

Chapter 10: The evolution of co-operation

A fundamental question about all forms of collective animal behaviour is how they evolved. At various points in this book I have turned to arguments based on individuals adopting or evolving behaviours that increase their own fitness to explain or make predictions about group behaviour. For example, group size distribution was described in terms of individuals attempting to join a group of a size which maximises their fitness (chapter 2); foraging birds were described as balancing searching for food themselves with copying others (chapter 3); consensus decision-making and synchronisation were described in terms of individuals co-ordinating so they can benefit from acting together (chapters 4 and 7). While such functional arguments are not the only way to understand the behaviour of groups (and indeed have played a secondary role to mechanistic explanations in the other chapters of this book) they are an essential part of biology. This chapter gives an overview of how functional reasoning can be applied to collective animal behaviour.

The theory of natural selection is grounded in the idea that those individuals exhibiting a behaviour that provides them with higher than average fitness pass their genes, and thus their particular behaviour, on to future generations. It is this idea that provides the basic assumption of evolutionary game theory models: those individuals adopting a strategy that provides them with higher than average fitness will increase in the population, while those with lower than average fitness will decrease. Despite the simplicity of this underlying assumption, these models have proved extremely powerful in predicting when co-operation between animals will evolve (Dugatkin & Reeve 1998; Maynard Smith 1982). As a result of this success, a vast literature has arisen on the evolution of co-operation, both theoretical and experimental. The size of this literature makes it difficult to give a concise account of how different models and experiments relate to one another. There is however an increasing consensus of how co-operation should be discussed in evolutionary biology (Clutton-Brock 2002; Foster et al. 2006; Lehmann & Keller 2006a; Lehmann & Keller 2006b; West et al. 2007). In this chapter, I follow this consensus, and categorise how collective and co-operative behaviour can evolve between non-relatives in four different ways: through parasitism, mutualism, synergism and repeated interactions. In doing so I discuss evolutionary game theory, which has become a central modelling tool in understanding co-operation between non-relatives.

A complication arises when we consider the evolution of altruism. Altruism is defined as individuals paying a cost greater than any resulting benefit to their average lifetime reproductive success while providing a benefit to the lifetime reproductive success of others. Under natural selection, an individual adopting an altruistic strategy will suffer a drop in fitness and will be less likely to directly contribute offspring to the next generation. By introducing genetic relatedness into evolutionary game theory models, however, we see that the inclusive fitness equation or Hamilton's rule provides a good predictor of when altruism can evolve. Altruism can evolve when individuals help their relatives, thus indirectly passing their genes to future generations. This chapter should thus provide a broad classification of many of the collective behaviours discussed in this book as arising from a combination of four distinct forms of co-operation (parasitism, mutualism, synergy and repeated interactions) and altruism arising from inclusive fitness.

10.1 Evolutionary game theory

Evolutionary game theory models describe how selection acts on behavioural strategies. Over repeated generations a strategy will increase in the population if it receives a higher than average payoff or fecundity, but decrease if it receives a lower than average payoff or fecundity. The biological interpretation of these arguments is different when considering evolution through natural selection or individual decision-making (Dugatkin & Reeve 1998). Under Darwinian natural selection, we assume that a genotype encodes for a particular strategy throughout an individual's lifetime and ask whether this genotype will increase or decrease in the population over generations. Those genotypes that provide lower fecundity die without reproducing and those providing higher fecundity produce offspring that fill their places. Under individual decision-making, we consider an individual that can change its strategy during its lifetime in response to its mistakes. This individual plays its strategy. If it then gets a payoff higher than average it keeps the strategy; otherwise, if its payoff is lower than average, it changes strategy towards one which will improve its payoff.

Evolutionary game theory assumes the following lifecycle for individuals. An infinitely large population of individuals undergo the following stages:

1. Individuals form groups of size N . Individuals are distributed entirely at random between groups, i.e. dispersal is global.
2. Each individual has a behavioural strategy, s_i . The payoff or fecundity of each individual i is determined by its own strategy and of all other

individuals in the group, i.e. $f(s_i, s_1 \dots s_N)$. No reproduction occurs while individuals are within the group.

3. Individuals then leave the group and a law of selection is applied to them: each strategy's contribution to the next generation is proportional to its fecundity relative to the average fecundity of the entire population (i.e. not just those in the group of size N). This contribution to the next generation is also known as the individual's fitness.

In many evolutionary game theory models it is often further assumed that $N=2$ (Maynard Smith 1982). For $N>2$ these assumptions are the same as "group selection" models (Nunney 1985; Wilson 1983). Provided $N<\infty$, the probability of repeatedly interacting with the same individual on consecutive generations is zero, as is the interaction probability for two related individuals, e.g. individuals with the same parents.

Some of these assumptions can be relaxed and the results of these models remain the same. For example, Nunney (1985) shows that a population interacting with local neighbours in a continuous space, rather than discrete groups, but dispersing globally before reproduction gives similar predictions as those from the lifecycle above. What cannot be relaxed however is the assumption that there is a zero probability of interacting with the same individual twice or of interacting with relatives. These two cases are dealt with separately below in the sections on repeated interactions and inclusive fitness, respectively. For a good discussion and justification of the other assumptions underlying these models see Grafen (1984), Nunney (1985), and Dugatkin & Reeve (1994).

Evolutionary games where $N=2$ and individuals choose between two distinct strategies can be expressed in terms of a payoff table. Table 10.1 gives such a payoff table for interactions between 'co-operators' and 'defectors'. In this game we assume that a focal co-operator pays a cost C and confers a benefit B to its partner. If the partner defects the focal co-operator still receives a *direct benefit* D , while if the partner co-operates it receives the benefit B plus an *extra benefit* E . A focal defector pays no cost but receives the benefit B if its partner is a co-operator. The model resulting from the above assumptions and this payoff table is analysed in Box 10.A.

The distinction between the different types of benefits B , D and E lies at the heart of understanding the evolution of co-operation. Figure 10.1 shows how the relationship between the parameters D and E and the cost C determines the

evolutionarily stable proportion of co-operators and defectors in the population. The figure gives four qualitatively distinct evolutionary outcomes, which I call parasitism, mutualism, synergy and 'failed' altruism. I now examine each of these scenarios separately, also discussing how the results for these two-player, two-strategy games relate to N -player and continuous strategy games. It turns out that two player games provide quite a general classification of the possible forms of co-operation, even when we add these further complications.

Box 10.A Two-player discrete strategy evolutionary games.

In addition to the assumptions in section 10.1, I further assume that $N=2$ and there are only two discrete behavioural strategies, called co-operate and defect. Interactions in these pairs give a payoff, i.e. fecundity or number of offspring produced, that can be expressed in a two by two payoff table (table 10.1). A focal individual adopting the co-operative strategy always pays a cost C and always confers a benefit B to its partner. Depending on the strategy of the partner the focal individual receives a *direct benefit* D , if the partner defects, or an *extra benefit* E , if the partner co-operates. A focal defector pays no cost but receives the benefit B if its partner is a co-operator. Note that B , C , D and E are assumed to be positive constants that do not depend upon the frequency of co-operators or defectors in the population as a whole.

The payoff to an individual depends on the strategy of their partner. Since pairs are selected entirely at random from an infinite population, the probability that a focal individual interacts with a co-operative individual is equal to the proportion of co-operators in the population as a whole, denoted x . Thus the expected payoff of a co-operator is

$$x(B + E - C) + (1 - x)(D - C)$$

and the expected payoff of a defector is

$$xB$$

These are the frequency-dependent payoff functions, which determine how many offspring individuals produce. The average payoff of the population is

$$x(x(B + E - C) + (1 - x)(D - C)) + (1 - x)xB$$

Natural selection implies that the proportion of individuals with higher than average payoff will increase and the proportion of individuals with lower than average payoff will decrease in the population. Thus the fitness of an individual can be defined as being equal to its payoff divided by the average payoff of population, i.e.

$$f(x) = \frac{x(B + E - C) + (1 - x)(D - C)}{x(x(B + E - C) + (1 - x)(D - C)) + (1 - x)xB}$$

We can then write the following expression for the rate of change of the proportion of co-operators in the population

$$\frac{dx}{dt} = x(\text{payoff of a co-operator} - \text{average payoff of population})$$

which in this case is equal to

$$\frac{dx}{dt} = g(x) = x(1 - x)(x(E - C) + (1 - x)(D - C))$$

Solving this equation as $t \rightarrow \infty$, i.e. $\frac{dx}{dt} = 0$, gives us the conditions under which co-operation is selected for.

Solving $\frac{dx}{dt} = 0$ gives steady states at $x_* = 0$, $x_* = 1$ and $x_* = (C-D)/(E-D)$. These correspond respectively to a population of all defectors, all co-operators and a mixture between co-operators and defectors. By differentiating $g(x)$ with respect to x and evaluating at x_* we find the conditions under which the addition of a small number of co-operators or defectors will lead to the growth of that strategy in the population away from the steady state. These steady states are said to be evolutionarily stable to such small perturbations whenever $g'(x_*) < 0$. For example, $g'(0) = D - C$ so all defect is stable if $C > D$. Similarly, all co-operate is stable if $E > C$. The mixed strategy steady state, $x_* = (C-D)/(E-D)$, exists (i.e. lies between 0 and 1) and is evolutionarily stable if $D > C > E$. The mixed strategy steady state exists but is unstable if $E > C > D$. These results are summarised and categorised in figure 10.1.

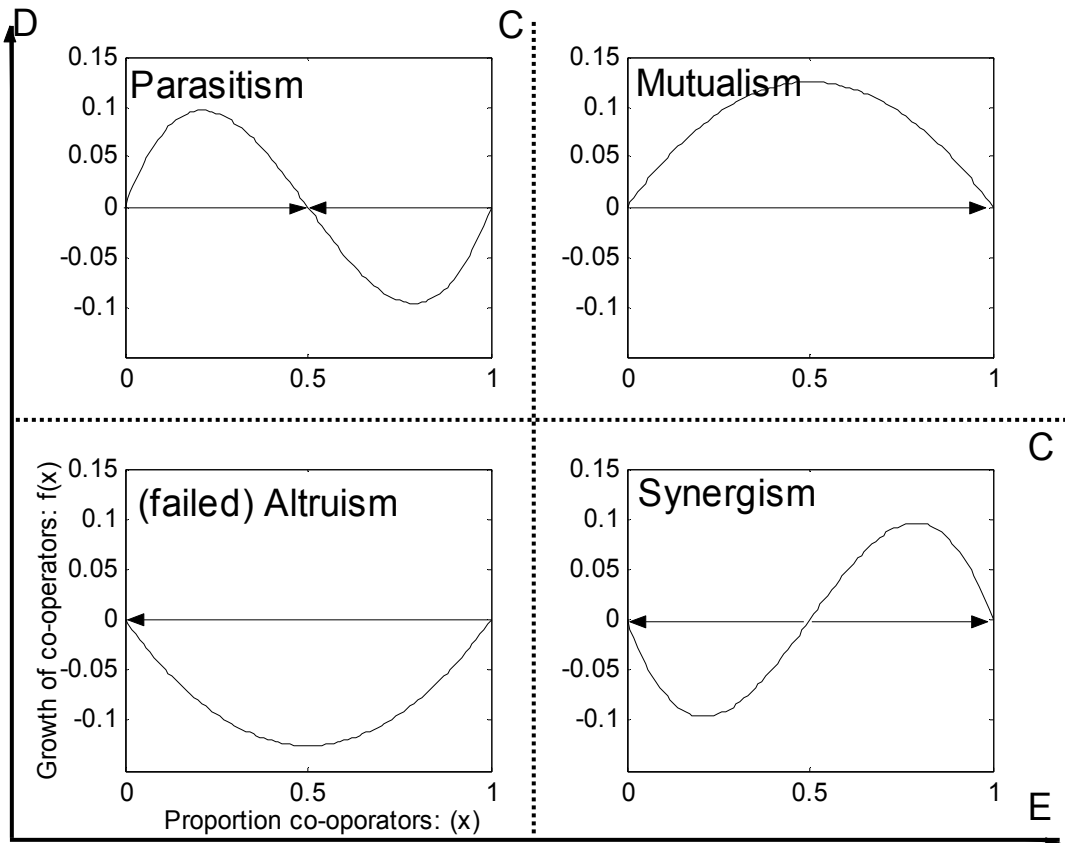


Figure 10.1 Outcome of two player evolutionary game. Shows how the cost benefit parameters, C , D and E determine the evolutionary stable strategies in two player, two strategy games. The axis of the main figure are the benefits E and D . Each of the four panels within the figure shows one of four qualitatively different forms of the fitness function $f(x)$, given by equation ? in Box 10. **Error! Reference source not found.** The arrows show how given a particular proportion of co-operators in the population whether co-operation will increase or decrease. If $D > C$ but $E < C$ then the outcome is parasitism, a single evolutionary stable state where the population consists of a mixture of co-operators and defectors; if $D > C$ and $E > C$ then the outcome is mutualism, the evolutionary stable state is to all co-operate; if $D < C$ and $E > C$ then the outcome is synergism, where there are two evolutionary stable states one corresponding to all defect and one to all co-operate; finally, if $D < C$ and $E < C$ then the evolutionary stable state is to all defect. This last case is classified as failed altruism, because were individuals to co-operate in such a situation their actions would be altruistic and result in negative direct fitness for an individual.

10.2 Parasitism

It's Sunday morning, the living room is a mess, but the newspaper has just arrived. You and your partner have the option of either tidying the living room or sitting down and enjoying the lifestyle supplement. The decision can be interpreted in terms of the costs and benefits in Table 10.1. Assume that D is the direct benefit you gain from having a tidy living room, but you pay cost C in the time you spend tidying. Further assume that the benefit of having a tidy room is greater than the cost of cleaning, so $D > C$. On the other hand if your partner tidies up

then you can defect and get the benefit B of reading and having a tidy room. If you tidy up together then you save some time E , but because tidying up always takes some positive amount of time, $C > E$.

Identifying $D > C > E$ in figure 10.1 shows that a unique evolutionarily stable state exists for individuals that wake up in a different house with a different partner every Sunday. A single defector in a population of co-operators benefits from the work of the others, leading to an increase in defection. However, when everyone defects then they all receive the lowest possible payoff and it pays for an individual to co-operate instead. The proportion of individuals who co-operate evolves to $(C-D)/(E-D)$. At this point no one individual can do better on average by changing strategy. The lack of stability of both pure co-operation and pure defect leads to a 'compromise' of some individuals that work and some which 'freeride' or parasitize the work of others.

Social parasitism in groups

Social parasitism is of relevance to a large number of social and biological situations. The producer-scrounger game discussed in chapter 3 in Box 3.B describes a situation where a focal foraging bird can either search for food itself (produce) or watch others search and share their finds (scrounge). The strategy an individual chooses depends on the strategies of others: it pays to scrounge in a population of pure producers and it pays to produce in a population of pure scroungers. The evolutionarily stable strategy, of a mixture between some scroungers and some producers, is widely observed in the behaviour of foraging birds (section 3.5).

A term that often arises in describing N -player social parasitism is the 'Tragedy of the Commons' (Hardin 1968). Imagine a commons pasture on which any individual can place their cattle. This pasture has a limited capacity and having cattle above this capacity causes the pasture to deteriorate in quality. A tragedy arises because the benefit of adding a cow goes directly to the owner, while the cost in deterioration is shared amongst everyone. Thus even when the pasture is over capacity, it can pay for an individual to add a cow. Ultimately, the pasture will yield an average payoff below that yielded at capacity. The same is true in the producer-scrounger game: at the evolutionarily stable state the average intake of each forager is lower than it would be if they all searched independently. And again for vigilance: flocks are less vigilant than would optimise their energy intake (Fernandez-Juricic et al. 2004b). And again in the

group size paradox in chapter 2: it pays solitary individuals to join a group even if this group size will become suboptimal for those already in it.

Hardin uses the 'Tragedy of the Commons' to argue for regulation of individual freedom, since individual freedom will lead to disadvantage for all. While this argument has some validity, the outcome of social parasitism is not as bleak as it might at first seem. Provided there remain direct benefits for co-operating, i.e. $D > C$, then the tragedy does not lead to everyone receiving the lowest possible payoff of zero. When no-one else is producing or being vigilant, it always pays to co-operate. The fact that in a fully co-operative society it pays for some individuals to defect does not imply that everyone will defect. We should not be surprised when we see animals or humans co-operating despite parasitism of their efforts by others, but should first consider what direct benefits are gained by the co-operators (Clutton-Brock 2002; Griffin & West 2002).

Continuous strategies

In the model in Box 10.A we assume that there are two distinct strategies, one corresponding to co-operate and the other to defect and study how the proportion of the population adopting each strategy evolves. Another approach is to have a continuous strategy, expressing the level of investment made in co-operation. For example, we can think of each individual in our population of living room occupants as investing a level, p , in tidying up. It is tempting to think of the proportion of the population and the level of investment as being the same thing, and conclude that individuals with continuous strategies will adopt a level of co-operation equal to the evolutionarily stable strategy in the discrete strategy game. The problem with this interpretation is that a population where all individuals have strategy p_* can be, although not always is, invaded by two strategies that lie either side of p_* (Vincent & Brown 1984). Ultimately, a sequence of invasions may lead to the population again dividing into having two distinct strategies, one which always co-operates and one which always defects.

The above observation is important because animals are often faced with decisions not simply as whether to co-operate or not, but how much they want to invest in a co-operative behaviour, such as guarding or searching for commonly exploitable food. The importance of continuous strategies has been widely acknowledged since evolutionary game theory was first developed

(Zeeman 1981). Geritz et al. (1998) provided an elegant framework for determining when evolution will converge on a particular strategy. Box 10.B shows how this framework, often referred to as adaptive dynamics, can be used to classify an evolutionary game with continuous strategies in which a focal individual who invests p in a co-operative behaviour receives payoff

$$P(p, q) = B(p, q) - C(p)$$

Box 10.B Two-player continuous strategy games.

Doebli et al. (2004) describe an evolutionary game theory model where $N=2$ and players meet in pairs. Each player invests a level between 0 and 1 in co-operative behaviour. They assume p is the level investment of the focal individual, q is the level of investment by its partner and the focal individual receives payoff

$$P(p, q) = B(p, q) - C(p)$$

Consider a population in which every individual adopts the same strategy q , apart from a rare mutant that adopts strategy p which is slightly different than q . Whether a mutant invades and the resident strategy changes through time is determined by the selection gradient

$$\begin{aligned} D(q) &= \frac{\partial}{\partial p} (P(p, q) - P(q, q)) \Big|_{p=q} \\ &= B'(q, q) - C'(q) \end{aligned}$$

Steady states q_* of $D(q_*) = 0$ are said to be convergent stable when $D'(q_*) < 0$, where the derivative here is taken with respect to q . The convergent stable states are the points towards which the population first evolves in figure 10.2.

The fact that q_* is convergent stable does not imply that it is the final resting point of evolution. It is possible that two mutants, one with a slightly larger investment than q_* and one with a slightly smaller investment than q_* can simultaneously invade the resident q_* . Doebli et al. (2004) show that the convergent stable state q_* is the final stopping point of evolution

if $B''(q_*, q_*) - C''(q_*) < 0$, where the partial derivative here is taken twice with respect to p and then evaluated at $p=q=q_*$. In this case, every individual makes the same investment q_* in co-operation (e.g. figure 10.2b). If $B''(q_*, q_*) - C''(q_*) > 0$ then q_* is said to be an evolutionary branching point and the population separates into individuals that adopt more co-operative and less co-operative strategies (e.g. figure 10.2a). Details of the above classification of steady states is given by Geritz et al. (1998).

Doebli et al. (2004) looked at non-linear benefit and cost functions of the form

$$B(p, q) = b_2(p + q)^2 + b_1(p + q)$$

$$C(p) = c_2 p^2 + c_1 p$$

and found a steady state at $q_* = (c_1 - b_1)/(4b_2 - 2c_2)$, which exists and is convergent stable when $4b_2 - 2c_2 < c_1 - b_1 < 0$ and which exists and is evolutionarily stable when $c_2 > b_2$. By insisting that investment in co-operation must be between 0 and 1, two further steady states are created at these points.

There are five different evolutionary outcomes of this model, examples of which are shown in figure 10.2. Four of the five outcomes are identical to the outcomes of the discrete strategy model shown in figure 10.1: some individuals defect, some co-operate (figure 10.2a; parasitism); depending on initial population all defect or all co-operate (figure 10.2c; synergy); all defect (figure 10.2d; failed altruism); or all co-operate (figure 10.2e; mutualism). The additional outcome (figure 10.2b) arises when the mixed strategy q_* is evolutionarily stable. In this case individuals balance their own investment in co-operation so as not to be exploited. The continuous model thus distinguishes the cases where the population all adopts the same strategy or splits between two strategies, but in both these cases the interactions are social parasitic.




Figure 10.2 Outcomes of a simulation of a continuous investment two player game. The top row shows the evolutionary dynamics of the strategy distribution; darker shades indicate higher frequencies of a strategy. The singular strategies (dashed vertical lines) are indicated where appropriate. The bottom row shows the cost $C(x)$ (dotted line); benefit $B(2x)$ (dashed line); and the mean payoff $B(2x) - C(x)$ (solid line). The dash-dotted vertical line indicates maximal mean payoffs. (A) Social parasitism with evolutionary branching ; (B) Social parasitism with an evolutionarily stable singular strategy; (C) Synergy with two evolutionarily stable strategies. Depending on the initial conditions, the population either evolves to full defection or to full cooperation (two distinct simulations shown); (D) Failed altruism; (E) Mutualism. See Doebli et al. (2004) for details and parameter values.

where q is the level of investment by its partner (Doebli et al. 2004). Figure 10.2 shows that the population will first move towards adopting a single strategy. If this strategy is not pure defect or pure co-operate then, depending on the exact form of $B(p, q)$ and $C(p)$, the population will either branch into some proportion which invest more (co-operators), and some proportion which invest less (defectors), in co-operation than those at the mixed singular strategy (figure 2a). Alternatively, the population will remain at this mixed strategy where all individuals invest the same amount in co-operation (figure 10.2b). For continuous strategy games it is this final outcome which is the evolutionarily stable strategy and cannot be invaded by mutant individuals adopting small changes to their strategy.

The main biological message of the continuous strategy game is that even when individuals have an opportunity to tune a level of investment in a co-operative activity, we can expect discrete strategies. The co-operators produce most of the benefit and pay most of the cost and the defectors parasitize their efforts

(Doebeli & Hauert 2005; Doebeli et al. 2004). This could explain why some individuals seem to specialise in apparently costly behaviours such as searching for food (Barnard & Sibly 1981) while others apparently reap only the benefits of these behaviours. It is important to note, however, that the co-operators in such scenarios still derive a direct benefit from co-operation. By definition, the co-operating individuals cannot increase their payoff by defecting and are thus acting in their own selfish interests.

10.3 Mutualism

If there are positive benefits to performing a co-operative behaviour independent of the actions of a partner then co-operation can evolve despite the existence of associated costs. If in table 10.1 there is an extra benefit that cancels out the cost, i.e. $E=C$, the payoff for both co-operating is $B+E-C=B$. Now there is no longer a positive incentive to defect. If E continues to increase then the total benefit becomes more than the sum of the separate actions and there are now mutual benefits to both co-operating, i.e. $B+E-C > B$. The extra benefit of working together outweighs any cost. In figure 10.1, I label this scenario mutualism. Both E and D are greater than C and under all circumstances it pays to co-operate.

Mutualism provides an explanation of many different behaviours by animal groups: animals aggregate in the hope that another individual in a group will be eaten by a predator (chapter 2), pigeons flying together benefit from each other's directional information (chapter 5), and so on. In many ways, mutualisms provide a null hypothesis for co-operative behaviour. If we see an animal performing a costly behaviour that benefits another individual, the first question we can ask is what benefits it gains itself from the action. If it gains irrespective of the actions of its partner then the interaction is mutualistic.

Despite providing a somewhat obvious reason for individuals to co-operate, mutualisms are sometimes overlooked. This can be because the benefits are not immediately clear or the costs are overestimated. Clutton-Brock (2002) emphasises that, when estimating these costs and benefits, the physiological state of the individual performing them must be taken into account. This state can differ between individuals. Thus while co-operative actions may have substantial energetic costs, they may be performed by individuals with high energy reserves. For example, as discussed in chapter 6, meerkats spend more time guarding their collective nest when they are well fed. While guarding may be costly to a hungry meerkat, it is beneficial to one with a full stomach (Clutton-

Brock et al. 2001; Clutton-Brock et al. 1999). It thus becomes mutually beneficial for meerkats to guard their nest and they do so irrespective of the guarding behaviour of other individuals (Bednekoff 1997; Foster 2004).

In mutualisms, the balance of power is often tilted toward one individual (Beekman et al. 2003). For example, carpenter bee queens usually found nests on their own, but are sometimes usurped through a violent struggle by another queen (Hogendoorn & Velthuis 1999). The usurper then lays all subsequent eggs. Ironically, the stage is now set for mutualism between these one time opponents. The usurper can benefit if the original foundress stays to guard the nest, but must provide an incentive for the foundress to stay. She can do this by not destroying all of the eggs laid by the foundress. While this small concession may provide only a small benefit to the foundress, it is greater than the probable benefit arising from an attempt to establish a new nest (Dunn & Richards 2003). Carpenter bee nest founding provides just one example of a "transactional concession" offered by a dominant individual to one or more sub-dominant. These concessions make it mutually beneficial for the sub-dominants to co-operate (Reeve & Keller 2001). Unlike in table 10.1, the dominant and sub-dominants have different payoff tables and the dominant uses concessions to manipulate the payoff table of the sub-dominants. This manipulation ensures that both have a benefit in co-operating, independent of the future actions of the other. Thus, while power and mutualism might appear unlikely partners in animal conflicts, as in human affairs, they lie at the heart of many co-operative group behaviours.

It is particularly important when considering mutualisms to observe that the benefit B has absolutely no role in the evolutionary dynamics (Box 10.A). This is somewhat counter-intuitive since it means that we can expect individuals to evolve to confer arbitrarily large benefits on others, giving a strong appearance of altruism. What should be borne in mind here is that for a particular act to benefit a partner usually incurs a cost to the actor. For example, giving your food to another individual might benefit them to degree B but will cost you to degree C . If the act of giving food provides no direct or extra benefit (i.e. $D=E=0$) then, however small the cost is of giving food and however large the benefit is to another, it should not evolve.

Despite the above limitation, we can expect to see acts where individuals confer huge benefits to others. In particular, acts which incur costs that are

slightly smaller than the direct benefits they provide, i.e. C is slightly smaller than D , may provide extremely large benefits to others. A human example of such a mutualism would be a confident swimmer jumping in a river to save a stranger's life. The confident swimmer is unlikely to drown (small C) and will reap a benefit in terms of reputation or simply from the vigorous training swim (D slightly bigger than C), but in doing so will provide an extremely large benefit B to the drowning stranger. However happy the drowning stranger is after the rescue, by the evolutionary game theory definition, this is no more altruistic than if the lifesaver had trained for, entered and won a swimming contest.

10.4 Synergism

Much of this book is about how non-linear interactions between individuals can produce patterns that a single individual could not achieve alone. The chemical trails of rats and ants, cliff swallow food calling (chapter 3), bark beetles attacking trees and caterpillars building tents (chapter 7) involve the mass co-operation of many individuals which ultimately improves individual efficiency. In this sense their co-operation could be described simply as mutualistic. However, there is an aspect of these co-operative behaviours that is not true of mutualisms as defined above: if we consider a single individual performing the behaviour in a group consisting solely of defectors it would pay a cost without receiving any benefit. For example, if a single tent caterpillar starts to build a tent and the others save the energy they would have expended producing silk then this focal individual pays all the cost and reaps only its small share of the benefit. If this cost outweighs the benefit, how then can such co-operation evolve?

For two-player interactions, this type of social dilemma can be formalized by setting the direct benefits of a co-operative behaviour to be less than the costs, i.e. $D < C$. If the focal individual knows its partner will co-operate then it is always better to co-operate, since as in mutualism, $B + E - C > B$. However, if the focal knows its partner will defect then, because $D < C$, it is better to also defect and avoid a negative payoff. In figure 10.1 this scenario is called synergism and has two possible evolutionarily stable states: one corresponding to everyone co-operating and another corresponding to everyone defecting. Which evolutionarily stable state the population evolves to depends on the initial conditions. If the population initially contains more than $(C - D) / (E - D)$ co-operators then evolution will lead to full co-operation, otherwise evolution will lead to full defection.

A prediction of the synergism model is that costly behaviour can evolve even if, when interacting with other individuals that defect, the focal individual gets no benefit from co-operating. This point is not always given full consideration when discussing the evolution of costly signals. For example, cliff swallows call to signal the location of insect swarms thus paying a, probably small, cost in lost food but providing nearby foraging partners a positive benefit in finding food. Brown et al. (1991) suggest, quite correctly, that swallows may have evolved call signalling because “even if other birds do not also call, the caller could benefit through local enhancement simply by watching the nearby group members as some of them track the subsequent movement of the prey”. If this is the case, then there may be no cost to interacting with a defector, i.e. $D > C$ and $E > C$, and full co-operation always evolves through mutualism. However, a benefit of local enhancement to signalling is not a requirement for the evolution of food calling. Rather, the game theory model predicts that provided there is an extra benefit when both birds call that is greater than the cost of calling then co-operation can evolve, independent of any direct benefits in the absence of calling, i.e. $D < C$. It is plausible that such extra benefits exist for cliff swallows. Groups that contain individuals that always signal can continuously track the movement of insect swarms. When interacting with a co-operator, the focal individual gets the additional benefit, $E > C$, of being able to re-find its own discovery. Defection would reduce both the focal and the partner bird’s ability to find food.

More than the sum of its parts

The last paragraph takes a two-player game and suggests it may apply to multi-player interactions. Swallows don’t forage in pairs but rather in large groups. Under what circumstances can synergistic co-operation persist in larger groups? Box 10.C describes a continuous strategy game with group size N in which each individual can make an investment p_i in co-operation. This investment incurs a constant cost, but group productivity increases as a function of the co-operative investment. This productivity is shared equally between individuals, so benefit to an individual increases as productivity divided by number of group members. Co-operation is evolutionarily stable for large groups in this model provided that group productivity increases with at least the square of the group size, or equivalently provided that the benefit per individual increases at least linearly with group size.

Figures 10.3 and 10.4 show how the evolutionarily stable states change with group size for two different productivity functions. When productivity grows superlinearly with group size then maximum co-operation is always a potential evolutionary outcome (figure 10.3). When productivity grows first superlinearly, but then saturates then maximum co-operation is stable for an intermediate range of group sizes (figure 10.4). This last observation is important, since realistic productivity functions must saturate at some point. Group productivity cannot continue to increase indefinitely with group size since at some point members of the large group must compete for resources (Foster 2004). Despite these diminishing returns for very large groups, co-operation is stable in intermediate-sized groups. This observation is true generally for productivity functions which increase super-linearly with group size at first and then saturate later (Sumpter & Brännström 2008).




Figure 10.3: Model of synergy described in Box 10.B with productivity that increases with the cube of group size, i.e. $g(P) = bP^3$. Bifurcation plots showing the location and stability of interior singular strategies and boundary points as a function of group size N . We choose $c/b=100$ so that for very small groups there is no benefit to co-operation, i.e. $p=0$ is the only stable strategy. With increasing group size a repelling interior singular strategy emerges and both no investment ($p=0$) and maximal investment ($p=1$) are locally stable strategies. Arrows indicate for which initial investment in co-operation these strategies will evolve (see (Sumpter & Brännström 2008) for details).

Examples discussed in earlier chapters such as cliff swallow foraging, tent building by caterpillars and ant foraging involve either a group productivity that increases super-linearly with group size or, equivalently, a benefit per individual increasing linearly with group size. A key property of many synergistic interactions is the use of signals to spread information. Cliff sparrows use vocal signals (Brown et al. 1991); Norway rats (Galef & Buckley 1996), naked mole rats (Judd & Sherman 1996) and ants deposit residual trails; social insects use an array of different type of dances and other signals. The evolution of these signals is intimately linked with positive feedback. Signalling by a focal individual improves other group members' chances of discovering food and since these group members are also signallers then this improves the chance of rediscovering the same food or finding other nearby sources. The positive feedback continues and group productivity increases as more than the sum of the group's parts (see chapter 3). Such

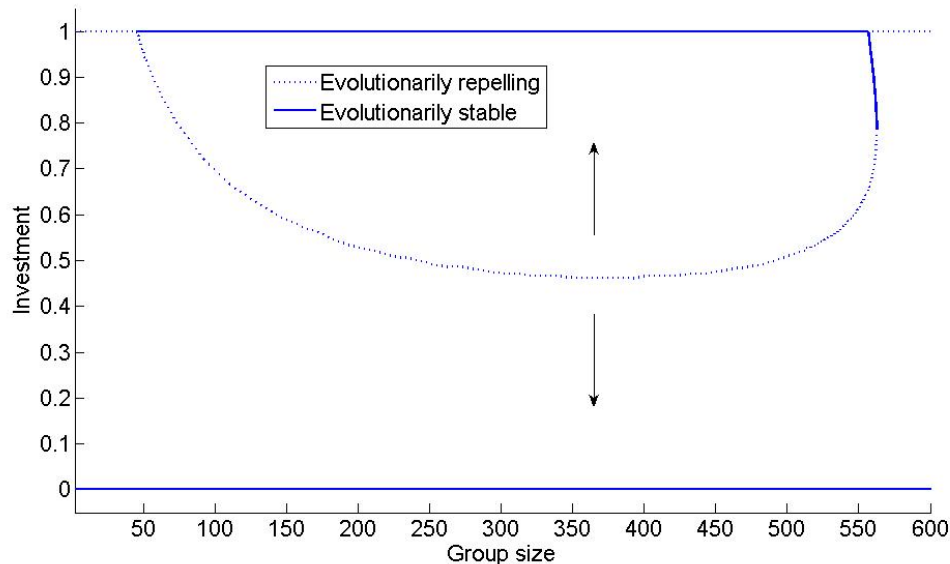


Figure 10.4: Model of synergy described in Box 10.B with group productivity that first increases with the cube of group size but later saturates to linear increase, i.e

$$g(P) = bT^2 \frac{P^3}{T^2 + P^2} .$$

Parameters are $c/b=5$ and $T=40$. Bifurcation plot showing the

location and stability of interior singular strategies and boundary points as a function of group size N . As in figure 10.3 with increasing group size a repelling interior singular strategy emerges and both no investment ($p=0$) and maximal

investment ($p=1$) are locally stable strategies. In this case however as group sizes increases further $p=1$ becomes unstable and a strategy corresponding to an intermediate investment in communication becomes stable. As group size increases still further the intermediate investment strategy disappears and $p=0$ is the only stable state. The arrows from points indicate for which initial investment in co-operation the various stable strategies will evolve (see (Sumpter & Brännström 2008) for details).

group productivity is precisely that needed for co-operation to evolve in the model in Box 10.C.

Positive feedback and self-organisation are sometimes proposed as alternatives to natural selection in explaining the evolution of co-operation. The argument is that because self-organised systems are more than the sums of their parts they cannot be understood simply in terms of the selfish individual units of which they are composed. By

Box 10.C Synergy in groups of size N.

Consider a population that on each generation randomly aggregates in isolated groups of size N . Each individual can choose to invest an amount $p_i \in [0,1]$ in a co-operative behaviour. The benefit to each individual, $g(\sum_{j=1}^N p_j)/N$, is assumed to be a function of the overall productivity of the group members, g , divided by the total number of group members. Assume that this function is the same for all group members. Thus the payoff for an individual i is

$$\frac{g(N, \sum_{j=1}^N p_j)}{N} - p_i c$$

where c is the cost of the co-operative behaviour. This model is an example of a structured-deme model (Nunney 1985; Wilson 1983).

Let's start by assuming that productivity increases as some power α of the level of co-operation of, i.e. $g(P) = bP^\alpha$. We now follow the method outlined by Doebeli et al. (2004). Assume that all individuals have the same strategy q apart from a mutant with strategy p . The selection gradient is then

$$D(q) = \frac{\partial}{\partial p} \left(\frac{b}{N} (p + (N-1)q)^\alpha - cp \right) \Big|_{p=q} = \frac{b\alpha}{N} (Nq)^{\alpha-1} - c$$

Since we insist that investment is between 0 and 1, we can evaluate the selection gradient at these two extremes in order to see whether they are stable strategies. Evaluating $D(0) = -c$ tells us that the all defect is an evolutionarily stable state. Similarly, $D(1) = b\alpha N^{\alpha-2} - c$ tells us that the all co-operate is also evolutionarily stable, provided $b\alpha N^{\alpha-2} > c$. When all co-operate is stable there exists, although we don't determine it explicitly here, a single steady state q_* between these two extremes which is not convergent stable. This steady state acts as a repellent: when initially $q > q_*$ then $q \rightarrow 1$ and when initially $q < q_*$ then $q \rightarrow 0$. Qualitatively, the situation is the same as in the two player discrete game discussed in the text: both all co-operate and all defect are evolutionarily stable.

The condition for synergistic co-operation in this model is $b\alpha N^{\alpha-2} > c$. If $\alpha < 2$ then as group size increases the cost an individual is willing to pay in co-operating decreases. For example, when $\alpha = 1$ we recover $b/N > c$. If $\alpha \geq 2$, however, then as group size increases the cost an individual is willing to pay tends toward a positive but finite limit. In particular, when $\alpha = 2$, co-operation is stable if $2b > c$ independent of N . Figure 10.3 shows how the steady states change with group size for $\alpha = 3$.

Figure 10.4 shows similar analysis for $g(P) = T^2 \frac{P^3}{T^2 + P^2}$. This productivity function initially grows cubically, but when group size exceeds T the growth becomes more linear. For large P growth is purely linear. Here there are three different parameter regimes. For very small group sizes all individuals evolve to invest nothing ($p=0$) in co-operation, but as group size increases the strategy of full investment ($p=1$) becomes stable. At intermediate group sizes the full investment becomes unstable and a compromise of partial investment becomes stable. As group size increases still further all communication becomes evolutionarily unstable and $p=0$ is the only evolutionarily stable state.

implication, the group becomes the unit upon which selection acts, and we need to consider how evolution functions on multiple levels (Fletcher & Doebeli 2006; Fletcher et al. 2006). The synergistic model I propose here shows that self-organisation is consistent with selection on the level of the selfish individuals. Positive feedback can generate synergistic effects, and synergisms benefit the individual who is part of the group. The selection pressure on the individual is to co-operate, since it benefits from group membership. While the mechanisms whereby synergism is generated are often complex and need disentangling, this does not imply that their evolution cannot be understood at the level of the individual parts.

Synergism depends crucially on frequency dependence: co-operation can persist in an established population of co-operators, but cannot establish itself in an already established group of defectors. This is not always explicitly spelt out when co-operation is discussed. For example, West et al. (2006; 2007) classify social behaviours in a table with effect on actor in rows and effect on recipient in the columns. In particular, West et al. (2007) make a valiant effort to clear up confusion about co-operation and altruism arising from inclusive fitness (see section 10.6 below). However, in their classification they do not differentiate between direct and extra benefits and implicitly assume that $E=D$. Instead, they extend the model in Box 10.A by allowing B to be negative (see West et al. 2007 tables 2 and, for the more general case $N>2$, table 3). Figure 10.5 interprets this classification in terms of the evolutionary game theory model in Box 10.A. In contrast to figure 10.1, B is now on the x-axis instead of D , but B has no effect on evolutionary outcome. Equating E and D means that parasitism and synergism cannot occur, and the only possible evolutionarily stable states are those where either all individuals co-operate or all defect. Conversely, in the model in figure 10.1, when $E>C$ co-operation can evolve without altruism, even when $C>D$ (negative effect of co-operation on actor at low, but not high, frequencies of co-operators) and $B>0$ (positive effect on recipient at all frequencies).

To be fair to West and co-workers, their classification is entirely valid and was designed to make the distinction between direct benefits and altruism. However, it is usually confusion between extra and direct benefits, and not misunderstanding of inclusive fitness, that leads to erroneous claims that 'altruism' can arise from group selection or self-organisation (Queller 1985; Queller & Strassmann 2006). Altruism is never predicted to occur in models with the assumptions I gave at the start of this chapter, while synergism can occur given a high enough initial proportion of co-operators.

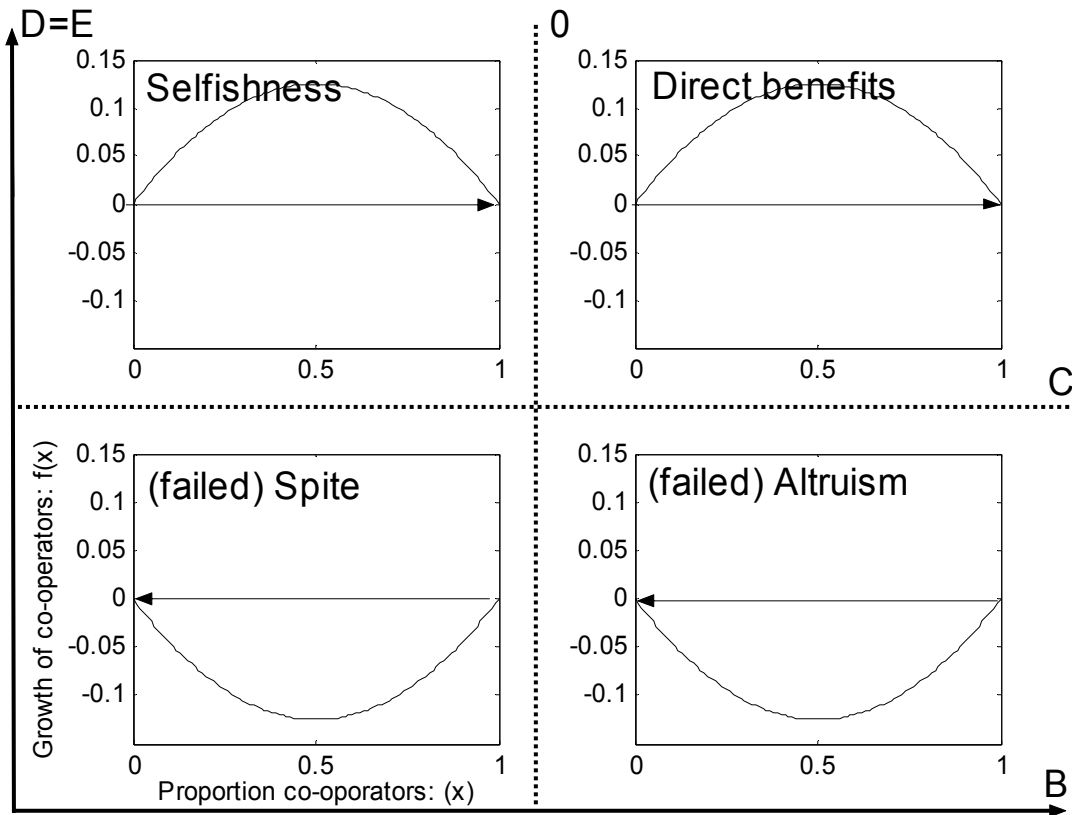


Figure 10.5 Classification of social behaviours for the two player game proposed by West et al. (2006; 2007). Their model assumes that there is a cost to co-operating, $C > 0$, there is a direct benefit to self of co-operating independent of whether the other player defects, $D = E > 0$, and that a co-operator provides a benefit B to another individual. In the model given in table 2 of West et al. (2007) they further assume that $B = D$, but in the general model for types of co-operation given implicitly in table 1 of West et al. (2007) they allow B to have any value, positive or negative. Given these assumptions we can substitute these parameters in to equation 10.A.1 in box 10. **Error! Reference source not found.** The figure then shows how the cost benefit parameters, C , B and D , determine the evolutionary stable strategies. As before, B plays no role in determining the evolutionary stable strategy.

One last point should be made about synergism. The model in Box 10.C depends on a 'fair' division of the payoffs. The fact that per capita productivity increases linearly with the number of individuals does not necessarily mean that the evolution of a particular form of co-operation is due to synergism. A central assumption that must also be tested is whether the benefits are on average shared equally between group members. If one individual has an opportunity to take consistently more than its share at the expense of the others then co-operation can fail. Unequal shares clearly occur in social insect colonies, where it is the small number of queens who produce the majority of offspring. In this case

we need the additional explanation, in the form of increased inclusive fitness or some form of transactional concession to maintain co-operation.

10.5 Repeated interactions

When there is a positive probability of interacting with the same individual again it can be beneficial to an individual to adopt a strategy that allows for co-operation. This observation is formalized in what is known as the iterated prisoner's dilemma (Axelrod & Hamilton 1981; Trivers 1971). In the iterated prisoner's dilemma, each individual plays a game a fixed number of times with payoffs as in table 10.1 with $D=E=0$ and C and B positive. Over a wide range of conditions, a strategy known as Tit-for-Tat is evolutionarily stable for this game. A focal Tit-for-Tat individual co-operates on the first interaction and then on the next interaction adopts the same strategy as its partner did on the previous interaction. So if the partner defects, so too does the focal Tit-For-Tat individual. If two Tit-for-Tat individuals meet each other then they always co-operate. This interaction allows co-operation to evolve under the 'threat' that defection by one individual will result in a break-off of co-operation and both individuals losing out.

There are many extensions to the iterated prisoner's dilemma, but in order for individuals to co-operate in these there must exist either a positive probability of interacting with the same individual (e.g. Killingback & Doebeli 2002; Schuessler 1989; e.g. Wahl & Nowak 1999) or indirect knowledge of a partner's previous strategic choices (Nowak & Sigmund 1998). Lehmann & Keller (2006a) provide a strong argument that, for all such models, a Tit-for-Tat like strategy is evolutionarily stable provided that

$$mB > C$$

where m is the probability that a further interaction will occur with the same individual multiplied by the probability that the focal individual knows the result of the partner's last interaction. For this case, the payoff table for Tit-for-Tat playing against an Always Defect strategy is given in table 10.2. These payoffs are the same as those for synergistic interactions as in table 10.1, with

$$E = \frac{m}{1-m}(B - C) \text{ and } D = 0. \text{ The iterated prisoner's dilemma thus leads to either}$$

everyone being Tit-for-Tat or everyone being Always Defect, depending on the initial proportion of each in the population. A more complete description of strategies in the repeated prisoner's dilemma, including analysis of other potential alternative strategies, can be found in Hofbauer & Sigmund (1998).

While there is some evidence that previous interactions and reputations are important in human interactions (Fehr & Fischbacher 2003; Fehr & Fischbacher 2004; Skryms 2004) and in low cost activities by animals such as grooming and feeding, these have not been convincingly demonstrated in high cost co-operative breeding (Clutton-Brock 2002). Furthermore, unless there are strongly preferential interactions within groups, we would expect that $m \propto 1/N$ and as group size increases the chance of repeated interactions will decrease. Unlike synergy, co-operation through repeated interactions relies on either small group sizes or a large capacity for remembering previous interactions. These observations have led several authors to question the general significance of repeated interactions alone in explaining the evolution of co-operation in many of the mass collective behaviors of animals (Clutton-Brock 2002; Richner & Heeb 1996). Indeed, repeated interactions do not play a significant role in explaining most of the collective behaviours I have discussed in this book. This is simply because the most interesting collective behaviours involve large numbers of individuals.

There are however various ways in which repeated interactions can enhance synergism or mutualism. For example, individuals with a choice of which group to join could first interact with group members to ascertain their behavioral strategy and then choose to join groups of co-operators rather than defectors. If payoff increases with group size then attracting other group members would provide an extra incentive for group members to co-operate with potential joiners (Kokko et al. 2002; Wilson & Dugatkin 1997). It is however difficult to envisage how such assortative interactions can lead to co-operation without the possibility of further repeated interactions or some form of synergism, mutualism or parasitism. For example, let's assume co-operators can evolve a mechanism for identifying other co-operators and preferentially interact with them. But then what is to stop defectors evolving the same mechanism and exploiting co-operators? The only answer is to insist that the mechanism for identifying the co-operative feature is only found in those individuals that carry the feature (Dawkins 1976). While this remains a possibility, it is not a particularly general explanation of co-operation (Lehmann & Keller 2006a; West et al. 2006).

10.6 Inclusive fitness

Altruism is a behaviour that increases the direct fitness, or average lifetime reproductive success, of another individual while decreasing the direct fitness of the actor (West et al. 2006). In the models discussed so far, direct fitness is

proportional to the payoff or fecundity of an individual minus the average payoff of the population (Box 10.A). This definition automatically excludes evolutionarily stable strategies from being altruistic. If an individual can increase its direct fitness by changing strategy then a state is not evolutionarily stable. Altruism cannot evolve in such a setting (figure 10.1).

Hamilton's rule

Hamilton proposed a simple yet far-reaching rule for the evolution of altruism. Assume that p and q are the respective behavioural strategies of a focal individual and its partner. Larger values of p and q can be thought of as corresponding to more investment in a co-operative behaviour. Hamilton argued that if r is the co-efficient for relatedness between the focal individual and its partner, $b(p,q)$ is the total increase in direct fitness conferred on the partner by the focal individual and $c(p,q)$ is the cost of the behaviour to the actor in terms of direct fitness then the focal individual's investment in helping will increase if the *inclusive fitness* of the focal individual

$$rb(p,q) - c(p,q)$$

is greater than zero. Inclusive fitness includes both the benefits of interaction that an individual gains directly in its own fitness plus those it gains indirectly through increased fitness of its relatives. The key idea underlying inclusive fitness is that, because related individuals have a higher probability of sharing the same genotype, genotypes with higher inclusive fitness are more likely to be transmitted to the next generation.

By interpreting p and q as the focal individual's and its partner's investment in co-operation, we can calculate the benefits and costs for interactions according to the payoffs in table 10.1. Using a method proposed by Taylor and Frank (1996) we find

$$b(p,q) = pE - pD + B \text{ and } -c(p,q) = qE + (1-q)D - C$$

for evolutionary game theory models. Assuming that $D=E=0$, then Hamilton's rule for an increase in helping is simply

$$rB > C \quad (10.1)$$

If this inequality holds, a focal individual's probability of helping p will tend to 1, independent of the strategy of its partner. From the viewpoint of direct fitness this is altruism. When everyone co-operates, $q=1$, a focal individual could increase its

direct fitness by defecting. In doing so however it would reduce its inclusive fitness. Hamilton's insight was that in Darwinian natural selection, where the gene is the unit on which selection acts, it is inclusive fitness and not payoff, fecundity or direct fitness which is maximised.

The relationship in equation 10.1 is a very specific version of Hamilton's rule. In the rule's more general form, the benefit $b(p,q)$ and cost $c(p,q)$ are functions of the strategies of the focal individual and the average strategy in the population. For example, if we relax the assumption that $D=E=0$ then Hamilton's rule is

$$r(pE - pD + B) > qE - (1 - q)D - C$$

and a whole range of evolutionarily stable states for interactions between relatives arises. The rule now depends on the frequency of co-operators in the population. In general, we see that relatedness between individuals leads to an increase in co-operation: parameter combinations that would not have led to co-operation now move towards parasitism and synergy; and instances of parasitism and synergy become mutualisms (see table 10.3 for a summary of the effect of inclusive fitness).

Hamilton's rule has proven extremely successful in making predictions about co-operation and conflict in animal societies where between individual relatedness is positive (Griffin & West 2003). Helping relatives at a direct cost to oneself is a feature of all levels of biological organisation: from bacteria to humans (Keller 1999). One of its most successful and intellectually interesting applications has been in understanding the evolution of insect societies (Bourke & Franks 1995; Queller & Strassmann 1998; Ratnieks et al. 2006). At the most basic level, the division of the colony into queens and workers is evidence of altruism. Non-reproducing workers help their mothers raise sisters who become queens and brothers who mate with queens from other colonies, thus passing their genes on to the next generation. A consistent failure to reproduce by a class of individuals is something that simply cannot evolve through natural selection in the absence of the gains in inclusive fitness implied by Hamilton's rule. Once the worker class is established, natural selection acts to increase the efficiency of the workers' interactions as colonies compete to produce the most queens and drones. It is here that the sophisticated forms of co-operation seen in honey bees, ants and termites are seen to evolve (chapter 9).

Although the members of insect societies are usually related, the relatedness within these groups are often lower than predicted by their supposed family structure (Heinze et al. 2001; Korb & Heinze 2004). These observations have even led some to question whether relatedness has any importance at all in explaining co-operation (Costa 2006; Wilson & Holldobler 2005). At an extreme, colonies of unicolonial ant species, such as the Argentine ant, contain many unrelated queens (Pedersen et al. 2006). The question is how co-operation persists with such low levels or zero relatedness between group members? One explanation is a combination of synergism and altruism through indirect fitness benefits. Workers in insect societies gain inclusive fitness from their mother queens who in turn, even if unrelated to other queens, have an incentive to produce more workers and increase group efficiency. Greater insight into these questions could be gained by determining how colony performance increases with group size. Most importantly, synergistic interactions and helping of relatives should not be viewed as distinct explanations of co-operation, but different parts of a common explanation (Foster et al. 2006; West et al. 2007).

Family groups and spatially structured populations

The correlation between relatedness and co-operation is far from the complete story of Hamilton's rule or the application of inclusive fitness theory. Depending on assumed lifecycle of individuals the predictions of inclusive fitness theory change. The derivation of equation 10.1 is based on the assumptions I made at the start of the chapter about the lifecycle of individuals. A question then arises as to where the between individual relatedness originates from in such a model? At the start of this chapter, I assumed that individuals disperse widely and interact with others chosen randomly from the whole population. In the previous section I assumed that they interact more often than not with relatives. If these assumptions are not to contradict each other they need some clarification.

One way in which interactions can take place between relatives while simultaneously allowing global dispersal is to assume that groups initially consist of a single foundress, which then reproduces to found a 'family' group of size N . We can then consider the interactions between the offspring of the foundress, who are positively related, and determine how much they should invest in co-operation. This approach can be further generalised by having n foundresses, or foundresses who have mated with n different males, where n is a lot smaller than the group size of offspring N . Multiple foundresses will reduce average relatedness in proportion to $1/n$ but relatedness will still have a positive effect on

co-operation. One or a small numbers of foundresses building a larger nest of related offspring is common to the lifecycle of highly co-operative insect societies (Bourke & Franks 1995) and of birds species (Emlen 1997; Komdeur & Hatchwell 1999). Social insects (Downs & Ratnieks 1999; Wilson 1971) and birds (Sharp et al. 2005) have evolved mechanisms that ensure that they interact primarily with relatives. It is often this family structure which leads to altruistic interactions in these groups.

Without the introduction of family-based population structure, global dispersal means that within group relatedness decreases to zero. On the other hand, local dispersal of individuals will always lead to some level of positive relatedness between neighbours. For example, assume that each group is formed on an island. At the end of each generation all the adults die and produce offspring, some of which remain on the island and others of which disperse to a randomly chosen island. There is a positive probability the individuals remaining on the same island share the same mother. This assumption is the basis of a mathematical model of local dispersal known as the continuous island model (Rousset 2004). Another way to introduce local dispersal or spatial population structure is to consider a ring of islands, each linked to one nearest neighbour. Dispersal then occurs between neighbouring islands, and nearby individuals become more related than those living at far away islands. This spatial structure is an example of a stepping stone model (Kimura & Weiss 1964; Rousset 2004; Wright 1943).

Although local dispersal increases relatedness, it does not necessarily increase co-operation. Taylor (1992a; 1992b) showed that in both the island model and the stepping stone models, despite relatedness between individuals, the predicted level of co-operation is entirely independent of migration. This is because as migration decreases, not only does relatedness increase but so too does competition between relatives. In terms of Hamilton's rule, migration effects not only r but also $c(p,q)$ and $b(p,q)$, and it effects them in such a way that the inclusive fitness $r b(p,q) - c(p,q)$ remains constant (Rousset, 2004). These theoretical ideas have been tested experimentally in fig wasps, where fighting between males was uncorrelated with their relatedness, but was negatively correlated with the level of competition they faced for future mating opportunities (West et al. 2001). In general, Taylor's result has deep implications for how relatedness is used to predict altruism. It shows that Hamilton's rule does not immediately imply that co-operation increases with relatedness.

While local migration alone does not necessarily promote co-operation, other details of the population's lifecycle do. Inclusive fitness theory predicts the conditions for co-operation over a wide range of assumptions about the lifecycle, including details of overlapping generations; different forms of migration and niche construction (Irwin & Taylor 2000; Irwin & Taylor 2001; Lehmann 2007; Taylor & Irwin 2000). In these models the inclusive fitness provides the condition under which co-operation can evolve (Grafen 1985; Lehmann & Keller 2006a; Rousset 2004). This breadth of application of inclusive fitness is sometimes underappreciated, probably because the simplicity of the inclusive fitness equation is deceiving. For any lifecycle the components r , $b(p,q)$ and $c(p,q)$ must all be calculated and, although there are extensive tools from population genetics that allows this to be done (Rousset 2004), it is not always straightforward to apply them.

A failure to see the proper connection to inclusive fitness has led to a confusing literature on how spatially local interactions can promote altruism (Lehmann & Keller 2006a; Lehmann & Keller 2006b). Various simulation models, collectively known as evolutionary graph theory, have been proposed that show how local interactions can lead to altruism (Hauert 2002; Lieberman et al. 2005; Nowak & May 1992; Ohtsuki et al. 2006; Santos & Pacheco 2006). In these models, individuals play a prisoner's dilemma-like (i.e. D and E both less than C) game with local neighbours on a graph or social network. Depending on the form of the network and the details of the interaction rules, individuals evolve to have a greater or lesser tendency to 'co-operate'. These models do not usually follow the inclusive fitness approach of calculating relatedness between individuals. This proves problematic, since local migration always implies a positive relatedness. For example, Lehmann et al. (2007) showed that some basic results from evolutionary graph theory are special cases of the application of inclusive fitness theory. Altruism in many of these models arises simply from increased interactions with relatives, and not because spatially local interactions are a novel route to co-operation.

Although inclusive fitness theory provides some powerful tools for disentangling the evolution of co-operation in spatially structured populations, it is important to recognize that conditions promoting co-operation in particular biological and social systems cannot be derived without a good understanding of the details of the systems' life cycle (Ratnieks 2006). Models of the evolution of co-operation

are often based on assumptions that adults produce large numbers of offspring, juveniles disperse independently of each other and that there are the same number of individuals in all groups. Changes to these assumptions lead to changes in the underlying population structure and changes in predictions about the direction of selection. A multitude of different theoretical studies have investigated how co-operation evolves under different assumptions about population structure (e.g., Taylor, 1992a,b; Taylor and Frank, 1996; Frank, 1998; Taylor and Irwin, 2000; Irwin and Taylor, 2001; Reuter and Keller, 2001; van Ballen & Rand, 1998; Le Galliard et al., 2003; Roze and Rousset, 2004; Lehmann et al., 2006). What this rich variety of theoretical studies reveals is that there is no simple mathematical formula for the evolution of altruism, but that the inclusive fitness framework provides a powerful toolset to disentangle the factors that promote co-operation.

10.7 Unifying explanations?

In the literature on social or group behaviour there is no clear agreement on the exact definition of terms such as co-operation, social behaviour, helping, mutual benefit, direct benefit, group selection, or even altruism. Different words have been defined by different authors to mean the same thing, and the same words have been defined to mean different things (West et al. 2006; West et al. 2007). An advantage of classifying co-operation as the outcome of different parameter values in a single model, as I have done here, is that the assumptions are clearly stated and the predictions are seen to be logical consequences of the model. I have of course chosen words to label different sets of assumptions and predictions and in doing so I chose mostly to follow (Clutton-Brock 2002), although instead of using his label of 'group augmentation' (Kokko et al. 2001) I use Maynard Smith & Szathmáry's (1995) label synergism. However, labelling is not important. While these labels can change, the outcome of these two-player, two-strategy games will not. Given the same set of assumptions these models will make the same predictions.

In addition to analysing two-player, two-strategy games, I have argued that many more complex game theoretic and inclusive fitness models can be classified under the same broad headings—parasitism, mutualism, synergism, repeated interactions and altruism—that I have given the simpler models. There remain a large number of detailed models of animal co-operation that I have not covered here. These introduce a wide variety of concepts such as rent; punishment; partner choice; diminishing returns; etc. Each of these models is

useful in describing particular situations and reveals how details of lifecycle produce changes in co-operation. I would however argue that in a broad sense they can be usefully classified as belonging to one or, depending on parameter values, a combination of the five headings I use above. Such a classification allows comparison between different systems. For example, we can say things like “starling foraging and vigilance are both examples of social parasitism, while cliff swallow foraging appears to be synergistic”. We may have built detailed game theoretic models that capture the details of each of these systems but the classification allows us to quickly summarise what we have found out about these systems.

There is one last distinction that is worth drawing in the study of collective animal behaviour: that is between behaviours which are strictly co-operative and those which are co-ordinating. The distinction is based on whether the cost to the focal individual has evolved or is simply a by-product of the individual undertaking an activity that is otherwise beneficial to itself (Clutton-Brock 2002). This is the same distinction that is made between cues and signals (Maynard Smith & Harper 2005). Cues can co-ordinate activities, but they have not evolved as a result of a benefit to the individuals which make them (chapter 3). Signals will only evolve when animals co-operate, and co-ordination always relies on cues. The distinction makes it straightforward to, for example, rule out social parasitism as an explanation of group foraging involving signals. The distinction is however less straightforward for mutualisms, where it is often difficult to distinguish whether a behaviour has evolved to provide mutual benefit or whether the benefit was simply serendipitous.

There is richness in biological interactions that can never be captured by broad classifications. Indeed, simply answering questions such as how do birds flock, how do fish avoid predators, how do ants build their foraging networks, and how do cockroaches choose where to live with respect to their position on figure 10.1 does not tell us much about these fascinating problems. What makes these problems interesting usually has little to do with a broad notion of why they evolved. So while I have ended this book with an explanation of the functional or ‘why’ questions in relation to collective animal behaviour, the intention is not to finally unify the earlier chapters. Rather the aim of this chapter has been to clarify how functional reasoning is applied to these systems. The hope is that, together with the mechanistic description from earlier in the book, we can use the combination of these approaches to build up a full understanding of collective animal behaviour.