How populations cohere: five rules for cooperation

Martin A. Nowak and Karl Sigmund

Subsequent chapters in this volume deal with populations as dynamic entities in time and space. Populations are, of course, made up of individuals, and the parameters which characterize aggregate behavior—population growth rate and so on—ultimately derive from the behavioral ecology and life-history strategies of these constituent individuals. In evolutionary terms, the properties of populations can only be understood in terms of individuals, which comes down to studying how life-history choices (and consequent gene-frequency distributions) are shaped by environmental forces.

Many important aspects of group behavior—from alarm calls of birds and mammals to the complex institutions that have enabled human societies to flourish—pose problems of how cooperative behavior can evolve and be maintained. The puzzle was emphasized by Darwin, and remains the subject of active research today.

In this book, we leave the large subject of individual organisms' behavioral ecology and life-history choices to texts in that field (e.g. Krebs and Davies, 1997). Instead, we lead with a survey of work, much of it very recent, on five different kinds of mechanism whereby cooperative behavior may be maintained in a population, despite the inherent difficulty that cheats may prosper by enjoying the benefits of cooperation without paying the associated costs.

Cooperation means that a donor pays a cost, c, for a recipient to get a benefit, b. In evolutionary biology, cost and benefit are measured in terms of fitness. While mutation and selection represent the main forces of evolutionary dynamics, cooperation is a fundamental principle that is required for every level of biological organization. Individual cells rely on cooperation among their components. Multicellular organisms exist because of cooperation among their cells. Social insects are masters of cooperation. Most aspects of human society are based on mechanisms that promote cooperation. Whenever evolution constructs something entirely new (such as multicellularity or human language), cooperation is needed. Evolutionary construction is based on cooperation.

The five rules for cooperation which we examine in this chapter are: kin selection, direct reciprocity, indirect reciprocity, graph selection, and group selection. Each of these can promote cooperation if specific conditions are fulfilled.

2.1 Kin selection

The heated conversation took place in an unheated British pub over some pints of warm bitter. Suddenly J.B.S. Haldane remarked, 'I will jump into the river to save two brothers or eight cousins.' The founding father of population genetics and dedicated communist in his spare time never bothered to write up this insight. The witness of the revelation was Haldane's eager pupil, the young John Maynard Smith. But given John's high regard for entertaining stories and good beer, can we trust his memory?

The insight that Haldane might have had in the pub was precisely formulated by William Hamilton. He wrote a PhD thesis on this topic, submitted a long paper to the *Journal of Theoretical Biology*, and spent much of the next decade in the Brazilian

jungle. This was arguably the most important paper in evolutionary biology in the second half of the twentieth century (Hamilton, 1964a, 1964b, 1998). The theory was termed kin selection by Maynard Smith (1964). The crucial equation is the following. Cooperation among relatives can be favored by natural selection if the coefficient of genetic relatedness, r, between the donor and the recipient exceeds the cost/benefit ratio of the altruistic act:

$$r > c/b \tag{2.1}$$

Kin-selection theory has been tested in numerous experimental studies. Indeed, many cooperative acts among animals occur between close kin (Frank, 1998; Hamilton, 1998). The exact relationship between kin selection and other mechanisms such as group selection and spatial reciprocity, however, remains unclear. A recent study even suggests that much of cooperation in social insects is due to group selection rather than kin selection (Wilson and Hölldobler, 2005). Note that kin selection is more likely to work in quite small groups; in large groups, unless highly inbred, the average value of r will be tiny.

2.2 Direct reciprocity

In 1971, Robert Trivers published a landmark paper entitled 'The evolution of reciprocal altruism' (Trivers, 1971). Trivers analyzed the question how natural selection could lead to cooperation between unrelated individuals. He discusses three biological examples: cleaning symbiosis in fish, warning calls in birds, and human interactions. Trivers cites Luce and Raiffa (1957) and Rapoport and Chammah (1965) for the Prisoner's Dilemma, which is a game where two players have the option to cooperate or to defect. If both cooperate they receive the reward, R. If both defect they receive the punishment, P. If one cooperates and the other defects, then the cooperator receives the sucker's payoff, S, while the defector receives the temptation, T. The Prisoner's Dilemma is defined by the ranking $T > R > P > S$.

Would you cooperate or defect? Assuming the other person will cooperate it is better to defect, because $T > R$. Assuming the other person will defect it is also better to defect, because $P > S$. Hence, no matter what the other person will do it is best to defect. If both players analyze the game in this rational way then they will end up defecting. The dilemma is that they both could have received a higher payoff if they had chosen to cooperate. But cooperation is irrational.

We can also imagine a population of cooperators and defectors and assume that the payoff for each player is determined by many random interactions with others. Let x denote the frequency of cooperators and $1 - x$ the frequency of defectors. The expected payoff for a cooperator is $f_C = Rx + S(1 - x)$. The expected payoff for a defector is $f_D = Tx + P(1 - x)$. Therefore, for any x, defectors have a higher payoff than cooperators. In evolutionary game theory, payoff is interpreted as fitness. Successful strategies reproduce faster and outcompete less successful ones. Reproduction can be cultural or genetic. In the non-repeated Prisoner's Dilemma, in a well-mixed population, defectors will outcompete cooperators. Natural selection favors defectors.

Cooperation becomes an option if the game is repeated. Suppose there are m rounds. Let us compare two strategies, always defect (ALLD), and GRIM, which cooperates on the first move, then cooperates as long as the opponent cooperates, but permanently switches to defection if the opponent defects once. The expected payoff for GRIM versus GRIM is nR. The expected payoff for ALLD versus GRIM is $T + (m - 1)P$. If $nR > T + (m - 1)P$ then ALLD cannot spread in a GRIM population when rare. This is an argument of evolutionary stability. Interestingly, Trivers (1971) quotes 'Hamilton (pers. commun.)' for this idea.

A small problem with the above analysis is that given a known number of rounds it is best to defect in the last round and by backwards induction it is also best to defect in the penultimate round and so on. Therefore, it is more natural to consider a repeated game with a probability w of having another round. In this case, the expected number of rounds is $1/(1 - w)$, and GRIM is stable against invasion by ALLD provided $w > (T - R)/(T - P)$.

We can also formulate the Prisoner's Dilemma as follows. The cooperator helps at a cost, c, and

the other individual receives a benefit, b. Defectors do not help. Therefore we have $T = b$, $R = b - c$, $P = 0$, and $S = -c$. The family of games that is described by the parameters b and c is a subset of all possible Prisoner's Dilemma games as long as $b > c$. For the repeated Prisoner's Dilemma, we find that ALLD cannot invade GRIM if

$$w > c/b \qquad (2.2)$$

The probability of having another round must exceed the cost/benefit ratio of the altruistic act (Axelrod and Hamilton, 1981; Axelrod, 1984). Notice, however, the implicit assumption here that the payoff for future rounds is not discounted (i.e. distant benefits count as much as present ones). In evolutionary reality, this is unlikely. We can address this by incorporating an appropriate discount factor in w (May, 1987), but note, from eqn 2, that this makes cooperation less likely.

Thus, the repeated Prisoner's Dilemma allows cooperation, but the question arises: what is a good strategy for playing this game? This question was posed by the political scientist, Robert Axelrod. In 1979, he decided to conduct a tournament of computer programs playing the repeated Prisoner's Dilemma. He received 14 entries, of which the surprise winner was tit-for-tat (TFT), the simplest of all strategies that were submitted. TFT cooperates in the first move, and then does whatever the opponent did in the previous round. TFT cooperates if you cooperate, TFT defects if you defect. It was submitted by the game theorist Anatol Rapoport (who is also the co-author of the book *Prisoner's Dilemma*; Rapoport and Chammah, 1965). Axelrod analyzed the events of the tournament, published a detailed account and invited people to submit strategies for a second championship. This time he received 63 entries. John Maynard Smith submitted tit-for-two-tats, a variant of TFT which defects only after the opponent has defected twice in a row. Only one person, Rapoport, submitted TFT, and it won again. At this time, TFT was considered to be the undisputed world champion in the heroic world of the repeated Prisoner's Dilemma.

But one weakness became apparent very soon (Molander, 1985). TFT cannot correct mistakes. The tournaments were conducted without strategic noise. In a real world, trembling hands and fuzzy minds cause erroneous moves. If two TFT players interact with each other, a single mistake leads to a long sequence of alternating defection and cooperation. In the long run two TFT players get the same low payoff as two players who flip coins for every move in order to decide whether to cooperate or to defect. Errors destroy TFT.

Our own investigations in this area began after reading a News and Views article in *Nature* where the author made three important points: first, he often leaves university meetings with a renewed appreciation for the problem of how natural selection can favor cooperative acts given that selfish individuals gain from cheating; second, strategies in the repeated Prisoner's Dilemma should not be error-free but subjected to noise; third, evolutionary stability should be tested not against single invaders but against heterogeneous ensembles of invaders (May, 1987). This was the motivation for the following work.

In 1989, we conducted evolutionary tournaments. Instead of inviting experts to submit programs, we asked mutation and selection to explore (some portion of) the strategy space of the repeated Prisoner's Dilemma in the presence of noise. The initial random ensemble of strategies was quickly dominated by ALLD. If the opposition is random, it is best to defect. A large portion of the population began to adopt the ALLD strategy and everything seemed lost. But after some time, a small cluster of players adopted a strategy very close to TFT. If this cluster is sufficiently large, then it can increase in abundance, and the entire population swings from ALLD to TFT. Reciprocity (and therefore cooperation) has emerged. We can show that TFT is the best catalyst for the emergence of cooperation. But TFT's moment of glory was brief and fleeting. In all cases, TFT was rapidly replaced by another strategy. On close inspection, this strategy turned out to be generous tit-for-tat (GTFT), which always cooperates if the opponent has cooperated on the previous move, but sometimes (probabilistically) even cooperates when the opponent has defected. Natural selection had discovered forgiveness (Nowak and Sigmund, 1992).

After many generations, however, GTFT is undermined by unconditional cooperators, ALLC. In a society where everybody is nice (using GTFT), there is almost no need to remember how to retaliate against a defection. A biological trait that is not used is likely to be lost by random drift. Birds that escape to islands without predators lose the ability to fly. Similarly, a GTFT population is softened and turns into an ALLC population.

Once most people play ALLC, there is an open invitation for ALLD to seize power. This is precisely what happens. The evolutionary dynamics run in cycles: from ALLD to TFT to GTFT to ALLC and back to ALLD. These oscillations of cooperative and defective societies are a fundamental part of all our observations regarding the evolution of cooperation. Most models of cooperation show such oscillations. Cooperation is never a final state of evolutionary dynamics. Instead it is always lost to defection after some time and has to be re-established. These oscillations are also reminiscent of alternating episodes of war and peace in human history (Figure 2.1).

A subsequent set of simulations, exploring a larger strategy space, led to a surprise (Nowak and Sigmund, 1993). The fundamental oscillations were interrupted by another strategy which seems to be able to hold its 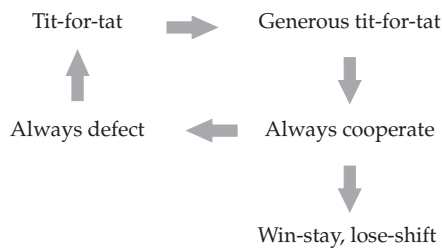ground for a very long period of time. Most surprisingly, this strategy is based on the extremely simple principle of win-stay, lose-shift (WSLS). If my payoff is R or T then I will continue with the same move next round. If I have cooperated then I will cooperate again, if I have defected then I will defect again. If my payoff is only S or P then I will switch to the other move next round. If I have cooperated then I will defect, if I have defected then I will cooperate (Figure 2.2).

If two WSLS strategists play each other, they cooperate most of the time. If a defection occurs accidentally, then in the next move both will defect. Hereafter both will cooperate again. WSLS is a simple deterministic machine to correct stochastic noise. While TFT cannot correct mistakes, both GTFT and WSLS can. But WSLS has an additional ace in its hand. When WSLS plays ALLC it will discover after some time that ALLC does not retaliate. After an accidental defection, WSLS will switch to permanent defection. Therefore, a population of WSLS players does not drift to ALLC. Cooperation based on WSLS is more stable than cooperation based on TFT-like strategies.

Figure 2.1 Evolutionary cycles of cooperation and defection. A small cluster of tit-for-tat (TFT) players or even a lineage starting from a single TFT player in a finite population can invade an always defect (ALLD) population. In fact, TFT is the most efficient catalyst for the first emergence of cooperation in an ALLD population. But in a world of fuzzy minds and trembling hands, TFT is soon replaced by generous tit-for-tat (GTFT), which can re-establish cooperation after occasional mistakes. If everybody uses GTFT, then always cooperate (ALLC) is a neutral variant. Random drift leads to ALLC. An ALLC population invites invasion by ALLD. But ALLC is also dominated by win-stay, lose-shift (WSLS), which leads to more stable cooperation than TFT-like strategies.

Win-stay

C (3) .... **C**   **D** (5) .... **D**

C   **C**

Lose-shift

C (0) .... **D**   **D** (1) .... **C** (probabilistic)

D   **D**

Figure 2.2 Win-stay, lose-shift (WSLS) embodies a very simple principle. If you do well then continue with what you are doing. If you are not doing well, then try something else. Here we consider the Prisoner's Dilemma payoff values $R = 3$, $T = 5$, $P = 1$, and $S = 0$. If both players cooperate, you receive three points, and you continue to cooperate. If you defect against a cooperator, you receive five points, and you continue to defect. But if you cooperate with a defector, you receive no points, and therefore you will switch from cooperation to defection. If, on the other hand, you defect against a defector, you receive one point, and you will switch to cooperation. Your aspiration level is three points. If you get at least three points then you consider it a win and you will stay with your current choice. If you get less than three points, you consider it a loss and you will shift to another move. If $R > (T + P)/2$ (or $b/c > 2$) then WSLS is stable against invasion by ALLD. If this inequality does not hold, then our evolutionary simulations lead to a stochastic variant of WSLS, which cooperates after a DD move only with a certain probability. This stochastic variant of WSLS is then stable against invasion by ALLD.

The repeated Prisoner's Dilemma is mostly known as a story of TFT, but WSLS is a superior strategy in an evolutionary scenario with errors, mutation, and many generations (Fudenberg and Maskin, 1990; Nowak and Sigmund, 1993).

Incidentally, WSLS is stable against invasion by ALLD if $b/c > 2$. If instead $1 < b/c < 2$ then a stochastic variant of WSLS dominates the scene; this strategy cooperates after a mutual defection only with a certain probability. Of course, all strategies of direct reciprocity, such as TFT, GTFT, or WSLS can only lead to the evolution of cooperation if the fundamental inequality (eqn 2) is fulfilled.

2.3 Indirect reciprocity

Whereas direct reciprocity embodies the idea of you scratch my back and I scratch yours, indirect reciprocity suggests that you scratch my back and I scratch someone else's. Why should this work? Presumably I will not get scratched if it becomes known that I scratch nobody. Indirect reciprocity, in this view, is based on reputation (Nowak and Sigmund, 1998a, 1998b, 2005). But why should you care about what I do to a third person?

The main reason why economists and social scientists are interested in indirect reciprocity is because one-shot interactions between anonymous partners in a global market become increasingly frequent and tend to replace the traditional long-lasting associations and long-term interactions between relatives, neighbors, or members of the same village. Again, as for kin selection, it is a question of the size of the group. A substantial part of our life is spent in the company of strangers, and many transactions are no longer face to face. The growth of online auctions and other forms of e-commerce is based, to a considerable degree, on reputation and trust. The possibility to exploit such trust raises what economists call moral hazards. How effective is reputation, especially if information is only partial?

Evolutionary biologists, on the other hand, are interested in the emergence of human societies, which constitutes the last (up to now) of the major transitions in evolution. In contrast to other eusocial species, such as bees, ants, or termites, humans display a large amount of cooperation between non-relatives (Fehr and Fischbacher, 2003). A considerable part of human cooperation is based on moralistic emotions, such as anger directed towards cheaters or the warm inner glow felt after performing an altruistic action. Intriguingly, humans not only feel strongly about interactions that involve them directly, they also judge actions between third parties as evidenced by the contents of gossip. There are numerous experimental studies of indirect reciprocity based on reputation (Wedekind and Milinski, 2000; Milinski et al., 2002; Wedekind and Braithwaite, 2002; Seinen and Schram, 2006).

A simple model of indirect reciprocity (Nowak and Sigmund, 1998a, 1998b) assumes that within a well-mixed population, individuals meet randomly, one in the role of the potential donor, the other as potential recipient. Each individual experiences several rounds of this interaction in both roles, but never with the same partner twice. A player can follow either an unconditional strategy, such as always cooperate or always defect, or a conditional strategy, which discriminates among the potential recipients according to their past interactions. In a simple example, a discriminating donor helps a recipient if her score exceeds a certain threshold. A player's score is 0 at birth, increases whenever that player helps and decreases whenever the player withholds help. Individual-based simulations and direct calculations show that cooperation based on indirect reciprocity can evolve provided the probability, p, of knowing the social score of another person exceeds the cost/benefit ratio of the altruistic act:

$$p > c/b \tag{2.3}$$

The role of genetic relatedness that is crucial for kin selection is replaced by social acquaintanceship. In a fluid population, where most interactions are anonymous and people have no possibility of monitoring the social score of others, indirect reciprocity has no chance. But in a socially viscous population, where people know each other's reputation, cooperation by indirect reciprocity can thrive (Nowak and Sigmund, 1998a).

In a world of binary moral judgements (Nowak and Sigmund, 1998b; Leimar and Hammerstein,

2001; Fishman, 2003; Panchanathan and Boyd, 2003; Brandt and Sigmund, 2004, 2005), there are four ways of assessing donors in terms of first-order assessment: always consider them as good, always consider them as bad, consider them as good if they refuse to give, or consider them as good if they give. Only this last option makes sense. Second-order assessment also depends on the score of the receiver; for example, it can be deemed good to refuse help to a bad person. There are 16 second-order rules. Third-order assessment also depends on the score of the donor; for example, a good person refusing to help a bad person may remain good, but a bad person refusing to help a bad person remains bad. There are 256 third-order assessment rules. We display four of them in Figure 2.3.

With the scoring assessment rule, cooperation, C, always leads to a good reputation, G, whereas

Reputation of donor and recipient

Figure 2.3 Four assessment rules. Assessment rules specify how an observer judges an interaction between a potential donor and a recipient. Here we show four examples of assessment rules in a world of binary reputation, good (G) and bad (B). For scoring, cooperation (C) earns a good reputation and defection (D) earns a bad reputation. Standing is very similar to scoring; the only difference is that a good donor can defect against a bad recipient without losing his good reputation. Note that scoring is associated with costly punishment (Sigmund et al., 2001; Fehr and Gaechter, 2002), whereas for standing punishment of bad recipients is cost-free. For judging it is bad to help a bad recipient. Shunning assigns a bad reputation to any donor who interacts with a bad recipient.

defection, D, always leads to a bad reputation, B. Standing (Sugden, 1986) is like scoring, but it is not bad if a good donor defects against a bad recipient. With judging, in addition, it is bad to cooperate with a bad recipient. For another assessment rule, shunning, all donors who meet a bad recipient become bad, regardless of what action they choose. Shunning strikes us as grossly unfair, but it emerges as the winner in a computer tournament if errors in perception are included and if there are only a few rounds in the game (Takahashi and Mashima, 2003).

An action rule for indirect reciprocity prescribes giving or not giving, depending on the scores of both donor and recipient. For example, you may decide to help if the recipient's score is good or your own score is bad. Such an action might increase your own score and therefore increase the chance of receiving help in the future. There are 16 action rules.

If we view a strategy as the combination of an action rule and an assessment rule, we obtain 4096 strategies. In a remarkable calculation, Ohtsuki and Iwasa (2004, 2005) analyzed all 4096 strategies and proved that only eight of them are evolutionarily stable under certain conditions and lead to cooperation (Figure 2.4).

Both standing and judging belong to the leading eight, but scoring and shunning are not. However, we expect that scoring has a similar role in indirect reciprocity to that of TFT in direct reciprocity. Neither strategy is evolutionarily stable, but their simplicity and their ability to catalyze cooperation in adverse situations constitute their strength. In extended versions of indirect reciprocity, in which donors can sometimes deceive others about the reputation of the recipient, scoring is the foolproof concept of 'I believe what I see'. Scoring judges the action and ignores the stories. There is also experimental evidence that in certain situations humans follow scoring rather than standing (Milinski et al., 2001).

In human evolution, there must have been a tendency to move from the simple cooperation promoted by kin or group selection to the strategic subtleties of direct and indirect reciprocity. Direct reciprocity requires precise recognition of individual people, a memory of the various interactions one had with them in the past, and enough brain
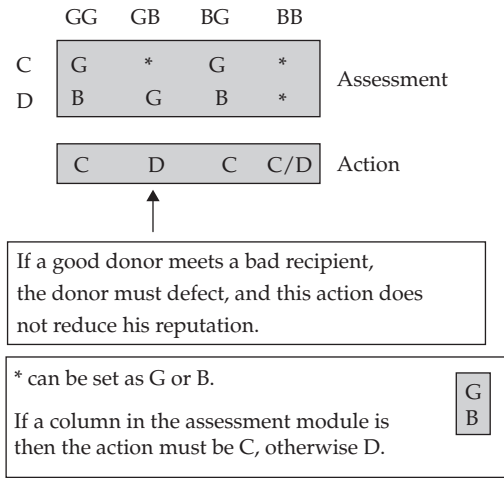
	GG	GB	BG	BB
C	G	*	G	*	Assessment
D	B	G	B	*

| C | D | C | C/D | Action |

If a good donor meets a bad recipient, the donor must defect, and this action does not reduce his reputation.

* can be set as G or B.

If a column in the assessment module is [G / B] then the action must be C, otherwise D.

Figure 2.4 Ohtsuki and Iwasa's leading eight. Ohtsuki and Iwasa (2004, 2005) have analyzed the combination of $2^8 = 256$ assessment modules with $2^4 = 16$ action modules. This is a total of 4096 strategies. They have found that eight of these strategies can be evolutionarily stable and lead to cooperation, provided that everybody agrees on each other's reputation. (In general, uncertainty and incomplete information might lead to private lists of the reputation of others.) The three asterisks in the assessment module indicate a free choice between G and B. There are therefore $2^3 = 8$ different assessment rules which make up the leading eight. The action module is built as follows: if the column in the assessment module is G and B, then the corresponding action is C, otherwise the action is D. Note that standing and judging are members of the leading eight, but that scoring and shunning are not.

Figure 2.5 Games on graphs. The members of a population occupy the vertices of a graph (or social network). The edges denote who interacts with whom. Here we consider the specific example of cooperators, C, competing with defectors, D. A cooperator pays a cost, c, for every link. Each neighbor of a cooperator receives a benefit, b. The payoffs of some individuals are indicated in the figure. The fitness of each individual is a constant, denoting the baseline fitness, plus the payoff of the game. For evolutionary dynamics, we assume that in each round a random player is chosen to die, and the neighbors compete for the empty site proportional to their fitness. A simple rule emerges: if $b/c > k$ then selection favors cooperators over defectors. Here k is the average number of neighbors per individual.

power to conduct multiple repeated games simultaneously. Indirect reciprocity, in addition, requires the individual to monitor interactions among other people, possibly judge the intentions that occur in such interactions, and keep up with the ever-changing social network of the group. Reputation of players may not only be determined by their own actions, but also by their associations with others.

We expect that indirect reciprocity has coevolved with human language. On the one hand, it is helpful to have names for other people and to receive information about how a person is perceived by others. On the other hand a complex language is needed, especially if there are intricate social interactions. The possibilities for games of manipulation, deceit, cooperation, and defection are limitless. It is likely that indirect reciprocity has
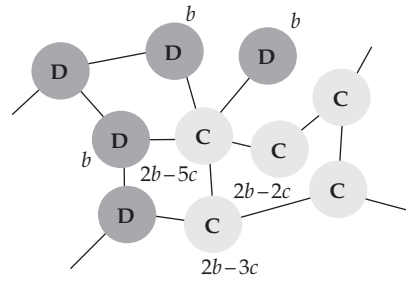
provided the very selective scenario that led to cerebral expansion in human evolution.

2.4 Graph selection

The traditional model of evolutionary game dynamics assumes that populations are well-mixed (Taylor and Jonker, 1978; Hofbauer and Sigmund, 1998). This means that interactions between any two players are equally likely. More realistically, however, the interactions between individuals are governed by spatial effects or social networks. Let us therefore assume that the individuals of a population occupy the vertices of a graph (Nakamaru et al., 1997, 1998; Skyrms and Pemantle, 2000; Abramson and Kuperman, 2001; Ebel and Bornholdt, 2002; Lieberman et al., 2005; Nakamaru and Iwasa, 2005; Santos et al., 2005; Santos and Pacheco, 2005). The edges of the graph determine who interacts with whom (Figure 2.5).

Consider a population of N individuals consisting of cooperators and defectors. A cooperator helps all individuals to whom it is connected, and pays a cost, c. If a cooperator is connected to k other individuals and i of those are cooperators, then its payoff is $bi - ck$. A defector does not provide any help, and therefore has no costs, but it

can receive the benefit from neighboring coopera-tors. If a defector is connected to k other indivi-duals and j of those are cooperators, then its payoff is bj. Evolutionary dynamics are described by an extremely simple stochastic process: at each time step, a random individual adopts the strategy of one of its neighbors proportional to their fitness.

We note that stochastic evolutionary game dynamics in finite populations is sensitive to the intensity of selection. In general, the reproductive success (fitness) of an individual is given by a constant, denoting the baseline fitness, plus the payoff that arises from the game under con-sideration. Strong selection means that the payoff is large compared with the baseline fitness; weak selection means the payoff is small compared with the baseline fitness. It turns out that many inter-esting results can be proven for weak selection, which is an observation also well known in population genetics.

The traditional, well-mixed population of evo-lutionary game theory is represented by the com-plete graph, where all vertices are connected, which means that all individuals interact equally often. In this special situation, cooperators are always opposed by natural selection. This is the fundamental intuition of classical evolutionary game theory. But what happens on other graphs?

We need to calculate the probability, ρ_C, that a single cooperator starting in a random position turns the whole population from defectors into cooperators. If selection neither favors nor opposes cooperation, then this probability is $1/N$, which is the fixation probability of a neutral mutant. If the fixation probability ρ_C is greater than $1/N$, then selection favors the emergence of cooperation. Similarly, we can calculate the fixation probability of defectors, ρ_D. A surprisingly simple rule deter-mines whether selection on graphs favors coopera-tion. If

$$b/c > k \qquad (2.4)$$

then cooperators have a fixation probability of greater than $1/N$ and defectors have a fixation probability of less than $1/N$. Thus, for graph selection to favor cooperation, the benefit/cost ratio of the altruistic act must exceed the average

degree, k, which is given by the average number of links per individual (Ohtsuki *et al.*, 2006). This relationship can be shown with the method of pair-approximation for regular graphs, where all individuals have exactly the same number of neighbors. Regular graphs include cycles, all kinds of spatial lattice, and random regular graphs. Moreover, computer simulations suggest that the rule $b/c > k$ also holds for non-regular graphs such as random graphs and scale-free networks. The rule holds in the limit of weak selection and $k \ll N$. For the complete graph, $k = N$, we always have $\rho_D > 1/N > \rho_C$. Preliminary studies suggest that eqn 2.4 also tends to hold for strong selection. The basic idea is that natural selection on graphs (in structured populations) can favor unconditional cooperation without any need for strategic com-plexity, reputation, or kin selection.

Games on graphs grew out of the earlier tradi-tion of spatial evolutionary game theory (Nowak and May, 1992; Herz, 1994; Killingback and Doebeli, 1996; Mitteldorf and Wilson, 2000; Hauert *et al.*, 2002; Le Galliard *et al.*, 2003; Hauert and Doebeli, 2004; Szabó and Vukov, 2004) and inves-tigations of spatial models in ecology (Durrett and Levin, 1994a, 1994b; Hassell *et al.*, 1994; Tilman and Kareiva, 1997; Neuhauser, 2001) and spatial mod-els in population genetics (Wright, 1931; Fisher and Ford, 1950; Maruyama, 1970; Slatkin, 1981; Barton, 1993; Pulliam, 1988; Whitlock, 2003).

2.5 Group selection

The enthusiastic approach of early group selec-tionists to explain all evolution of cooperation from this one perspective (Wynne-Edwards, 1962) has met with vigorous criticism (Williams, 1966) and even a denial of group selection for decades. Only an embattled minority of scientists defended the approach (Eshel, 1972; Wilson, 1975; Matessi and Jayakar, 1976; Wade, 1976; Uyenoyama and Feldman, 1980; Slatkin, 1981; Leigh, 1983; Szathmary and Demeter, 1987). Nowadays it seems clear that group selection can be a powerful mechanism to promote cooperation (Sober and Wilson, 1998; Keller, 1999; Michod, 1999; Swenson *et al.*, 2000; Kerr and Godfrey-Smith, 2002; Paulsson, 2002; Boyd and Richerson, 2002; Bowles and

Gintis, 2004; Traulsen et al., 2005). We only have to make sure that its basic requirements are fulfilled in a particular situation (Maynard Smith, 1976). Exactly what these requirements are can be illustrated with a simple model (Traulsen and Nowak, 2006).

Imagine a population of individuals subdivided into groups. For simplicity, we assume that the number of groups is constant and given by m. Each group contains between 1 and n individuals. The total population size can fluctuate between the bounds m and nm. Again, there are two types of individual, cooperators and defectors. Individuals interact with others in their group and thereby receive a payoff. At each time step a random individual from the entire population is chosen proportional to payoff in order to reproduce. The offspring is added to the same group. If the group size is less than or equal to n then nothing else happens. If the group size, however, exceeds n then with probability q the group splits into two. In this case, a random group is eliminated (in order to maintain a constant number of groups). With probability $1 - q$, the group does not divide, but instead a random individual from that group is eliminated (Figure 2.6).

This minimalist model of multilevel selection has some interesting features. Note that the evolutionary dynamics are entirely driven by individual fitness. Only individuals are assigned payoff values. Only individuals reproduce. Groups can stay together or split (divide) when reaching a certain size. Groups that contain fitter individuals reach the critical size faster and therefore split more often. This concept leads to selection among groups, although only individuals reproduce. The higher level selection emerges from lower level reproduction. Remarkably, the two levels of selection can oppose each other.

As before, we can compute the fixation probabilities, ρ_C and ρ_D, of cooperators and defectors to check whether selection favors one or the other. If we add a single cooperator to a population of defectors, then this cooperator must first take over a group. Subsequently the group of cooperators must take over the entire population. The first step is opposed by selection, the second step is favored by selection. Hence, we need to find out if the
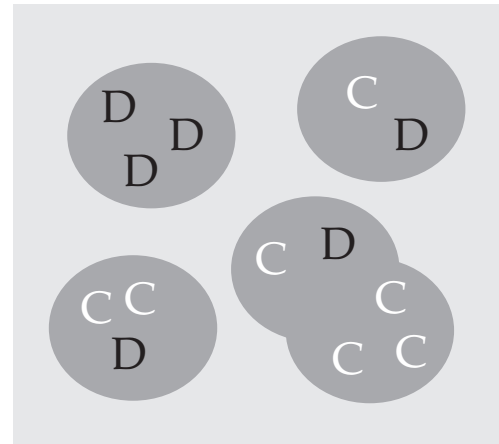

Figure 2.6 A simple model of group selection. A population consists of m groups of maximum size n. Individuals interact with others in their group in the context of an evolutionary game. Here we consider the game between cooperators, C, and defectors, D. For reproduction, individuals are chosen from the entire population with a probability proportional to their payoff. The offspring is added to the same group. If a group reaches the maximum size, n, then it either splits in two or a random individual from that group is eliminated. If a group splits, then a random individual dies, in order to keep the total population size constant. This metapopulation structure leads to the emergence of two levels of selection, although only individuals reproduce.

overall fixation probability is greater to or less than what we would obtain for a neutral mutant. An analytic calculation is possible in the interesting limit $q \ll 1$, where individuals reproduce much more rapidly than groups divide. In this case, most of the groups are at their maximum size and hence the total population size is almost constant and given by $N = nm$. We find that selection favors cooperators and opposes defectors, $\rho_C > 1/N > \rho_D$, if

$$b/c > 1 + n/(m - 2) \qquad (2.5a)$$

This result holds for weak selection. Smaller group sizes and larger numbers of competing groups favor cooperation. We also notice that the number of groups, m, must exceed 2. There is an intuitive reason for this threshold. Consider the case of $m = 2$ groups with $n = 2$ individuals. In a mixed group, the cooperator has payoff $-c$ and the defector has payoff b; the defector/cooperator difference is $b + c$. In a homogeneous group, two cooperators have payoff $b - c$, while two defectors

have a payoff of 0. Thus the disadvantage for cooperators in mixed groups cannot be compensated for by the advantage they have in homogeneous groups. Interestingly, however, for larger splitting probabilities, q, we find that cooperators can be favored even for $m = 2$ groups. The reason is the following: for very small q, the initial cooperator must reach fixation in a mixed group; but for larger q, a homogeneous cooperator group can also emerge if a mixed group splits, giving rise to a daughter group that has only cooperators. Thus, larger splitting probabilities make it easier for cooperation to emerge.

Let us also consider the effect of migration between groups. The average number of migrants accepted by a group during its lifetime is denoted by z. We find that selection favors cooperation provided that

$$b/c > 1 + z + n/m \qquad (2.5b)$$

In order to derive this condition we have assumed weak selection and $q \ll 1$, as before, but also that both the numbers of groups, m, and the maximum group size, n, are much larger than 1.

Group selection (or multilevel selection) is a powerful mechanism for the evolution of cooperation if there is a large number of relatively small groups and migration between groups is not too frequent.

2.6 Conclusion

We end by listing the five rules that we mentioned in the beginning. These rules represent laws of nature governing the natural selection of cooperation.

1. Kin selection leads to cooperation if $b/c > 1/r$, where r is the coefficient of genetic relatedness between donor and recipient.
2. Direct reciprocity leads to cooperation if $b/c > 1/w$, where w is the probability of playing another round in the repeated Prisoner's Dilemma.
3. Indirect reciprocity leads to cooperation if $b/c > 1/q$, where q is the probability of knowing the reputation of a recipient.
4. Graph selection (or network reciprocity) leads to cooperation if $b/c > k$, where k is the degree of the graph; that is, the average number of neighbors.
5. Group selection leads to cooperation if $b/c > 1 + z + n/m$, where z is the number of migrants accepted by a group during its lifetime, n is the group size, and m is the number of groups.

In all five theories, b is the benefit for the recipient and c the cost for the donor of an altruistic act.