	18.25	51.0	45.5
Benefit to rest of nest of being joined	28.15	-10.5	-6.4
When should a female join rather than become a solitary queen?	$r > 0.48$	Never	Never
When should a female join rather than die?	Always	$r < 0.46$	$r < 0.56$

one general point does. Why do we apply Hamilton's rule when Noonan's data are good enough to allow us simply to calculate number of offspring? In the spirit of section 3.3.1, should not the two methods give us the same answer and would not number of offspring be easier? One reason to apply Hamilton's rule is to see whether a trait is advantageous through an individual's own reproduction alone, or whether the effect on relatives' reproduction swings the balance. Another reason is that in counting number of offspring we must average over all animals who would have decided (i.e. had the genes for deciding) to join and those who would have decided to go it alone. But we do not know which all these animals are. The problem arises because not all animals are called on to make the decision, and so not all the animals' strategies are laid bare. We do not know who would have done what. Using the inclusive fitness approach, we can legitimately concentrate on only those animals who faced the decision. (The 'work' of deciding who would have done what is in effect done for us by the calculation of relatedness.) This is a crucial advantage to the inclusive fitness/Hamilton's rule approach when not all animals are faced with the decision of interest.

3.3.4 The validity of Hamilton's rule

Hamilton's rule holds good only under certain assumptions. There are different definitions of r , and the scope of the rule depends on the definition of r employed (Michod & Hamilton 1980; Seger 1981). Here we are concerned mainly with applications, and so restrict ourselves to forms of r that can be estimated from data. Charlesworth's derivation of Hamilton's rule (Charlesworth 1980) makes the roles of the assumptions clearest, and we follow his treatment in what follows. The latest, but perhaps not the last, word on the validity of inclusive fitness and Hamilton's rule is a review by Michod (1982).

Assumption 1: Additivity of costs and benefits

An animal that is helped m times, and helps n times, should experience a change of $mb - nc$ in its number of offspring as a result. This assumes that effects add. Addition will not always be the most plausible way for costs and benefits to combine. If the trait affects survival, then multiplication may be more appropriate (Charlesworth 1978). To see this, consider an animal that is exposed on two separate occasions to a 75% chance of dying. Each occasion quarters its fitness, and the overall effect is to reduce its fitness to one-sixteenth of what it would have been. The assumption of additivity is untroubling if the b and c are small, because then additivity will hold at least approximately. 'Small' is relative to average lifetime reproductive output.

In fact the assumption of additivity is really two assumptions in disguise. Two things go wrong if additivity is broken. First, an animal that pays one cost and receives one benefit does not have a net gain of $b - c$. This is partly just a measurement problem—if we define b and c as the average effects measured in number of offspring, then we may avoid it. Secondly, and more seriously, the cost and benefit of an action will depend on the genotype of the actor and recipient respectively. If two helps are not twice as good as one, then altruists will tend to lose out; they receive more than their share of helping and so receive more of the substandard second helps. This would invalidate the theory on which inclusive fitness is based at a quite fundamental level (Seger 1981).

Assumption 2: The gene frequency among potential donors and receivers is the same

A potential donor is an animal that is faced with the decision being investigated. A possible though unlikely exception is considered by Charlesworth (1978). Suppose that a dominant altruistic mutant gene caused all of its bearers to commit suicide for the benefit of their sibs. The sibs that benefited could certainly not be altruists! The altruistic allele would be inevitably selected against, indeed would have a

fitness of zero, no matter what the values of b , c and r . In a sense this assumption says that r has to mean what we think it means—the extent of genetic similarity at the locus of interest.

Assumption 3: Weak selection

This assumption hinges on r . The essential property of r in Charlesworth's derivation of Hamilton's rule is as follows (Charlesworth 1980). Let the set of animals S be the possible recipients of an act by an animal I , and let T be the set of all animals in the population. Let $p(A|Z)$ be the probability that an allele selected at random from a given locus in entity Z is allele A . (For example, $p(A|T)$ is the population proportion of A .) Then the r that is relevant for the decision is implicitly defined by:

$$p(A|S) = r \cdot p(A|I) + (1 - r) \cdot p(A|T)$$

In words, the possible recipients of the act are partly like I and partly like T ; and r measures how like they are to I . The condition that $r = 0$ means that the possible recipients are genetically representative of the population as a whole; $r = 1$ means that they are genetically of the same constitution as the potential donor.

A remarkable thing about this definition is that when there is no selection going on, r is the same for all alleles at all loci with the same inheritance pattern (e.g. autosomal, X-linked, Y-linked). Furthermore, it can be calculated from family trees. However, when selection is occurring, r will not be the same at all loci and will only be calculable approximately from family trees. Consider an animal and a class of its relatives at a period in the life cycle when selection is occurring. Selection is a systematic change in the relative proportions of animals with different genotypes; and so the genetic similarity of the animal to surviving members of the class of relatives must change as selection occurs. The assumption of weak selection is necessary to make ancestry an accurate enough guide to the 'true' r as defined above.

Weak selection should rarely be a problem in practice. If $rb - c$ is large in magnitude when calculated with an r derived from ancestry, then the deviation from the true value of r will not affect the sign of $rb - c$. Alternatively, if $rb - c$ is small in magnitude, then selection is weak and ancestry is a reasonable guide to the true r . Thus the direction of selection will be preserved, although it is true that the strength of selection may be misjudged.

It is interesting that these problems do not disappear if r is measured by electrophoresis. The deviation of the ancestral from the true r occurs at loci undergoing selection, and there is no reason to believe that the loci sampled electrophoretically are undergoing the same selection as the locus affecting the decision of interest.

3.4 GROUP SELECTION

Group selection has a bad reputation among many behavioural ecologists and evolutionists that obfuscates debates on some recent models that go under the name of group selection. In section 3.4.1 I discuss that kind of group selection which quite rightly elicits disapproval, and which is certainly still with us today. Section 3.4.2 deals briefly with the first generation of models of group selection, and section 3.4.3 deals more extensively with a quite different class of models of the natural selection of animals in groups.

3.4.1 The bad and dangerous

The reputation of group selection comes not from mathematical models, nor from deliberate discussion of group selection, but from a certain naivety practised by laymen and many biologists alike in their day-to-day thinking about the adaptedness of animals to their environment. An adaptation is 'for something'—and that certain naivety is to propose an adaptive explanation without stopping to think what that 'something' is. It is this lack of thought, not any deliberate and considered choice of the group as that 'something', which is the target of most accusations of 'group selectionism'.

This lack of thought is a target deserving of attack—and as sustained and effective an attack as we can muster. However, 'group selection' is a poor name for this lack of thought, because the 'good of the species' and the 'good of the ecosystem' are as prominent as the 'good of the group' in its effects (to be found in journalistic and academic publications).

The fundamental case for careful thought is that adaptations arise only by natural selection, and that natural selection does not normally promote adaptations for the good of any unit larger than the organism. Two excellent 'self-help' guides in careful thought are books by G. C. Williams (1966) and R. Dawkins (1976). (The organismal approach suggested here is not in conflict with the 'gene selectionism' of Dawkins (1982a, b). In his language, we are saying that the individual is usually a well-adapted vehicle for gene replication, while groups usually are not.) In the next two sections we will see possible exceptions to the general rule, but we must not allow them to distract us from a most important lesson about adaptations: a very convincing case is needed to explain why an adaptation should be for anything other than the organism. Lack of thought is the basis for that group selection which is mad, bad and dangerous to know, and the exceptions that follow are no licence for that laziness.

3.4.2 The old

The first generation of mathematical models for the exploration of group selection can be understood by reference to Fig. 3.1, which is

adapted from Maynard Smith (1976b), who gives references to examples of these models. A composite of those models works as follows. There are a large number of discrete locations each of which is capable of supporting one group. Migration between locations is restricted. There are two alleles, A_0 and A_1 , at a given locus. Animals with A_1 are more cooperative than those without, but at a cost to personal fitness. The consequences are that within any group A_0 quickly displaces A_1 , but that groups consisting of A_1 are better off than groups of A_0 . 'Better off' means either less likely to go extinct in a year, or able to produce more migrants to leave and try to colonize empty locations.

There are three kinds of location (neglecting the transient mixtures of A_0 and A_1): E (for empty), A_0 and A_1 . The traffic between these kinds occurs for a number of reasons. Extinction sends A_0 and A_1 to E. Migration to empty locations sends E to A_0 or A_1 . Migration to occupied locations sends A_1 to A_0 .

The question for any given model is exactly how the cooperation of A_1 animals affects extinction and number of migrants produced by the group, and whether migrants are allowed to join occupied locations or only empty ones. The factors promoting the spread of A_1 are (1) cooperation reducing extinction, (2) cooperation increasing number of migrants produced by the group, (3) migrants allowed to join only empty locations, and (4) small number of founders in a group. The last is important because the more founders there are, the more likely it is that there is at least one A_0 among them.

The final consensus on these models was that the conditions for A_1 to be successful were too stringent to be realistic. Wynne-Edwards, whose book *Animal Dispersion in Relation to Social Behaviour* (1962) sparked the whole controversy, wrote in 1978:

'but in the last 15 years many theoreticians have wrestled with it, and in particular with the specific problem of the evolution of altruism. The general consensus of theoretical biologists at present




Fig. 3.1. This shows the possible states of sites in an 'old' group selection model. The arrows represent possible transitions. E means empty; A_0 and A_1 refer to groups with only that allele at the locus of interest. Mixtures of A_0 and A_1 are considered too transient to matter. The figure is adapted from Maynard Smith 1976b.

is that credible models cannot be devised by which the slow march of group selection could overtake the much faster spread of selfish genes that bring gains in individual fitness. I therefore accept their opinion.' (Wynne-Edwards 1978).

Even before this was written, a new sort of group selection had been discussed that did not require unreasonable assumptions. The new group selection of Price (1970, 1972) and Hamilton (1975) is the topic of the next section.

3.4.3 The new

Many animals live in groups, and these groups may be grouped by natural features such as rivers or mountains, and even these super-groups may be grouped by, for example, the ocean. The focus of attention of the new group selection is on this hierarchy of grouping and the effect it might have on natural selection. As a preliminary caution, we must not allow the rhetoric of hierarchies and groups within groups to intoxicate us. A sober appraisal is required, for as we shall see it is possible to have hierarchies and groups within groups that have no effect whatsoever on the workings of natural selection. The next three parts deal in turn with altruism, the sex ratio and, briefly, a different approach to explaining the workings of the following models.

Altruism

The real hero of this section is Hamilton's rule, and indeed we follow Hamilton's treatment closely here (Hamilton 1975). To understand why grouping might have an effect, we first examine the definition of r given in section 3.3.2 above. It was that for an animal I, a set of other animals S, and the total population T, r from the animal to the set satisfies:

$$p(A|S) = r \cdot p(A|I) + (1 - r) \cdot p(A|T).$$

In words, the genetic composition of S is a weighted average of the genetic compositions of the animal I and the population T; r measures the weight given to the animal I's genotype in that average.

Taking S to be the rest of the group to which an animal belongs, what factors might cause r to be non-zero? There are only two (Maynard Smith 1976b). One is common ancestry, and the other is preferential assortment. Taking common ancestry first, the grouping of the population may make it difficult for us to ascertain properly all the relevant kin links. As an example, Hamilton points out that in a virtually closed group the genetic similarity builds up eventually to $1/(2M + 1)$, where M is the absolute number of migrants each generation, independent of the group's size. The ties in a large group 'make up in multiplicity what they lack in close degree' (Hamilton 1975).

So unsuspected kin links may increase r above initial expectations. An example of how to estimate relatedness in a grouped population is provided by Bertram (1976). By modelling the way lion prides are formed, he calculated the genetic similarity between the different kinds of animals in a pride—the females, the males and the young. Just using the observed kin relations in a pride would have underestimated the genetic similarities. Bertram's work is a good example of the new group selection applied to data, although it is expressed entirely in terms of kin selection. An important message from it is that animals in groups may not all be equally related, and the distinctions are very interesting.

The other way for r to be raised in groups besides common ancestry is preferential assortment. If altruists share a preference for certain habitats or microhabitats, then altruists will tend to be in groups with other altruists and that is all that is required to make r positive according to the definition above. There are a number of reasons why preferential assortment is not a plausible source of genetic similarity in nature. They are based on the fact that while ancestry provides relatedness that is the same for all loci, preferential assortment only causes relatedness at the loci that cause it, and at linked loci. It is unlikely that the locus for altruism is closely linked to the loci for habitat preference. Even if it were, there would be selection at unlinked loci to suppress the altruism; for while the r at the altruism locus may be positive, the r at unlinked loci is zero and Hamilton's rule applies equally to both sorts of loci. Finally, there would be selection for a 'free-rider' allele (if one arose) at a locus unlinked to the altruism locus. It would have the effect of creating the same habitat preference as that of altruists, whether or not its bearer was an altruist.

For all these reasons the only plausible cause of genetic similarity among group members is common ancestry. Furthermore, the only moving force in the new group selection is genetic similarity.

Hamilton (1975) discusses a model of towns, with low migration, in which relatedness and altruism build up together; he suggests this as a possible genetic basis for xenophobia. The model would apply equally well to non-human groups. It is the most considered attempt so far to show how the new group selection might be important in nature. *I wish now to discuss two important caveats in interpreting this model.* The model itself works because, as noted above, the relatedness in a community that receives M migrants each generation builds up to $1/(2M + 1)$.

The first caveat is that the migration rates must be low to achieve a noticeable relatedness. One migrant every two generations produces sibling level relatedness; eight per generation produces second cousin level relatedness. In animal groups in which all of one sex disperses we expect very little effect.

The second caveat is more fundamental. When we say in the ordi-

nary, ungrouped, model of kin selection, that an act has effects b and c , do we really mean that there are $b - c$ extra offspring in the population as a whole surviving to maturity and breeding? Perhaps, but probably not. Normally we would expect that a winter bottle-neck or some other factor would limit total population size in a way unaffected by the act of altruism. We probably mean to say that the donor loses c , the recipient gains b , and then because the total population number remains the same the population as a whole loses $b - c$. If we let d represent the general decrement associated with the act (and d will equal $b - c$ if the act does not increase carrying capacity), and let r_e be the average relatedness of the potential donor to those suffering the general decrement, then we can write Hamilton's rule more fully as:

$$rb - c - r_e d > 0$$

Formally this is just expanding b to include all the effects of the act. The subscript 'e' means economic, describing the forces that regulate population size. Hamilton (1964) discussed this near conservation of fitness in the population as a whole.

If the population as a whole suffers the general decrement, then the expansion of Hamilton's rule is unnecessary. By the definition of r above, the relatedness to the population as a whole is zero and so the extra term falls out. However, in the case of grouped populations the extra term will often be very important. The most important factor in creating intra-group relatedness, isolation, will also cause the general decrement of an act to be felt by those with whom an animal is isolated. This means that r_e will tend to be proportional to the intra-group relatedness. In the extreme case where all the general decrement falls on the group, and the act does not increase carrying capacity, there will be no selection for intra-group altruism; for what the animal gives in bulk to one member, he takes away in dribs and drabs from all the others.

There are then two parallel factors at work in the new group selection. One is the pattern of relatedness among groups, and the other is the pattern of joint dependence of offspring of group members on the same resources. The resolution of these two, through the expansion of Hamilton's rule above, determines how grouping affects the workings of natural selection on altruism.

There remains one loose end to tie. It is obvious from this discussion that with random grouping there is no selection for altruism. In this case $r = 0$, and the spread of an allele is determined by the selfish criterion $c < 0$. Matessi and Jayakar (1976), Cohen and Eshel (1976) and Wilson (1975, 1980 and references therein) have claimed that there is selection for altruism in this case. This arises because they all defined altruism not in terms of absolute fitness as Hamilton (1964, 1972) did before them, but in terms of relative fitness within the group. This is discussed further at the end of this section.

Sex ratio

In Chapter 8, Maynard Smith explains the basic Fisherian model (Fisher 1958) and the idea of local mate competition. The first model showing an effect of grouping on sex ratio was constructed by Hamilton (1967). In this model, n inseminated females arrive at a site and lay their eggs, which develop and then mate among themselves. The males die, the inseminated females disperse and the next generation begins. The unbeatable sex ratio (for a diploid species) is that a fraction $(1/2)((n-1)/n)$ of resources should be allocated to producing sons.

This female bias can be seen as the result of two factors acting against the Fisherian force towards equality of investment. The first is diminishing returns to producing sons—that is, each additional son creates fewer and fewer additional grandoffspring for his mother. The second is that making a daughter increases the number of grandoffspring males produce for their mother, through increasing the available number of mates; creating sons does not affect daughters' success. This asymmetry arises because it is assumed that males contribute nothing to the care of offspring. We can add a third factor by supposing the n females to be related. Making a daughter would then have the additional benefit of increasing the mating success of sons of relatives. The sex ratio would therefore be even more female biased if the n females were related to each other.

These three factors are at work in a recent model of the sex ratio in grouped populations, by Bulmer and Taylor (1980). It is designed to account for female-biased sex ratios in the wood lemming. In the model, n foundresses arrive at a site, and g generations take place within the site before population-wide dispersal occurs and again groups of n foundresses form. Bulmer and Taylor consider separately the case where females of the dispersal generation mate within their own group before dispersal, and where they mate at random in the population. They sought that sex ratio which, once an allele producing it was common in the population, would not allow a dominant allele coding for any other sex ratio to invade. This stable sex ratio depended on the parameters n and g of the model. It was always female biased. As the number of foundresses n increased, the bias diminished, corresponding to a relaxation in the diminishing returns of producing sons. The bias increased as g , the number of generations between dispersals, increased. How can this be explained?

During the g generations between dispersals, the number of animals in each site grows exponentially at a rate determined by the sex ratio produced by the females. This creates increasing returns to scale for producing daughters when measured in number of eventual dispersers. It also causes relatedness to increase as the g generations proceed. If each of n unrelated foundresses has two daughters and one son, then each member of the next generation has as sibs a fraction

$2/(3n - 1)$ of the rest of the population. Now the exponential growth of the group means that producing daughters benefits all members of the group, and of course a female only cares about the effect on the reproduction of the others to the extent that she is related to them.

This view of the model suggests two predictions about the model's behaviour. First, the sex ratio bias should be more extreme if variation in reproductive success among males or females exists, since this also increases the relatedness between group members. (In an extreme case where all offspring have the same parent of either sex, group members are at least half-siblings.) Secondly, the sex ratio bias should increase as the g generations proceed because the relatedness between group members increases. Bulmer and Taylor do not allow the females to choose their sex ratio according to the generation to which they belong, nor do they vary variance in reproductive success.

In a similar model, however, Wilson and Colwell (1981) do vary variance in reproductive success. Their model is haplodiploid, and their equilibrium sex ratio is a genetic polymorphism between an allele producing a 1:1 sex ratio, and an allele producing a sex ratio of varying degrees of female bias. They confirm the results of Bulmer and Taylor, and also show that the result of increasing the variance in male reproductive success is indeed to increase the female bias of the sex ratio.

So we see that the female bias in sex ratios that can occur in grouped populations can be attributed to diminishing returns or increasing returns to one sex, producing females being a source of reproductive success for sons, and producing females being a source of reproductive success for the sons of relatives.

Didactics

The most vocal proponents of the new group selection have been Wilson and Colwell. (Wilson 1980, and references therein; Wilson & Colwell 1981; Colwell 1981.) They have scandalized many by speaking positively of 'group selection'. The purpose of this section is to reconcile their chosen way of explaining why grouping has the effect it does with that of the previous sections. They concentrate on dividing selection into two parts: within-group selection, which is the local change in gene frequency within each group, and between-group selection which is the result of differential fecundity of groups. This difference of approach is purely didactic—there is no disagreement about matters of substance.

The different approach has, however, led to misunderstanding. Wilson and Colwell identify within-group selection with 'individual selection', and between-group selection with 'group selection'. Now individual selection already has a meaning which is quite different, and one I think is very valuable. An act is said to be favoured by individual selection when it spreads through its effect on the actor's

number of offspring alone. In terms of Hamilton's rule, this occurs when there is a negative cost, i.e. a simple gain, in terms of number of offspring to the actor. As r and b also enter into Hamilton's rule, this definition does not make individual selection synonymous with 'natural selection'. See the discussion of Noonan's study in section 3.3.3 above as an example.

Another source of misunderstanding arises from the use of the word 'altruism'. As we noted earlier, altruism will not evolve in simple, one-generation groups that are formed at random from the population. Matessi and Jayakar (1976) and Cohen and Eshel (1976), as well as Wilson (1975, 1980 and references therein), redefined altruism to refer to relative success within the group rather than absolute success. Relative success is the individual's number of offspring divided by the average number of offspring of members of his group. Absolute success is number of children (or number of children relative to the whole population). Under the 'relative' definition, 'altruism' can spread. Wilson calls the acts that are altruistic under the 'relative' definition, but not under the 'absolute' definition, 'weakly altruistic'. An alternative I prefer is 'a self-interested refusal to be spiteful'.

Wilson and Colwell's 'between-group variance' is very closely connected to relatedness to other group members, and the same points can be made using both concepts. In a haploid model with constant group size, they are connected by the following formulae:

$$v = v_b(1 + (n - 1)r) \quad r = \frac{1}{n - 1} \left[\frac{v}{v_b} - 1 \right]$$

where v_b is the variance that would arise from a binomial distribution of the same overall proportion of genotypes. Several differences may be noted between r and v . The expected degree of altruism depends on r by Hamilton's rule, and not on n . Knowing that $r = 0.22$ gives many biologists an understanding of the genetic closeness described; the knowledge that $n = 10$ and $v/v_b = 2.98$ is (at least for the present) less illuminating. When there is no effect of grouping, that is the grouping is random, $r = 0$; and when spite is expected (Hamilton 1975), r is negative. The value of v itself gives little indication of the effect of grouping; $v/v_b = 1$ for random grouping and < 1 when spite is expected. I think it will be admitted that r is the more useful and familiar measure of genetic similarity. Familiarity is important for clarifying the connections between the new group selection and what is already known about kin selection.

The connections are certainly there. Bertram (1976) correctly described his study as one in kin selection, and yet I used it as an example of the new group selection applied to data. Once the basis of the new group selection is understood (namely, genetic similarity due to kinship but where groups are clearly in evidence) most kin selectionists should realize they have been new group selectionists all their

lives. It is vital to remember, of course, that when the population is grouped there may be unsuspected kin links; and that in groups that last for a number of generations, relatedness builds up as the generations proceed.

3.5 CONCLUSIONS

1. The methods of behavioural ecology depend for their correctness on various genetic assumptions.
2. Number of offspring and inclusive fitness are two different valid measures of reproductive success.
3. Inclusive fitness has often been misdefined. It includes relatives' offspring only if the animal's help is responsible for their existence. It excludes those among the animal's own offspring that exist because of help received from others.
4. Hamilton's rule in the form ' $rb - c > 0$ ' should be used in applications of the inclusive fitness approach to data. As a general rule, inclusive fitness is applied wrongly to data but Hamilton's rule is applied correctly. This is because Hamilton's rule encourages use of the correct logic of differences.
5. New group selection models are most readily understood using Hamilton's rule. Genetic similarity as expressed in relatedness is the driving force towards altruism. The forces affecting sex ratio are increasing or decreasing returns to scale for one sex or the other, and the increase in male relatives' mating success that follows from producing a daughter.