

4

Evolutionary game theory

JOHN M. MCNAMARA AND FRANZ J. WEISSING

Overview

Evolutionary game theory may have done more to stimulate and refine research in animal behaviour than any other theoretical perspective. In this chapter, we will review some of the insights gained by applying game theory to animal behaviour. Our emphasis is on conceptual issues rather than on technical detail. We start by introducing some of the classical models, including the Hawk–Dove game and the Prisoner’s Dilemma game. Then we discuss in detail the main ingredients of a game-theoretical approach: strategies, payoffs and ‘solution concepts’ such as evolutionary stability. It should become clear that first-generation models like the Hawk–Dove game, while of enormous conceptual importance, have severe limitations when applied to real-world scenarios. We close with a sketch of what we see as the most important gaps in our knowledge, and the most relevant current developments in evolutionary game theory.

4.1 Introduction

Social behaviour involves the interaction of several individuals. Therefore within most social contexts the best thing to do depends on what others are doing. In other words, within social contexts selection is typically frequency-dependent (Ayala & Campbell 1974, Heino *et al.* 1998). Game theory was originally formulated to predict behaviour when there is frequency dependence in economics, for example competition between firms (von Neumann & Morgenstern 1944, Luce & Raiffa 1957). John Maynard Smith and George Price had the fundamental insight that this theory could also be used to predict the evolutionary outcome under frequency-dependent selection within biology (Maynard Smith & Price 1973, Maynard Smith 1982). Their idea was that, rather than following the evolution of a population over time, one could use ideas from game theory to characterise the

eventually stable endpoints of the evolutionary process. Their concept of an *evolutionarily stable strategy* (ESS) attempts to capture the properties of these endpoints. To understand this concept let us refer to a strategy as the *resident strategy* if almost all population members adopt this strategy. Then in intuitive terms, a given resident strategy is an ESS if no rare *mutant strategy* can invade this population under the action of natural selection. A necessary condition for this to be true is that no mutant strategy has greater fitness than the resident strategy. This latter condition is the familiar Nash equilibrium concept of economics (Nash 1950), with economic payoffs replaced by fitness payoffs.

The fact that an ESS corresponds to a Nash equilibrium has important conceptual implications. It implies that natural selection will shape social behaviour in such a way that it resembles the behaviour of *Homo economicus*, an agent whose decisions are guided by rational deliberations (Persky 1995). As a consequence, many insights from economic theory apply to animal behaviour, without having to assume that animals are ‘rational’ in any way (Hammerstein & Selten 1994). Since *Homo economicus* does not exist in our world, we may even have to face the ‘rationality paradox’ that economic theory describes animal behaviour better than human behaviour (Hammerstein 1996).

The ‘quasi-rationality’ of adaptive animal behaviour allows us to adopt many tools, ideas and insights from classical game theory. Concepts such as a payoff matrix, and examples such as the Prisoner’s Dilemma game have found their way into the biological sciences. Perhaps more important is the adoption of ‘strategic thinking’, which led biologists to realise that important aspects of behavioural programmes may not easily be observable and that seemingly minor aspects of the interaction structure can have major implications for the evolution of behaviour. Insights like the one that in many contexts signals have to be costly in order to be reliable (the ‘handicap principle’: Zahavi & Zahavi 1997) had independent origins in both biology and economics (Spence 1973).

In this chapter we focus on conceptual issues rather than on technical detail (for technical reviews see van Damme 1987, Reeve & Dugatkin 1998, Gintis 2000, McGill & Brown 2007). Our goal is to introduce the simple concepts that stimulated research, and to indicate where these concepts need further refinement to reflect more of the real-world complexity of behaviour.

4.2 Setting the scene: classical models

One of the early applications of evolutionary game theory was to the evolution of levels of aggression between individuals. Maynard Smith and Price (1973) considered a scenario in which two randomly selected population members contest a resource such as a mate, food item or territory. If an individual obtains the resource, the fitness of that individual is increased by an amount V . Each individual adopts one of two actions. Action ‘Dove’ specifies that the individual will display to opponent, but will run away if opponent attacks. Action ‘Hawk’ specifies that the individual will attack opponent and fight if opponent fights back. The possible outcomes are then: if both choose Dove each contestant is equally likely to obtain the resource; if one chooses Dove and the other chooses Hawk then the Hawk obtains the resource; if both choose Hawk each is equally likely to

Table 4.1. *The payoff structure of the Hawk–Dove game. Table entries give the fitness payoff to the focal player.*

	Opponent plays Dove	Opponent plays Hawk
Focal player plays Dove	$V/2$	0
Focal player plays Hawk	V	$(V - C)/2$

win the fight and hence the resource, but the loser has a reduction in fitness of C due to injuries sustained. The fitness consequences of the various combinations of actions are summarised in Table 4.1.

For this scenario, suppose that a strategy specifies the probability, p , that an individual plays Hawk in the contest. If the resident strategy is always to play Dove ($p = 0$) the best response for a rare mutant is to always play Hawk ($p = 1$). This is because the mutant always gets the resource in the contest since it attacks residents who then run away. In contrast, consider the situation in which the resident strategy is to always play Hawk ($p = 1$). Then if the cost of injury is greater than the value of the resource ($V < C$) the best response of a mutant is to always play Dove ($p = 0$), since this avoids the cost of injury. This illustrates the basic frequency dependence in this situation: the best thing to do depends on the actions of others. It is easy to show that when $V < C$ the strategy $p^* = V/C$ has the property that a best response to the resident strategy p^* is also to adopt strategy p^* . In other words, p^* is a Nash-equilibrium strategy. However, this does not guarantee that p^* is an ESS. When the resident strategy is p^* any single mutant has the same fitness as a resident, and to verify that p^* is an ESS one must show that mutants are selected against when the frequency of mutants starts to increase, so that mutants can play mutants in the contest. This can be verified (Maynard Smith 1982), and it turns out that $p^* = V/C$ is the unique ESS when $V < C$, and $p^* = 1$ is the unique ESS when $V \geq C$.

Note that, although the reasoning behind an ESS suggests that a population that reaches an ESS will not evolve away from this strategy, it does not guarantee that evolution will actually lead to an ESS. This shortcoming of the original ESS concept will be discussed further below.

The Hawk–Dove game involves the competitive interaction of two population members. But the idea of an ESS can be extended to deal with n -player games – for example, to the analysis of dominance hierarchies (Chapters 7 and 14). The concept can also be used to analyse situations in which individuals ‘play the field’. This term refers to situations in which the fitness of an individual depends on some overall characteristic of the resident population. For example, in analysing evolutionarily stable sex ratios, as the proportion of females in the breeding population increases the advantage to a breeding male increases for two reasons: there are more females to mate with, and fewer males to compete with. Consequently the advantage of producing sons over producing daughters depends in a non-linear way on the proportion of resident members producing sons as opposed to producing daughters (Seger & Stubblefield 2002).

The evolution of male size when there is male–male competition is another case of playing the field. Here the fitness of a male of a given size depends in a typically non-linear way on the size distribution of resident males. Because of the non-linearity the fitness of the male cannot be expressed as the sum of fitness contributions from all the pair-wise interactions with other males, so that this situation cannot be reduced to a series of independent two-player games. As we discuss further below, in many settings even the Hawk–Dove should not really be considered as a two-player game but as playing the field. Specifically, when an individual must play a series of Hawk–Dove games, fitness may not be the sum of the fitness contributions from each game.

Although all individuals in a population are maximising their fitness at an ESS, one of the important messages of game theory is that mean population fitness is not usually maximised. This is because all individuals are doing the best for themselves, and this is not necessarily the best for the population as a whole. To illustrate this point consider the Hawk–Dove game. In this game the same resources are available per pair regardless of the strategies employed by population members; one pair member gets the resource of value V , the other does not. But if fighting occurs individuals lose fitness through injury. Thus mean population fitness is maximised by avoiding fights, i.e. by all population members playing Dove. This would not evolve, however, since in such a population an individual playing Hawk would gain an advantage. The ‘tragedy of the commons’ also illustrates this point; if there is some common good that all individuals can share there is always selection pressure to take more than a fair share, often resulting in the overuse and demise of the common good (Hardin 1968).

As a final example of this point consider the Prisoner’s Dilemma game (e.g. Axelrod & Hamilton 1981). This game is played between two opponents. Each has a choice between cooperating with their partner or defecting. The fitness payoff to an individual depends on this focal individual’s choice and that of the opponent (Table 4.2). In this illustration, regardless of the choice of action by the opponent it is best to Defect (since $5 > 3$ and $1 > 0$). The opponent has an identical payoff matrix and should do likewise. Thus both players Defect at the unique ESS (and each receives a fitness payoff of 1 unit for the case illustrated). In contrast, had they both cooperated they would each have obtained a higher payoff (3 units in the case illustrated). This property of the Prisoner’s Dilemma has made it a test-bed for models of the evolution of cooperation. The challenge is to understand how selfish behaviour (produced by the action of natural selection) can lead individuals to cooperate. In other words, when is it in an individual’s best interests to cooperate?

Take-home messages of section 4.2

- In social contexts, selection is typically frequency-dependent.
- In the case of frequency-dependent selection, mean population fitness is usually not maximised. To predict the outcome of selection, one should not look for fitness optima but for evolutionarily stable strategies (ESS). These maximise fitness conditional on the behaviour of other population members.
- Every ESS corresponds to a Nash equilibrium and can therefore be viewed as ‘quasi-rational’ behaviour.

Table 4.2. *Illustration of the payoff structure of the Prisoner's Dilemma game. Table entries give the fitness payoff to the focal player.*

	Opponent cooperates	Opponent defects
Focal player cooperates	3	0
Focal player defects	5	1

4.3 Strategies

An important insight of game theory is that of *strategic thinking*. A strategy is a rule for choosing which action to perform. For example, in the Hawk–Dove game the probabilistic rule – with probability $1 - p$ play Dove, with probability p play Hawk – is an example of a mixed strategy. Strategies are rules that are contingent on circumstances, and the choice of action can depend on an organism's state, its role, the information that it has, etc. For example, in a variant of the Hawk–Dove game two individuals contest a territory. One individual is assigned the role of territory owner, the other that of intruder. Although each player can choose to play Hawk or Dove, as in the standard game, the set of strategies is enlarged. The rule, if intruder play Dove, if owner play Hawk, is an example of a strategy that is contingent on ownership (this strategy is referred to as Bourgeois: Maynard Smith & Parker 1976). In the Repeated Prisoner's Dilemma game (e.g. Gintis 2000) two individuals play the Prisoner's Dilemma several times against each other. A strategy for this repeated game specifies the choice of action in the current round as a function of the whole history of the game up to that point. For example, the rule (known as Tit-for-Tat), choose the same action as opponent chose on the last round, is an example of a simple strategy where the action taken is contingent on the opponent's behaviour in the past.

Game-theory payoffs are assigned to the combinations of actions, but the analysis of the game is in terms of strategies and not actions. In evolutionary terms this is the right level of analysis, since genes may be viewed as 'recipes' that prescribe how an organism should act contingent on the current conditions. As we have seen before, it can be useful in social contexts to randomise one's behaviour. Strategies including an element of randomisation (for instance, play Hawk with probability 0.75 and Dove with probability 0.25) are called *mixed* strategies, whereas strategies that prescribe actions in a deterministic way (like Tit-for-Tat or Bourgeois) are called *pure* strategies.

In formulating a game, the details of the model may matter enormously for the predicted outcome; in particular, it is crucial what strategies are allowed. The Hawk–Dove game (with $V < C$) illustrates this point. In the original game a (mixed) strategy specifies the probability, p , that an individual will play Hawk. There is a unique ESS at which individuals play Hawk with probability $p^* = V/C$. In the version of the Hawk–Dove game with territorial ownership a strategy is specified by two probabilities, p_0 and p_I , where p_0 is the probability of playing Hawk when in the role of territory owner and is the probability of playing Hawk when in the role of intruder. For the latter game the resident strategy $p_0 = V/C$, $p_I = V/C$ is not evolutionarily stable (Maynard Smith & Parker 1976). Instead

there are two ESSs, both of which are pure strategies. One ESS is Bourgeois, i.e. always play Hawk when owner and always play Dove when intruder ($p_O = 1, p_I = 0$). The other ESS is Anti-Bourgeois; i.e. always play Dove when owner and always play Hawk when intruder ($p_O = 0, p_I = 0$).

It is a general feature of asymmetric games that ESSs are pure strategies (Selten 1980, 1983), at least when payoffs are fixed (cf. Webb *et al.* 1999). The owner–intruder terminology suggests that there are relevant differences between the two types of players, for example with respect to fighting ability or the payoffs received. However, the above result also applies to situations where the roles are just arbitrary labels. In the absence of such labels, a mixed-strategy population is predicted, where fights occur with probability $(VC)^2$. In the presence of an otherwise irrelevant label, a pure-strategy population will result where fights are avoided due to the strategic convention (either Bourgeois or Anti-Bourgeois). Hence a seemingly trivial change in the formulation of a game can have major implications for the evolutionary outcome.

Models of parental effort provide a second illustration of the importance of the strategic perspective (see also Chapter 11). In the parental-effort model of Houston and Davies (1985) each parent chooses its level of effort independently of the other parent. There is no variation within each sex, but males and females may differ from each other. At evolutionary stability, effort of one parent depends on the effort expended by the other parent. Thus, if male effort increases female effort decreases, and vice versa. In this model it is implicitly assumed that parents do not respond to each other in real time. Instead, levels of effort are fixed and genetically determined, and change only occurs over evolutionary time (McNamara *et al.* 1999). This seems unrealistic; in real populations there is will always be within-sex variation in effort, so that it will be advantageous to respond directly to partner's effort. Efforts will then be negotiated using response rules. Thus models should be looking at the evolution of response rules rather than of efforts. In other words, we should be taking the response rules as genetically determined rather than the efforts, and at evolutionary stability the response rule used by males will depend on the response rule used by females, and vice versa. As McNamara *et al.* (1999) show, this can mean that the effort chosen as the result of the negotiation is not the same as the best effort given the (negotiated) effort of partner.

As a third example, consider a simple game between parents in which each decides whether to care for their common young or to desert. The points we wish to make are more fully discussed by McNamara and Houston (2002). The payoff matrices for males and females are given in Table 4.3. Here the differences in payoff between males and females might arise because males are better at care than females. We consider two versions of this game. In the first version, which we refer to as the *simultaneous choice* version, each parent chooses whether to care or desert without knowing the decision of the other parent. Neither parent can change its mind once the partner's decision is revealed. In this version a male's strategy specifies the probability he will desert; similarly, a female's strategy specifies the probability she will desert. As can be seen from the payoff matrices, if the female deserts it is best for the male to care, and if the male cares it is best for the female

Table 4.3 *Payoffs in the parental desertion game*

(a) Payoff to the male

	Female cares	Female deserts
Male cares	10	7
Male deserts	8	2

(b) Payoff to the female

	Male cares	Male deserts
Females cares	10	4
Female deserts	11	2

to desert. Thus male care and female desertion are best responses to each other and are in Nash equilibrium. In fact it is easy to see that this gives the unique ESS: at evolutionary stability males always care and females always desert.

In the second version of the desertion game the male is the first to choose whether to desert. If he decides to care he stays with the young, if he decides to desert he departs. The female then decides whether to care or desert, her decision being contingent on whether the male is still present or not. In this *sequential choice* version the strategy of a male again specifies his probability of desertion. In contrast, the strategy of a female is now a contingent rule that specifies the probability of desertion if the male cares, and the probability of desertion if he deserts. Suppose that in a population the resident male strategy is for a male to always desert, and the resident female strategy is to desert if the male cares and care if he deserts. Then it can be seen from Table 4.3 that no mutant male adopting a different strategy can do better. If instead of deserting he cares, the female will desert and he will be worse off. Similarly, no mutant female adopting a different strategy can do better. Thus the population is in Nash equilibrium. Note that in this population females are never observed to desert because males never care, but the threat of partner desertion keeps the male away from caring. Thus an unobserved aspect of the female strategy is crucial to the game.

So is this resident population (males always desert and females care if the male deserts and desert if he cares) evolutionarily stable? The problem is that, since all males desert, there is no selection pressure acting on what the female would do if the male were to care. Thus the female strategy that specifies care under all circumstances does equally well, and can increase by drift. The resident strategy is therefore not an ESS according to the original definition of Maynard Smith and Price (1973). However, suppose that males occasionally care by mistake. Then it does matter what the female would do if the male were to care, and the female strategy of always caring is strictly worse than the strategy of caring if and only if the male deserts. Thus the population is evolutionarily stable if infrequent mistakes are assumed. As Selten (1983) emphasises, occasional mistakes can stabilise the solution of a game. Selten refers to an equilibrium that is stable under infrequent mistakes as a *limit ESS*.

In the simultaneous-choice version of the above desertion game, the payoff to the male at evolutionary stability is 7. The payoff to a male at evolutionary stability in the sequential-choice version is 8. Thus the male gains an advantage by being the first to choose. However, it is the information that each parent has when making its decision, rather than the sequence of events, that is crucial. When the female makes her decision she has reliable information on the male's action. When the male makes his decision he does not know the action of the female. It is this informational asymmetry that gives the males the advantage. Interestingly, it is just the uninformed party (the males) that profits from the information asymmetry. This is in striking contrast to (frequency-independent) optimisation problems, where extra information always provides an advantage. McNamara *et al.* (2006) explicitly consider whether it is better to give information, to receive it or to be ignorant in a two-player game.

Similar conclusions have been drawn in 'playing the field' contexts, such as determining the evolutionarily stable sex ratio of offspring. If a mother does not make her sex-ratio decision contingent on her state or environmental conditions, she should invest equal amounts of resources into male and female offspring (*equal-allocation principle*: e.g. Seger & Stubblefield 2002). If male and female offspring are equally costly to produce, this should lead to a 50:50 sex ratio. Whenever male and female offspring are differentially affected by the mother's state or environmental conditions, the mother should make her sex-ratio decision contingent on these parameters (Trivers & Willard 1973). If she does, the population sex ratio at the ESS is no longer 50:50 (Frank & Swingland 1988), showing again that the results obtained in one type of game context cannot easily be generalised to another context, even if the differences between the models seem of minor importance. As in the desertion game considered above, information asymmetries are also of crucial importance for evolutionarily stable sex ratios. For example, Pen and Taylor (2005) show that in a queen-worker conflict over the sex ratio the workers can much better achieve their preferred sex ratio at ESS if they are not informed about the mother's decision than if they have this information. Of course, having more information can also provide an advantage. For example, Pen and Weissing (2002) show that in the conflict between the male and the female parent over the sex ratio of their offspring the informed party 'wins' the conflict.

At an ESS, individuals should make their behaviour contingent on the available information. As a consequence, it can be advantageous to manipulate the information available to an opponent. In some cases (e.g. the desertion game) a player can put an opponent at a disadvantage by giving the opponent reliable information about the player's own action (cf. Brams 1983). In other cases, it is better to hide one's intentions, or to actively deceive the opponent. It has been argued (e.g. Trivers 1985) that this intention and deception may even lead to the evolution of self-deception, since organisms can deceive their opponents more efficiently if they also deceive themselves.

In the sequential version of the desertion game, the fact that a male is not present gives the female reliable information about the male's action. This benefits the male. So can the female (who chooses second) gain an advantage by giving her partner reliable information on what she will do? In particular, suppose that the female threatens to desert if the male

deserts. If this threat is credible it will give the female an advantage, since the male will be forced to care. So is the threat credible? To analyse this, consider a population in which the female strategy is to always desert regardless of the male's action, and the male strategy is always to care. These strategies are again best responses to one another. In particular, the male is forced to care by the female's threat that she will desert if he deserts. However, mistakes by the males now destabilise the equilibrium. If a male is not present (either because he makes a mistake and deserts, or because he tries to care but is killed), then it is not optimal for the female to carry out her threat. Thus a genuine threat will not evolve; i.e. the threat by the female is not credible and is not reliable information in this game.

One way to give an opponent reliable information is to handicap oneself, so limiting one's choice of options (Schelling 1960, Elster 2000). In a version of the desertion game, Barta *et al.* (2002) allow females to choose their energy reserves before they (and their mate) decide whether to care for the young or desert. Females choose reserves that are too low for them to be able to care alone. This gives the male the reliable information that the female will desert if he deserts. He is therefore forced to care.

Strategic thinking has important implications for the analysis of a game. Consider the Repeated Prisoner's Dilemma game, where the players can make their choice of action in the current round dependent on the history of the game. After one round, there are already four possible histories corresponding to the four combinations of the actions Cooperate and Defect that can be chosen by the two players in the first round. Accordingly, there are $2^4 = 16$ pure 'local' strategies prescribing the choice of action in the second round. In the third round, there are already $4^2 = 16$ possible histories and $2^{16} = 65\,536$ pure local strategies. Obviously, even in relatively simple social contexts, the strategy space can reach astronomical proportions.

Perhaps more importantly, the number of Nash equilibria can also be huge. This is formalised in the so-called Folk Theorem of game theory (e.g. van Damme 1987), which states that for almost all games *every* outcome that is feasible and 'individually rational' in the one-shot game can be realised as the Nash equilibrium outcome of the repeated version of the game, provided that the number of repetitions is sufficiently large. For the Prisoner's Dilemma game in Table 4.2, this implies that *any* payoff outcome x with $1 \leq x \leq 3$ can be realised by a Nash equilibrium. The discussion in the literature often focuses on the extreme cases, i.e. on 'uncooperative' Nash strategies like Always Defect leading to the outcome $x = 1$, and on 'cooperative' Nash strategies like Tit-for-Tat leading to the outcome $x = 3$. The Folk Theorem shows that there is much more to it. In fact, it is one of the major challenges of evolutionary game theory to single out those Nash equilibria that are most 'reasonable' from an evolutionary perspective.

Take-home messages of section 4.3

- To understand social behaviour, one should apply *strategic thinking*, that is, focus on integrated behavioural programmes rather than on singular actions.
- In social contexts, at evolutionary stability alternative actions often have the same payoff, and the strategy of population members is to choose between the actions according to specified probabilities.

- Not all relevant aspects of a strategy may be directly observable. Strategies specify what to do under every eventuality, but in playing a game some circumstances may not be encountered because of the choice of actions by contestants. This does not mean that what a player would have done in these circumstances is irrelevant. Indeed it may be crucial to the outcome of the game, because it may be decisive in deciding whether to take actions which will lead to these circumstances.
- Seemingly irrelevant details concerning the strategic options available may matter enormously for the outcome of a social interaction. Asymmetries, even if they do not affect fighting ability or payoffs, may be crucial for solving a conflict.
- Having extra information is not always advantageous. A player can put an opponent at a disadvantage by giving the opponent reliable information about the player's action.
- Having extra options available is not always advantageous. One way to give an opponent reliable information is to handicap oneself, so limiting one's choice of options.
- Even relatively simple social interactions will often have a huge strategy set and a huge number of Nash equilibria. The challenge is not to find one of these equilibria, but to single out those that are 'reasonable' from an evolutionary perspective.

4.4 Payoffs

Game theory is usually based on a cost–benefit analysis in which the payoff to an organism depends on its strategy and that of other population members. In classical game models payoffs are proxies for fitness. They are meant to represent how much the outcome of the game increase fitness. However, the link between outcomes and fitness is often not considered carefully or spelt out.

Fitness is a quantity assigned to strategies, not to individual actions. The fitness of a strategy is some appropriate measure of the mean number of descendants left in the future by an organism following the strategy. Technically, this can be quantified by the asymptotic growth rate of a cohort of individuals following this strategy, which corresponds to the leading eigenvalue of a so-called population projection matrix (Metz *et al.* 1992, Caswell 2001). Note, however, that the projection matrix is not usually that for the whole population. Rather it is the projection matrix for a rare mutant within the population. Fitness is then the rate of invasion of the mutant into the resident population. This fitness measure is often not intuitive and may be difficult to apply in practice (e.g. Pen & Weissing 2002). In practice, evolutionary considerations are therefore usually based on alternative measures for the evolutionary success of a strategy.

In the simplest cases, expected lifetime reproductive success (i.e. the mean lifetime number of surviving offspring produced) is a good substitute measure of fitness. However, when offspring differ in their ability to spread the genes of their parents, this simple measure may not be adequate. For example, consider the sex-ratio problem. For this scenario a strategy for a female specifies the proportion of her offspring that are male. Just adding up the number of male and female offspring would not be an adequate fitness measure. To see this, assume that the population sex ratio in the offspring generation is male-biased. In such a situation, females are in short supply and a female offspring will on average leave more descendants (and, hence, spread its parents' genes more efficiently) than a male

offspring. Adding up male and female offspring in a fitness measure would therefore be like adding up apples and oranges.

The concept of reproductive value, which was introduced by Fisher (1930), quantifies the relative evolutionary importance of individuals in different states (Grafen 2006). Reproductive values may be viewed as weighing factors that allow to compare individuals in different states. For example, in the sex-ratio context above, the reproductive value of a male offspring is inversely proportional to the fraction s of males in the population, $v_m = 1/s$, while the reproductive value of a female offspring is inversely proportional to the fraction of females, $v_f = 1/(1-s)$. It can be shown that under many circumstances the weighted sum $W = n_m v_m + n_f v_f$ of male and female offspring, each offspring being weighted by its reproductive value, is a fitness measure that quantifies the evolutionary success of a strategy (e.g. Pen & Weissing 2002).

Reproductive values can be rigorously derived (Caswell 2001, Grafen 2006) for all kinds of situations where individuals differ in state, where 'state' may represent sex, age, body size, energy reserves, dominance status, previous actions taken in a game, and the like. For an individual, reproductive value quantifies its dependence of future descendants on current state. Under many circumstances an organism maximises its fitness by always behaving to maximise its reproductive value (McNamara 1993). The most important advantage of reproductive value is that it allows a cost–benefit analysis involving a comparison of individuals in different states. Should a mother, for example, defend her n kids if she runs a mortality risk μ but will save her kids from mortal danger with probability p ? With the help of reproductive values, the answer is straightforward: nest defence will be selected if the loss in terms of reproductive value of the mother, μv_{mother} , is outweighed by the gain in terms of reproductive value of her kids, npv_{kids} .

Thus reproductive value, or some equivalent, can be used as a surrogate currency for fitness in game payoffs. We can use the Hawk–Dove game to illustrate the use of reproductive value. Suppose that males are searching for females to mate with. We assume that if a male encounters a female that is not contested by another male, he mates with her. If the female is contested by another male he must decide to play Hawk or Dove in a contest with this male. If the male wins the contest he gets to mate with the female, if he loses a Hawk versus Hawk fight he dies with probability z . In this game the payoff V for winning the contest corresponds to the value of a mating, which we can normalise to 1. Let R denote the reproductive value of the male after the contest. Then the cost of losing a fight is $C = zR$, since this is the expected loss in reproductive value as a result of losing a fight. Classical game theory would take C as given and predict an ESS probability of playing Hawk of $p^* = \min \{V/C, 1\}$. This approach may suffice for some purposes, but it is an approach that isolates each contest from its ecological and social setting, and can mislead. The problem is that R is not usually a given quantity. R corresponds to the future number of matings obtained by a male before he dies. This quantity depends on the number of uncontested females in the environment. It also depends on the fighting strategies employed by the focal male and other males in future contests, and hence depends on the solution of the game. In other words, the solution of the game depends on R , but R depends on the solution of

the game! To take this complication into account, it is better not to consider each contest in isolation. Instead, one proceeds as follows (e.g. Taylor & Frank 1996, van Boven & Weissing 2004).

We consider a situation in which all population members play some given strategy in contests (the resident strategy). For such a monomorphic population, the reproductive value R can be calculated by standard means. Given this value of R , one can now check whether any alternative mutant strategy would have a higher fitness than the resident and, thus, be able to invade. One then seeks a resident strategy that is not invadable and, hence, a Nash-equilibrium strategy and a possible ESS. For the case of the Repeated Hawk–Dove game, Houston and McNamara (1991) demonstrate that for some ecological parameters there can be two ESSs. At each, the R (and hence $C = zR$) that emerges is such that the behaviour in each contest conforms to the prediction of the classical Hawk–Dove game with $V = 1$ and this value of C . At one ESS individuals always play Hawk in contests, resulting in high mortality, low R and hence low C . This is then consistent with $p^* = 1$. At the other ESS individuals do not always fight, resulting in a higher R and hence C , and consistent with the lower value of p^* obtained.

In classical applications of game theory, the payoff parameters were usually assumed to be externally given. The above example illustrates that it is often more natural to assume that these parameters are intrinsically generated by feedbacks with the population strategy. Takings such feedbacks into account can lead to qualitatively different and sometimes surprising results (see van Boven & Weissing 2004 for an example). Many applications of game theory also neglect the ecological embedding of social interactions (Mylius & Diekmann 1995). The fitness of a resident population corresponds to the asymptotic growth rate of the population (Metz *et al.* 1992). In ecological equilibrium, this asymptotic growth rate has to be equal to zero, since the population would go extinct in case of a negative value and it would go to infinity in case of a positive value. In order to achieve a zero growth rate, not all payoff parameters can be fixed and externally given. In an ecologically realistic setting, at least some of the payoff parameters must be density-dependent, to allow population regulation. One might think that the way in which population regulation is achieved (e.g. by reduced survival or by reduced fecundity at high densities) is of marginal importance for the outcome of evolution. This, however, is not the case, since the mechanism of density regulation may differentially affect reproductive values and, hence, the outcome of evolutionary cost–benefit analyses. Consider, for example, a situation where density regulation acts via increased juvenile mortality at high population densities. At ecological equilibrium, the reproductive value of juveniles is relatively low, implying that adults should invest relatively much into their own survival rather than into the production and survival of offspring. The opposite would be the case if density regulation were to act via reduced survival of adults. Many examples, ranging from sex-ratio evolution (Pen & Weissing 2002) to cooperative breeding (Pen & Weissing 2000), demonstrate that neglecting the mechanism of density regulation can lead to highly misleading conclusions.

Take-home messages of section 4.4

- The payoffs assigned to strategies are proxies for fitness. In situations where individuals can differ in ‘state’ (e.g. sex, age, size, energy reserves, dominance status), the quantification of fitness is often difficult and not straightforward.
- In such situations, the reproductive-value concept allows a cost–benefit analysis, where individuals in different states are weighed according to their efficiency in spreading their genes to future generations.
- The solution of a game reflects the payoffs, but the payoffs often reflect the behaviour in the population and, hence, the solution of the game. Taking such feedbacks into account can strongly affect the evolutionary predictions, both quantitatively and qualitatively.
- All social interactions are embedded in an ecological context (Chapter 18). Neglecting this context and the corresponding density dependence of payoff parameters can lead to highly misleading evolutionary conclusions.

4.5 Evolutionary analysis

The evolutionary analysis of a social interaction typically starts with the specification of the set of feasible strategies. As we have seen above, the strategy set reflects assumptions on the interaction structure, the information available to the interacting agents, and the actions available in any possible situation. The strategy set also reflects all kinds of limitations and constraints at the sensory, cognitive or behavioural level. In a next step, fitness payoffs are assigned to the strategies. Since fitness is frequency-dependent in the case of social interactions, the fitness of an organism depends not only on its own strategy p , but also on the ‘population strategy’ u , the distribution of strategies in the population. Hence the selective forces acting in the population are characterised by a fitness function $W(p,u)$. As we have seen above, the definition of fitness is not always obvious, and the function W will reflect life-history considerations and ecological factors such as the mechanism of density dependence. Once the fitness function has been obtained, it can be used to make predictions concerning the expected outcome of natural selection. This is, however, less straightforward than one might think.

In the case of frequency-independent selection (where fitness $W(p)$ depends only on an organism’s own strategy p and not on the strategy distribution in the population), it is often useful to imagine evolution as a hill-climbing exercise on the ‘fitness landscape’ generated by the function W . From one generation to the next, fitness is expected to increase, until a (local) maximum is reached (Wright 1932). Such a strategy p^* , for which $W(p^*) > W(p)$ holds for all $p \neq p^*$ in the vicinity of p^* , may then be viewed as a potential outcome of evolution. The fitness-landscape metaphor can also be applied to frequency-dependent selection, but now the selective forces acting on behaviour are characterised by a bivariate function $W(p,u)$. For any given population strategy u , there is again the fitness landscape, but this landscape changes with any change of the population strategy (Fig. 4.1). Hence social evolution corresponds to a climb on a fluctuating fitness landscape, and the fluc-

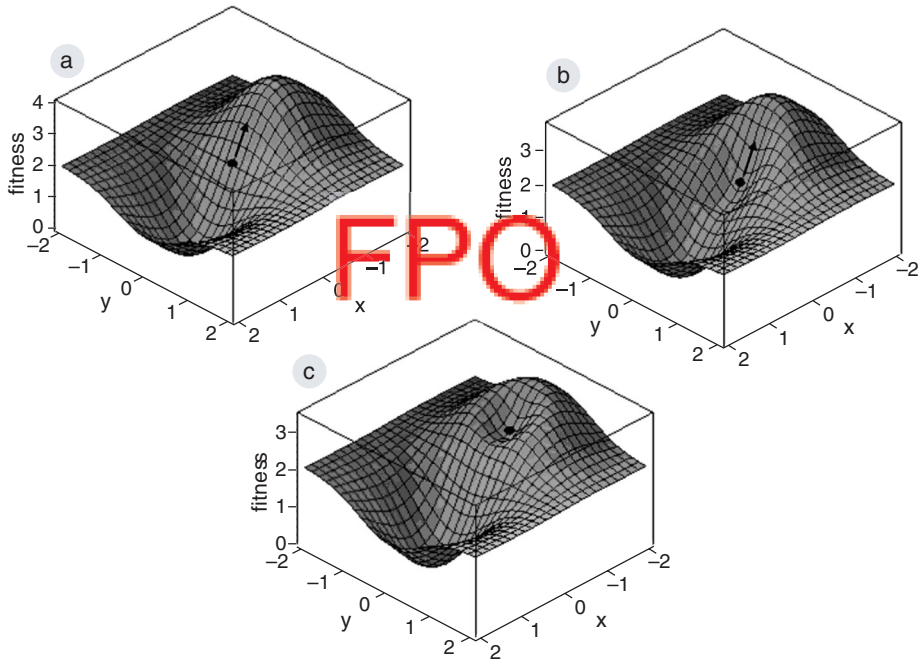


Figure 4.1. Representation of frequency-dependent selection as a climb on a fluctuating fitness landscape. The three panels depict the fitness landscape at three points in time. The corresponding population strategy u_t is indicated by a black dot. At each point of time, the population strategy changes in the direction of steepest ascent (arrow). A change in population strategy from (a) u_t to (b) u_{t+1} induces, however, a change in the fitness landscape. In the example, the fitness landscape is slightly depressed in the vicinity of the population strategy. As a consequence, the fitness of u_{t+1} in (b) is not necessarily higher than the fitness of u_t in (a). A depression of fitness near the population strategy often occurs in models where individuals with a strategy close to the established strategy in the population suffer most from intraspecific competition. In such a case, the population may even end up in a local fitness minimum (c), although the population changed in an uphill direction throughout the whole trajectory.

tuations themselves are the result of selection. As a consequence, the outcome of social evolution is not always obvious, and sometimes even counterintuitive.

The metaphor of a fluctuating fitness landscape explains, for example, the well-known fact that mean population fitness is usually not maximised in the case of frequency-dependent selection. Suppose that the population strategy is u_t at time t and that selection shifts the population to a new state u_{t+1} . Then the new state will typically have a higher fitness than the old state, but only with respect to the ‘old’ fitness landscape $W(p, u_t)$ generated by u_t . There is no a-priori reason why u_{t+1} should also have a higher fitness with respect to the new fitness landscape $W(p, u_{t+1})$. In fact, mean population fitness will often deteriorate in time, and in some cases the population may even converge to a state where mean population fitness is minimised (as illustrated in Fig. 4.1c).

Many seemingly similar concepts have been developed in order to predict the outcome of evolution in a system where the selective forces are characterised by a frequency-dependent fitness function $W(p,u)$. Here we only mention a few, in order to indicate why a single concept does not capture all aspects. In 1973, Maynard Smith and Price introduced the concept of an evolutionarily stable strategy, which is based on the idea that the stable endpoints of evolution correspond to resident strategies that cannot be invaded by alternative ‘mutant’ strategies. To this end, they considered a monomorphic resident population employing strategy p^* that is challenged by a rare mutant strategy p that occurs with a small frequency ε . Then the population strategy is given by $u\varepsilon = (1 - \varepsilon)p^* + \varepsilon p$. The residents will be immune against invasion by p if their fitness $W(p^*,u\varepsilon)$ exceeds $W(p,u\varepsilon)$, the fitness of the mutants. Accordingly, p^* is considered an ESS if $W(p^*,(1 - \varepsilon)p^* + \varepsilon p) > W(p,(1 - \varepsilon)p^* + \varepsilon p)$ for all $p \neq p^*$ and sufficiently small mutant frequency $\varepsilon > 0$.

Taking the limit $\varepsilon \rightarrow 0$ in the above definition of an ESS, we obtain the Nash-equilibrium condition from classical game theory: $W(p^*, p^*) \geq W(p, p^*)$ for all $p \neq p^*$. In words, no strategy is better against p^* than p^* itself; p^* is a best response to itself. Alternatively, p^* corresponds to a local maximum of $W(p, p^*)$, the fitness landscape generated by p^* itself. Hence every ESS is a Nash-equilibrium strategy. In practical applications, authors often only check the Nash condition, without demonstrating that the more stringent ESS condition of Maynard Smith and Price is also satisfied. Virtually all games considered in the literature have a least one Nash equilibrium (van Damme 1987), and specifying the Nash equilibria of a game is typically a much simpler task than specifying all ESSs. Nash-equilibrium strategies may be viewed as candidate ESSs, but one should be aware that the Nash-equilibrium condition is rather weak. In fact, many Nash equilibria are not ‘reasonable’ as the outcome of evolution.

As an example, consider the coordination game in Table 4.4. Two players have to choose independently between two options L and R, and they only get a positive payoff if they both choose the same option. In the human world, one might think of driving one’s car on the left (L) or the right (R) side of the road. In this example, there are two Nash equilibria in pure strategies, which both happen to be an ESS: (1) always play L, or (2) always play R. However, the game also has a third Nash strategy \tilde{p} : play each option with probability $\frac{1}{2}$. Whenever this strategy is established in a population, it does not matter whether one chooses L or R. Whatever one’s choice is, one can expect to get the payoffs 1 and 0 with probability $\frac{1}{2}$. Therefore $W(p, \tilde{p}) = \frac{1}{2}$ for all strategies p , implying $W(\tilde{p}, \tilde{p}) = W(\tilde{p}, \tilde{p})$. Hence the Nash condition is satisfied for \tilde{p} (be it in a weak sense), but it can be shown that \tilde{p} is not an ESS. Intuitively, it is obvious that the corresponding population is not very stable. Whenever there is a slight majority for L in the population, it is individually advantageous to play L as well. Accordingly, one would expect a rapid evolution towards the pure strategy L. The same applies vice versa whenever there is a majority for R. This is confirmed in our human world: there are many countries where everybody drives on the left side of the road, and many other countries where everybody drives on the right side. In contrast, there are few examples where individuals decide at random on which side of the road to drive!

Table 4.4. *Payoff structure of a coordination game.*

	Opponent plays L	Opponent plays R
Focal player plays L	1	0
Focal player plays R	0	1

In contrast to the ‘weak’ mixed-strategy Nash equilibrium \tilde{p} , the two pure-strategy equilibria of the coordination game are ‘strict’ Nash strategies. This means that the Nash condition is satisfied in the strict sense, $W(p^*, p^*) > W(p, p^*)$, for all $p \neq p^*$. Strict Nash strategies have many desirable properties. In particular, every strict Nash strategy is an ESS. However, whereas all games have at least one (and often many) Nash strategies, most games do not have a strict Nash equilibrium. Moreover, most games with a rich interaction structure also do not have a single ESS *sensu* Maynard Smith and Price. For this reason, the ESS concept has been weakened in several ways, leading to concepts such as *direct ESS* or *limit ESS* (Selten 1983, van Damme 1987). Perhaps most importantly, many authors now use (often implicitly) a local definition of the Nash property and the ESS condition. For example, p^* is considered a strict Nash strategy if $W(p^*, p^*) > W(p, p^*)$ holds for all $p \neq p^*$ in the vicinity of p^* . Such local conditions have the advantage that they can be checked relatively easily by standard methods from differential calculus. Biologically, the local version of such stability conditions corresponds to the assumption that a resident population of p^* -strategists is only challenged by mutant strategies that differ only slightly from p^* . If all mutations have a small effect, this is indeed plausible. However, one has to be aware that this assumption is not unproblematic, since even point mutations often have a macroscopic effect on the phenotype (for further discussion see Wolf *et al.* 2008).

The example of the coordination game illustrates that the real question is often not which strategy is invasion-proof, but which of the invasion-proof strategies will actually be achieved in the course of evolution. It is rather obvious that driving on the left and on the right side of the road are alternative ESSs, but the question remains why driving on the left side is the outcome in some countries, while driving on the right side is the outcome in others. In the case of the coordination game, we already gave an intuitive answer in terms of historical contingency: whenever there is the slightest bias in one direction, it is in the self-interest of the individual to follow this bias, leading to a self-reinforcing process of evolution in the direction on the initial bias.

In systems with a simple strategic structure, arguments like this can be formalised as follows (Taylor 1996, Geritz *et al.* 1998). Take any Nash strategy p^* . The question is whether evolution will lead towards this strategy. To this end, consider a resident population with strategy \hat{p} that is a small distance away from p^* . The \hat{p} -population will typically be not evolutionarily stable, implying that certain mutant strategies can invade. If the successful mutants are closer to p^* than \hat{p} , it is plausible to assume that evolution through successful invasion and replacement of the resident by the invader will eventually converge

towards p^* . A strategy p^* with the property that evolution by gene substitution will converge towards p^* (because the successful invaders are closer to p^* than the former resident \hat{p}) will be called an *evolutionary attractor* (but there are many alternative terms for this property in the literature, including ‘population stability’ and ‘convergence stability’). If, on the other hand, successful mutants are those that are further away from p^* than \hat{p} , one can assume that evolution will lead away from p^* . In such a case, p^* is called an *evolutionary repeller*.

Interestingly, the property of being evolutionarily stable against invasion attempts by mutants is, at least to a certain extent, independent of the property of being an evolutionary attractor or repeller. All four combinations are possible (Geritz *et al.* 1998). A strategy that is both an evolutionary attractor and evolutionarily stable is called a *continuously stable strategy* or CSS (Eshel 1983; see also Chapter 6). However, an ESS *sensu* Maynard Smith and Price can be an evolutionary repeller, i.e. an ESS may not be attainable by gene substitution events. Such a scenario has been dubbed ‘Garden of Eden’ (a desirable state that is not attainable). It is also possible that an evolutionary attractor is not evolutionarily stable. In such a case, called *evolutionary branching*, evolution will lead to a fitness minimum, which is a point of disruptive selection. Evolutionary branching plays an important role in models of sympatric speciation (Chapter 19; Dieckmann & Doebeli 1999, van Doorn *et al.* 2004). Finally, it is of course also possible that an evolutionary repeller is not evolutionarily stable. The rich field of dynamic behaviour in the case of frequency-dependent selection has led to the emergence of a new research field called *adaptive dynamics*, which reanalyses many classical models in a more stringent and consistent way (Geritz *et al.* 1998, Dieckmann 2004, McGill & Brown 2007).

The arguments given above mainly apply to contexts with a simple, one-dimensional strategy space, where a strategy p corresponds to a univariate continuous trait like a sex ratio, a switching time between activities, or a preference for a certain type of resource. In the case of social evolution, strategies are often conditional, and hence of a much more complicated structure. In such a situation, verbal or semi-quantitative arguments are usually not sufficient to predict the outcome of frequency-dependent selection. The so-called *gradient method* (also called *best response method* (McNamara *et al.* 1997, Pen *et al.* 1999) or *genetic algorithm method* (Crowley 2001, Hamblin & Hurd 2007) is an efficient numerical technique to calculate those strategies that are plausible outcomes of evolution. Basically, this method mimics the walk uphill across the (fluctuating) fitness landscape generated by the fitness function $W(p,u)$. Starting at a population strategy u_0 , this technique determines a subsequent strategy u_1 by taking a small step in the direction of steepest ascent on the fitness landscape (the fitness gradient). Repeating this procedure yields a sequence of population strategies u_t that will often converge to a Nash strategy p^* , which is both an evolutionary attractor and an invasion-proof strategy. The method can be further refined by choosing an appropriate step size and by including small errors in decision making (McNamara *et al.* 1997, van Doorn *et al.* 2003a, 2003b). Applications of the gradient method have revealed that, as in the case of univariate strategies, evolutionarily stable strategies are not necessarily dynamically attainable. It turns out that classical non-dynamical

game theory has often focused on equilibria that are biologically not relevant, since they are not approached by any evolutionary trajectory (e.g. Hamblin & Hurd 2007).

The gradient method is particularly useful in the case of social interactions that, because of the underlying strategic complexity, have a multitude of Nash equilibria. Van Doorn *et al.* (2003a), for example, apply this technique to the Repeated Hawk–Dove game, which, like the Repeated Prisoner’s Dilemma, has a huge number of equilibria. They start by considering the primordial situation where the interacting organisms do not make use of their memory and hence do not employ refined strategies. Under this assumption, the population will converge to $p^* = V/C$, the ESS of the one-shot Hawk–Dove game. Subsequently, van Doorn and colleagues add strategic complexity to the model in a stepwise manner, corresponding to the improvement of behavioural architecture in the course of evolution. Taking the evolutionary outcome for the previous level of complexity as their point of departure, they use the gradient method to determine the most plausible behaviour for the next level of complexity. Irrespective of the exact sequence in which strategic complexity was added, only two types of equilibria result from this type of analysis. One of these equilibria is particularly interesting, since it corresponds to the ‘winner–loser’ effect often observed in agonistic interactions (Chapters 7 and 14). At this equilibrium, the winners of previous fights tend to be aggressive (play Hawk) in the future, while the losers tend to be peaceful (play Dove). This is true even if winning or losing is equally likely for all individuals and hence does not reflect differences in fighting ability. Accordingly, this model provides a plausible explanation for the evolution of social dominance that does not rely on differences in payoffs, information or resource-holding potential.

Besides being a potent technique for determining potential evolutionary attractors, the gradient method has the advantage that it is a useful heuristic, since it reflects a popular way to imagine the evolutionary process. It is, however, important to realise that the metaphor of natural selection corresponding to a hill-climbing exercise on a fitness landscape has severe limitations. While virtually all methods of evolutionary game theory are based on fitness considerations, the trajectory and outcome of evolution is not solely determined by fitness differences among strategies. In sexually reproducing populations, genetic processes such as Mendelian segregation and recombination may affect the outcome of evolution in two important respects. First, genetic processes may impose constraints on the strategy combinations that are feasible in a sexual population. Consider, for example, an ESS that can only be realised by a heterozygous genotype (a genotype Aa consisting of two different alleles A and a). Aa individuals do not transmit their genotype, but either an A -allele or an a -allele to their offspring. As a consequence, although Aa offspring are produced, so too are homozygous AA and aa offspring, even if the latter have a rather low fitness. Hence, because of Mendelian segregation, a purely heterozygous ESS population is not feasible, despite potentially strong fitness differences between genotypes. Second, genetic processes can alter the direction of evolution. Again, this is related to the fact that in a sexual population

genotypes (i.e. strategies) are not directly transmitted from the parents to their offspring. Instead, the genes making up a genotype are reshuffled by recombination, and the parents only transmit half of their diploid set of genes to their offspring. Processes like meiosis and recombination have the tendency to break up coadapted sets of genes, and therefore often counteract the tendency of natural selection to increase fitness. In fact, examples can be constructed where fitness is not maximised but minimised, due to processes at the genetic level (Moran 1964).

There are various modelling approaches that attempt to integrate selection forces caused by fitness differentials and processes at the genetic level in a coherent framework. The most important ones are the approach of population genetics that directly considers allele frequency changes at the genetic level (e.g. Cressman *et al.* 1996; see also Chapters 1 and 2); the approach of quantitative genetics that is based on selection gradients and genetic variances and covariances (Taper & Case 1992); and the canonical equation of adaptive dynamics that reflects evolution by gene substitution events (Dieckmann & Law 1996). Interestingly, all these approaches, though quite different in their underlying assumptions, can be described by very similar mathematical structures (Day 2005). When applied to specific examples, it turns out that genetic detail can be of considerable importance for the outcome of evolution. For example, as a result of small differences in genetic architecture, a population can either converge to an evolutionary equilibrium or show ongoing oscillations with large amplitude (e.g. van Doorn & Weissing 2006).

Because of results like this, focusing purely on fitness considerations (as is typically done in evolutionary game theory) can be problematic (Weissing 1996). Neglecting genetic detail may indeed lead to unreliable conclusions. This, however, confronts us with three challenges. First, as we have seen, frequency-dependent selection is already an intricate process, even in the absence of complications at the genetic level. Explicitly including non-trivial genetic assumptions will often make the models intractable. Second, the same kind of fitness scenario would necessitate a different analysis for different species, since the underlying genetics can not be assumed to be the same. As a consequence, it is often difficult to distil general conclusions from models including genetic detail. Third, and most importantly, the genetics underlying social behaviour is virtually unknown (but see Chapters 1, 2, 6 and 11). Despite enormous progress in fields such as ecological genomics, we predict that this will not change fundamentally in the next few decades. Including realistic genetic assumptions into models of frequency-dependent selection is therefore not an option for years to come.

Fortunately, the situation is less bleak than it may appear at first sight. Some methods of evolutionary game theory appear rather robust and do not depend on genetic details (Leimar 2001). It can also be shown that in the limiting case of weak selection the conclusions derived from fitness considerations are quite robust (Nagylaki *et al.* 1999). Finally, it can also be argued that genetic constraints and genetic processes interfering with adaptive evolution are themselves subject to selection and, as a consequence, will disappear in a long-term perspective (Hammerstein 1996, van Doorn & Dieckmann 2006, Galis & Metz 2007).

Take-home messages of section 4.5

- In the context of frequency-dependent selection, the often-used metaphor of an adaptive landscape has to be applied with caution. The fitness landscape is not constant but fluctuates due to the influence of natural selection. As a consequence, fitness can decrease over the generations. It is even possible that a population evolves to a fitness minimum.
- There are many different concepts that all try to capture the idea of evolutionary stability. Many applications focus on Nash-equilibrium strategies. Most games have, however, a multitude of Nash equilibria, and many of these equilibria lack evolutionary stability.
- An evolutionarily stable strategy is not necessarily attainable. In fact, evolution can lead away from an ESS ('Garden of Eden'). In contrast, evolution can lead to an evolutionarily unstable strategy, a so-called branching point. Such evolutionary branching may provide an explanation for processes like sympatric speciation.
- The *gradient method* (or *genetic algorithm method*) is a useful technique to determine the attainable and evolutionarily stable Nash-equilibrium strategies of an evolutionary game.

4.6 Conclusions and future directions*Whole systems, not just components*

We have noted that in the past there has been a tendency to consider simple games in isolation. An example is the parental desertion games introduced by Maynard Smith (1977). In his Model 2, each of two parents has to decide whether to desert their common young. Payoffs for desertion are given quantities, which potentially leads to a problem of consistency (Webb *et al.* 1999). The benefits of desertion for a male are meant to be in terms of re-mating opportunities. However, opportunities to re-mate necessarily depend on how many females are deserting, and hence the solution of the game. So, as with the Hawk–Dove game of Houston and McNamara (1991), payoffs determine behaviour, but behaviour feeds back to determine payoffs.

More generally, mating systems are characterised by many inherently linked games. For example, the mate-choice strategy of females should depend on both male genetic quality and male parental effort (see Chapters 10 and 18). However, if males trade off current effort against future mating opportunities, male effort should depend on female mate-choice strategies. Whether females choose to give extra-pair copulations (EPCs) should depend on how the social male reacts to loss in paternity, but this in turn should depend on his future mating opportunities and whether he can also gain EPCs, and hence on female strategies. Of course it is always possible to consider some component such as male parental effort, holding everything else fixed. In that way it is possible to investigate whether the level of male effort makes sense given the rest of the system. However, if the objective is to predict what mating system will evolve under given circumstances, one cannot consider components in isolation. The payoffs for one component are determined by what is going on in other components, and it is necessary to specify all feedbacks and links and to solve for all aspects at the same time.

These remarks probably apply to most social systems. A holistic approach is necessary if one is to understand what systems are possible and how they depend on environmental conditions. When feedback between different components is strong we might expect that there can be more than one stable, self-maintaining system in given circumstances. However, we might also expect that only a few types of system are possible in general. If these properties hold, then models could be used to investigate how one social system flips into another system as environmental conditions change. Because reversing the change would not necessarily restore the original system, one outcome of the investigation would be to reveal which phylogenetic trees are more likely than others (see Chapter 5).

The importance of variability

From the theory of sexual selection, it is known that costly female preferences can only evolve if there is sufficient (non-adaptive) variation in male traits (Chapter 10; Andersson 1994). This is not difficult to understand. Paying the costs of being choosy can only be advantageous if these costs are balanced by benefits. Such benefits only accrue if choosy females actually 'have a choice', i.e. if there is sufficient variation among males. The same basic principle also applies to many social situations, where the variance in a behavioural trait is often also important in determining how the mean value of the trait will evolve. For example, in a version of the Repeated Prisoner's Dilemma, McNamara *et al.* (2004) maintain variability in a population through mutation. They show that the direction of evolution is determined by the amount of mutation; cooperation only evolves above a critical mutation rate.

The ability to opt out of an interaction can radically change the predictions of a model. For example, in McNamara *et al.* (2008) individuals can break up a partnership in order to seek a more cooperative partner if their current partner is not cooperative. Whether they do so depends crucially on the variation in cooperation within the population. If future partners are all similar to the current one it is better to stick with the current partner and avoid the costs of seeking a new one, but if there is sufficient variation there are likely to be better partners and it is worth paying the cost of search. When this happens, uncooperative individuals lose their partners and must also pay the cost of seeking a new partner. Thus uncooperative individuals do badly and there will be selection for increased cooperation – so again increased variation leads to the evolution of cooperation. Note that markets of this sort tend to produce a non-random association of players even if they meet at random.

In deciding how to interact with another individual it may be possible to gain useful information on this individual by observing his or her past behaviour. However, observations waste time and energy, and so are likely to be costly. Thus if all individuals are similar it is not worth paying this cost because there is little to learn. In contrast, once there is sufficiently high variance in the trait of interest, it will be worth observing. Thus variation selects for social sensitivity. Once population members are socially sensitive this changes the selection pressure within the population: for example, individuals that are observed to be uncooperative with others may be shunned and will hence do worse.

Flexibility, personality and a complex world

Animals have to deal with a complex world. To do so they have rules of thumb which perform well on average, but are not optimal in every situation. The behavioural rule of thumb used by an individual is implemented via psychological mechanisms such as emotions and motivational states. Because of the variation that always exists in a population, different individuals are liable to have rules that are adjusted slightly differently. Furthermore, because animals following these rules are not completely flexible, individuals will display certain predictabilities in their behaviour. In other words, different individuals will have different personalities. These non-adaptive aspects of personality pose a challenge to evolutionary game theory. They mean that the idea of subgame perfection – i.e. animals always do the best given their current situation – must be abandoned. Now, instead, previous behaviour is indicative of current behaviour, leading to the establishment of reputation and generating the need for social sensitivity, as mentioned above.

Acknowledgements

This chapter was written when JMM was Tage Erlander Guest Professor at Gothenburg University.

Suggested readings

- Diekmann, O. (2004) A beginner's guide to adaptive dynamics. *Banach Center Publications* **63**, 47–86.
- Gintis, H. (2000) *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Houston, A. I. & McNamara, J. M. (1999) *Models of Adaptive Behaviour: An Approach Based on State*. Cambridge: Cambridge University Press.
- Reeve, H. K. & Dugatkin, L. A., eds (1998) *Game Theory and Animal Behavior*. Oxford: Oxford University Press.
- van Doorn, G. S., Hengeveld, G. M. & Weissing, F. J. (2003) The evolution of social dominance. *Behaviour*, **140**, 1305–1332, 1333–1358.

References

- Andersson, M. (1994) *Sexual Selection*. Princeton, NJ: Princeton University Press.
- Axelrod, R. & Hamilton, W. D. (1981) The evolution of cooperation. *Science*, **211**, 1390–1396.
- Ayala, F. J. & Campbell, C. A. (1974) Frequency-dependent selection. *Annual Review of Ecology and Systematics*, **5**, 115–138.
- Barta, Z., Houston, A. I., McNamara, J. M. & Székely, T. (2002) Sexual conflict about parental care: the role of reserves. *American Naturalist*, **159**, 687–705.
- Brams, S. (1983) *Superior Beings*. New York, NY: Springer.
- Caswell, H. (2001) *Matrix Population Models*, 2nd edn. Sunderland, MA: Sinauer Associates.
- Cressman, R., Hofbauer, J. & Hines, W. G. S. (1996) Evolutionary stability in strategic models of single-locus frequency-dependent viability selection. *Journal of Mathematical Biology*, **34**, 707–733.

- Crowley, P. H. (2001) Dangerous games and the emergence of social structure: evolving memory-based strategies for the generalized hawk–dove game. *Behavioral Ecology*, **12**, 753–760.
- Day, T. (2005) Modeling the ecological context of evolutionary change: déjà vu or something new? In: *Paradigms Lost: Routes to Theory Change*, ed. K. Cuddington & B. E. Beisner. New York, NY: Academic Press, pp. 273–309.
- Dieckmann, U. & Doebeli, M. (1999) On the origin of species by sympatric speciation. *Nature*, **400**, 354–357.
- Dieckmann, U. & Law, R. (1996) The dynamical theory of coevolution: A derivation from stochastic ecological processes. *Journal of Mathematical Biology*, **34**, 579–612.
- Dieckmann, U., Heino, M. & Parvinen, K. (2006) The adaptive dynamics of function-valued traits. *Journal of Theoretical Biology*, **241**, 370–389.
- Diekmann, O. (2004) A beginner's guide to adaptive dynamics. *Banach Center Publications* **63**, 47–86.
- Elster, J. (2000) *Ulysses Unbound*. Cambridge: Cambridge University Press.
- Ens, B. J., Weissing, F. J. & Drent, R. H. (1995) The despotic distribution and deferred maturity: two sides of the same coin. *American Naturalist*, **146**, 625–650.
- Eshel, I. (1983) Evolutionary and continuous stability. *Journal of Theoretical Biology*, **103**, 99–111.
- Eshel, I. & Sansone, E. (1995) Owner-intruder conflict, Grafen effect and self-assessment: the Bourgeois principle re-examined. *Journal of Theoretical Biology*, **177**, 341–356.
- Fisher, R. A. (1930) *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.
- Frank, S. A. & Swingland, I. R. (1988) Sex ratio under conditional sex expression. *Journal of Theoretical Biology*, **135**, 415–418.
- Galis F. & Metz J. A. J. (2007) Evolutionary novelties: the making and breaking of pleiotropic constraints. *Integrative and Comparative Biology*, **47**, 409–419.
- Geritz, S. A. H., Kisdi, E., Meszina, G. & Metz, J. A. J. (1998) Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology*, **12**, 35–57.
- Gintis, H. (2000) *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Grafen, A. (2006) A theory of Fisher's reproductive value. *Journal of Mathematical Biology*, **53**, 15–60.
- Hamblin, S. & Hurd, P. L. (2007) Genetic algorithms and non-ESS solutions to game theory models. *Animal Behaviour*, **74**, 1005–1018.
- Hammerstein, P. (1996) Darwinian adaptation, population genetics and the streetcar theory of evolution. *Journal of Mathematical Biology*, **34**, 511–532.
- Hammerstein, P. & Selten, R. (1994) Game theory and evolutionary biology. In: *Handbook of Game Theory with Economic Applications*, ed. R. J. Aumann & S. Hart. Amsterdam: Elsevier, Vol. 2, pp. 929–993.
- Hardin, G. (1968) The tragedy of the commons. *Science*, **162**, 1243–1248.
- Heino, M., Metz, J. A. J. & Kaitala V. (1998) The enigma of frequency-dependent selection. *Trends in Ecology and Evolution*, **13**, 367–370.
- Houston, A. I. & Davies, N. B. (1985) Evolution of cooperation and life history in dunnocks. In: *Behavioural Ecology: the Ecological Consequences of Adaptive Behaviour*, ed. R. Sibly & R. H. Smith. Oxford: Blackwell, pp. 471–487.
- Houston, A. I. & McNamara, J. M. (1991) Evolutionarily stable strategies in the repeated hawk–dove game. *Behavioral Ecology*, **2**, 219–227.
- Leimar, O. (1996) Life-history analysis of the Trivers and Willard sex-ratio problem. *Behavioral Ecology*, **7**, 316–325.
- Leimar, O. (2001) Evolutionary change and Darwinian demons. *Selection*, **2**, 65–72.

- Leimar, O., Hammerstein, P. & van Dooren T. J. M. (2006) A new perspective on developmental plasticity and the principles of adaptive morph determination. *American Naturalist*, **167**, 367–376.
- Luce, R. D. & Raiffa, H. (1957) *Games and Decisions*. New York, NY: Wiley.
- Maynard Smith, J. (1977) Parental investment: a prospective analysis. *Animal Behaviour*, **25**, 1–9.
- Maynard Smith, J. (1982) *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Maynard Smith, J. & Parker, G. A. (1976) The logic of asymmetric contests. *Animal Behaviour*, **24**, 159–177.
- Maynard Smith, J. & Price, G. R. (1973) The logic of animal conflict. *Nature*, **246**, 15–18.
- McGill, B. J. & Brown, J. S. (2007) Evolutionary game theory and adaptive dynamics of continuous traits. *Annual Review of Ecology, Evolution and Systematics*, **38**, 403–435.
- McNamara, J. M. (1993) Evolutionary paths in strategy space: an improvement algorithm for life-history strategies. *Journal of Theoretical Biology*, **161**, 23–37.
- McNamara, J. M. & Houston, A. I. (1996) State dependent life histories. *Nature*, **380**, 215–221.
- McNamara, J. M. & Houston, A. I. (2002) Credible threats and promises. *Philosophical Transactions of the Royal Society B*, **357**, 1607–1616.
- McNamara, J. M., Webb, J. N., Collins, E. J., Székely, T. & Houston A. I. (1997) A general technique for computing evolutionarily stable strategies based on errors in decision-making. *Journal of Theoretical Biology*, **189**, 211–225.
- McNamara, J. M., Gasson, C. E. & Houston, A. I. (1999) Incorporating rules for responding into evolutionary games. *Nature*, **401**, 368–371.
- McNamara, J. M., Barta, Z. & Houston, A. I. (2004) Variation in behaviour promotes cooperation in the prisoner's dilemma game. *Nature*, **428**, 745–748.
- McNamara, J. M., Wilson, E. & Houston, A. I. (2006) Is it better to give information, receive it or be ignorant in a two-player game? *Behavioral Ecology*, **17**, 441–451.
- McNamara, J. M., Barta, Z., Fromhage, L. & Houston, A. I. (2008) The coevolution of choosiness and cooperation. *Nature*, **451**, 189–192.
- Metz, J. A. J., Nisbet, R. M. & Geritz, S. A. H. (1992) How should we define 'fitness' for general ecological scenarios? *Trends in Ecology and Evolution*, **7**, 198–202.
- Moran, P. A. P. (1964) On the nonexistence of adaptive topographies. *Annals of Human Genetics*, **27**, 338–343.
- Mylius, S. D. & Diekmann (1995) On evolutionarily stable life histories, optimization and the need to be specific about density dependence. *Oikos*, **74**, 218–224.
- Nash, J. (1950) Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the USA*, **36**, 48–49.
- Nagylaki, T., Hofbauer J. & Brunovský, P. (1999) Convergence of multilocus systems under weak epistasis or weak selection. *Journal of Mathematical Biology*, **38**, 103–133.
- Pen, I. & Taylor, P. D. (2005) Modelling information exchange in worker-queen conflict over sex allocation. *Proceedings of the Royal Society B*, **272**, 2403–2408.
- Pen, I. & Weissing, F. J. (2000) Towards a unified theory of cooperative breeding: the role of ecology and life history re-examined. *Proceedings of the Royal Society B*, **267**, 2411–2418.
- Pen, I. & Weissing, F. J. (2002) Optimal sex allocation: steps towards a mechanistic theory. In: *Sex Ratios: Concepts and Research Methods*, ed. I. C. W. Hardy. Cambridge: Cambridge University Press, pp. 26–45.
- Pen, I., Daan, S. & Weissing, F. J. (1999) Seasonal sex ratio trend in the European kestrel: an evolutionarily stable strategy analysis. *American Naturalist*, **153**, 384–397.
- Persky, J. (1995) Retrospectives: the ethology of *Homo economicus*. *Journal of Economic Perspectives*, **9**, 221–231.

- Reeve, H. K. & Dugatkin, L. A., eds. (1998) *Game Theory and Animal Behavior*. Oxford: Oxford University Press.
- Schelling, T. C. (1960) *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Seger, J. & Stubblefield, J. W. (2002) Models of sex ratio evolution. In: *Sex Ratios: Concepts and Research Methods*, ed. I. C. W. Hardy. Cambridge: Cambridge University Press, pp. 1–25.
- Selten, R. (1980) A note on evolutionarily stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology*, **84**, 93–101.
- Selten, R. (1983) Evolutionary stability in extensive two-person games. *Mathematical Social Sciences*, **5**, 269–363.
- Spence, A. M. (1973) Job market signaling. *Quarterly Journal of Economics*, **87**, 355–374.
- Taper, M. L. & Case, T. J. (1992) Models of character displacement and the theoretical robustness of taxon cycles. *Evolution*, **46**, 317–333.
- Taylor, P. D. (1996) Inclusive fitness arguments in genetic models of behaviour. *Journal of Mathematical Biology*, **34**, 654–674.
- Taylor, P. D. & Frank S. A. (1996) How to make a kin selection model. *Journal of Theoretical Biology*, **180**, 27–37.
- Trivers, R. L. (1985) *Social Evolution*. Menlo Park, CA: Benjamin/Cummings.
- Trivers, R. L. & Willard, D. E. (1973) Natural selection on parental ability to vary the sex ratio of offspring. *Science*, **179**, 90–92.
- van Boven, M. & Weissing, F. J. (2004) The evolutionary economics of immunity. *American Naturalist*, **163**, 277–294.
- van de Pol, M., Pen, I., Heg, D. & Weissing, F. J. (2007) Variation in habitat choice and delayed reproduction: adaptive queuing strategies or individual quality differences? *American Naturalist*, **170**, 530–541.
- van Damme, E. (1987) *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag.
- van Doorn, G. S. & Dieckmann, U. (2006) The long-term evolution of multilocus traits under frequency-dependent disruptive selection. *Evolution*, **60**, 2226–2238.
- van Doorn, G. S. & Weissing, F. J. (2006) Sexual conflict and the evolution of female preferences for indicators of male quality. *American Naturalist*, **168**, 743–757.
- van Doorn, G. S., Hengeveld, G. M. & Weissing, F. J. (2003a) The evolution of social dominance. I. Two-player models. *Behaviour*, **140**, 1305–1332.
- van Doorn, G. S., Hengeveld, G. M. & Weissing, F. J. (2003b) The evolution of social dominance. II. Multi-player models. *Behaviour*, **140**, 1333–1358.
- van Doorn, G. S., Dieckmann, U. & Weissing, F. J. (2004) Sympatric speciation by sexual selection: a critical re-evaluation. *American Naturalist*, **163**, 709–725.
- von Neumann, J. & Morgenstern, O. (1944) *Theory of Games and Economic Behaviour*. Princeton, NJ: Princeton University Press.
- Webb, J. N., Houston, A. I., McNamara, J. M. & Székely, T. (1999) Multiple patterns of parental care. *Animal Behaviour*, **58**, 983–999.
- Weissing, F. J. (1996) Genetic versus phenotypic models of selection: can genetics be neglected in a long-term perspective? *Journal of Mathematical Biology*, **34**, 533–555.
- Wolf, M., van Doorn, G. S., Leimar, O. & Weissing, F. J. (2007) Life-history trade-offs favour the evolution of animal personalities. *Nature*, **447**, 581–584.
- Wolf, M., van Doorn, G. S., Leimar, O. & Weissing, F. J. (2008) Do animal personalities emerge? *Nature*, **451**, E9–10.
- Wright, S. (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress of Genetics*, pp. 355–366.
- Zahavi, A. & Zahavi, A. (1997) *The Handicap Principle: a Missing Piece of Darwin's Puzzle*. Oxford: Oxford University Press.