

again and again. The players know that they interact only once, so that retaliation is impossible.

Snowdrift: two players each receive an endowment, on provision that they pay a fee to the experimenter that is lower than the endowment. They must decide whether they are willing to pay the fee or not, knowing that if both are willing, each of them pays only half.

Chapter Two

Game Dynamics and Social Learning

2.1 GAMES

It can be difficult to decide what is best. The task can be fraught with uncertainties (as when an investor wants to optimize a portfolio), or it can be computationally demanding (as when a traveling salesman has to find the shortest route through 87 towns). A peculiar complication arises in interactions between two (or more) decision-makers with different views about what is best. This is the realm of game theory.

As an example, consider two players I and II engaged in the following, admittedly childish game. At a given signal, each holds up one or two fingers. If the resulting sum is odd, player I wins. If the sum is even, player II wins. Clearly there is no outcome that can satisfy both players. One of the players would always prefer to switch to the other alternative. Situations with a similar structure abound in social interactions.

Let us formalize this. Suppose that player I has to choose between n options, or *strategies*, which we denote by e_1, \dots, e_n , and player II between m strategies f_1, \dots, f_m . If I chooses e_i and II chooses f_j , then player I obtains a *payoff* a_{ij} and player II obtains b_{ij} . The game, then, is described by two $n \times m$ payoff matrices A and B : alternatively, we can describe it by one matrix whose element, in the i -th row and j -th column, is the pair (a_{ij}, b_{ij}) of payoff values. The payoff is measured on a utility scale consistent with the players' preferences. In biological games, it can be some measure of Darwinian fitness reflecting reproductive success. For simplicity, we stick to monetary payoffs.

The two players in our example could bet, say, one dollar. Each player has two strategies, *even* and *odd*, which correspond to e_1 and e_2 for player I and f_1 and f_2 for player II, and the payoff matrix is

$$\begin{pmatrix} (-1, 1) & (1, -1) \\ (1, -1) & (-1, 1) \end{pmatrix}. \quad (2.1)$$

If the outcome is $(-1, 1)$, player I (who chooses the row of the payoff matrix) would have done better to choose the other strategy; if the outcome is $(1, -1)$, it is player II, the column player, who would have done better to switch. If players could out-guess each other, they would be trapped in a vicious circle.




Figure 2.1 The simplices S_2 , S_3 , and S_4 . On the top row, S_2 and S_3 are imbedded in R^2 and R^3 respectively.

2.2 MIXED STRATEGIES

For both players, it is clearly important not to have their decision anticipated by the co-player. A good way to achieve this is to randomize, i.e., to let chance decide. Suppose that player I opts to play strategy e_i with probability x_i . This *mixed* strategy is thus given by a stochastic vector $\mathbf{x} = (x_1, \dots, x_n)$ (with $x_i \geq 0$ and $x_1 + \dots + x_n = 1$). We denote the set of all such mixed strategies by S_n : this is a simplex in R^n , spanned by the unit vectors e_i of the standard base, which are said to be the *pure* strategies, and correspond to the original set of alternatives, see figure 2.1. (All components of e_i are 0 except the i -th component, which is 1.)

Similarly, a mixed strategy for player II is an element \mathbf{y} of the unit simplex S_m . If player I uses the pure strategy e_i and player II uses strategy \mathbf{y} , then the payoff for player I (or more precisely, its expected value) is

$$(A\mathbf{y})_i = \sum_{j=1}^m a_{ij} y_j. \quad (2.2)$$

If player I uses the mixed strategy \mathbf{x} , and II uses \mathbf{y} , the payoff for player I is

$$\mathbf{x} \cdot A\mathbf{y} = \sum_i x_i (A\mathbf{y})_i = \sum_{i,j} a_{ij} x_i y_j, \quad (2.3)$$

and the payoff for player II, similarly, is

$$\mathbf{x} \cdot B\mathbf{y} = \sum_{i,j} b_{ij} x_i y_j. \quad (2.4)$$

(The dot on the left hand side denotes the dot product, or Euclidean product, of two vectors.)

If player I knows, by any chance, the strategy \mathbf{y} of the co-player, then player I should use a strategy that is a *best reply* to \mathbf{y} . The set of best replies is the set

$$BR(\mathbf{y}) = \arg \max_{\mathbf{x}} \mathbf{x} \cdot A\mathbf{y}, \quad (2.5)$$

i.e., the set of all $\mathbf{x} \in S_n$ such that

$$\mathbf{z} \cdot A\mathbf{y} \leq \mathbf{x} \cdot A\mathbf{y} \quad (2.6)$$

holds for all $\mathbf{z} \in S_n$. Player I has no incentive to deviate from \mathbf{x} .

Since the function $\mathbf{z} \mapsto \mathbf{z} \cdot A\mathbf{y}$ is continuous and S_n is compact, the set of best replies is always non-empty. It is a convex set. Moreover, if \mathbf{x} belongs to $BR(\mathbf{y})$, so do all pure strategies in the *support* of \mathbf{x} , i.e., all e_i for which $x_i > 0$. Indeed, for all i ,

$$(A\mathbf{y})_i = e_i \cdot A\mathbf{y} \leq \mathbf{x} \cdot A\mathbf{y}. \quad (2.7)$$

If the inequality sign were strict for some i with $x_i > 0$, then $x_i (A\mathbf{y})_i < x_i (\mathbf{x} \cdot A\mathbf{y})$; summing over all $i = 1, \dots, n$ then leads to a contradiction. It follows that the set $BR(\mathbf{y})$ is a face of the simplex S_n . It is spanned by the pure strategies which are best replies to \mathbf{y} .

2.3 NASH EQUILIBRIUM

If player I has found a best reply to the strategy \mathbf{y} of player II, player I has no reason not to use it—as long as player II sticks to \mathbf{y} .

But will player II stick to \mathbf{y} ? Only if player II has no incentive to use another strategy, i.e., has also hit upon a best reply. Two strategies \mathbf{x} and \mathbf{y} are said to form a *Nash equilibrium* pair if each is a best reply to the other, i.e., if $\mathbf{x} \in BR(\mathbf{y})$ and $\mathbf{y} \in BR(\mathbf{x})$, or alternatively if

$$\mathbf{z} \cdot A\mathbf{y} \leq \mathbf{x} \cdot A\mathbf{y} \quad (2.8)$$

holds for all $\mathbf{z} \in S_n$, and

$$\mathbf{x} \cdot B\mathbf{w} \leq \mathbf{x} \cdot B\mathbf{y} \quad (2.9)$$

holds for all $\mathbf{w} \in S_m$. A Nash equilibrium pair (\mathbf{x}, \mathbf{y}) satisfies a minimal consistency requirement: no player has an incentive to deviate (as long as the other player does not deviate either).

A basic result states that Nash equilibrium pairs always exist for any game (A, B) . This will be proved in section 2.11. The result holds for vastly wider classes of games than considered so far; it holds for any number of players, any convex compact sets of strategies, any continuous payoff functions, and even beyond. But it would not hold if we had not allowed for mixed strategies: this is shown by the simple example from section 2.1 (betting on even or odd). In that case, the mixed strategies of choosing, with equal probability 1/2, an even or an odd number, obviously lead to an equilibrium pair: each player gains, on average, zero dollars, and none has an

incentive to deviate. On the other hand, if player I uses any other strategy (x_1, x_2) against the $(1/2, 1/2)$ of player II, player I would still have an expected payoff of zero. However, the *other* player would then have an incentive to deviate: whenever $x_1 > x_2$, the best reply for II would be to play $(1, 0)$. If player II did that, however, player I would do better to play $(0, 1)$, and the vicious circle would be in full swing.

In this example, (x, y) with $x = y = (1/2, 1/2)$ is the unique Nash equilibrium pair. We have seen that as long as player II chooses the equilibrium strategy y , player I has no reason to deviate from the equilibrium strategy x ; but that on the other hand, player I has no reason not to deviate, either. This would be different if (x, y) were a *strict* Nash equilibrium pair, i.e., if

$$z \cdot Ay < x \cdot Ay \quad (2.10)$$

holds for all $z \neq x$, and

$$x \cdot Bw < x \cdot By \quad (2.11)$$

holds for all $w \neq y$. In this case, i.e., when both best-reply sets are singletons, each player will be penalized for unilaterally deviating from the equilibrium.

Whereas every game admits a Nash equilibrium pair, many games admit no *strict* Nash equilibrium pair; the number game from section 2.1 is an example.

Moreover, even if there exists a strict Nash equilibrium, it can be a let-down, as the Prisoner's Dilemma example from section 1.3 shows. This game has a unique Nash equilibrium, which is strict: both players defect, i.e., $x = y = (0, 1)$. Each player, in that case, would be penalized for deviating unilaterally. If both players, however, were to deviate, and opt for cooperation, they would be better off.

A further caveat applies: for many games, there exists not one, but several equilibrium pairs. Which one should the players choose? They could, of course, sit down and talk it over, but this is not considered a solution. In many cases players cannot communicate—sometimes this is prohibited by explicit rules, and sometimes it is just a waste of breath. Consider the Snowdrift game from section 1.4, for instance. In that case, it is easy to see that (e_1, f_2) and (e_2, f_1) are two Nash equilibrium pairs. They look similar to a bystander, but certainly not to the players themselves. The strategy pair (e_1, f_2) means that player I shovels and player II leans back and relaxes. Player I will not like this, but has no incentive not to shovel—for refusing to shovel means spending the night in the car. Of course player I would prefer the other Nash equilibrium pair. But if player I aims at that other equilibrium, and consequently uses strategy e_2 , while player II stubbornly clings to the strategy f_2 corresponding to the equilibrium pair which is better for II, then both players end up with the strategy pair (e_2, f_2) (an uncomfortably cold night in the car), which is not a Nash equilibrium pair.

2.4 POPULATION GAMES

So far, we have considered games between two specific players trying to guess each other's strategy and find a best reply. Let us now shift perspective, and consider a *population* of players, each with a given strategy. From time to time, two players

meet randomly and play the game, using their strategies. We shall consider these strategies as behavioral programs. Such programs can be learned, or inherited, or imprinted in any other way. In a biological setting, strategies correspond to different types of individuals (or behavioral phenotypes).

In order to analyze this set-up, it is convenient to assume that all individuals in the population are indistinguishable, except in their way of interacting, i.e., that the players differ only by their type, or strategy. This applies well to certain games such as the Prisoner's Dilemma, where both players are on an equal footing; but for many other examples of social interactions, there is an inherent asymmetry—for instance, between buyers and sellers.

For simplicity, we start by considering only *symmetric* games. In the case of two-player games, this means that the game remains unchanged if I and II are permuted. In particular, the two players have the same set of strategies. Hence we assume that $n = m$ and $f_j = e_j$ for all j ; and if a player plays strategy e_i against someone using strategy e_j (which is the former f_j), then that player receives the same payoff, whether labeled I or II. Hence $a_{ij} = b_{ji}$, or in other words $B = A^T$ (the transpose of matrix A). Thus a symmetric game is specified by the pair (A, A^T) , and therefore is defined by a single, square payoff matrix A .

As we have seen with the Snowdrift example, a symmetric game can have asymmetric Nash equilibrium pairs. These are plainly irrelevant, as long as it is impossible to distinguish players I and II. Of interest are only symmetric Nash equilibrium pairs, i.e., pairs of strategies (x, y) with $x = y$. A *symmetric Nash equilibrium*, thus, is specified by *one* strategy x having the property that it is a best reply to itself (i.e., $x \in BR(x)$). In other words, we must have

$$z \cdot Ax \leq x \cdot Ax \quad (2.12)$$

for all $z \in S_n$. A symmetric *strict* Nash equilibrium is accordingly given by the condition

$$z \cdot Ax < x \cdot Ax \quad (2.13)$$

for all $z \neq x$.

We shall soon prove that every symmetric game admits a symmetric Nash equilibrium.

2.5 SYMMETRIZING A GAME

There is an obvious way to turn a non-symmetric game (A, B) into a symmetric game: simply by letting a coin toss decide who of the two players will be labeled player I. A strategy for this *symmetrized* game must therefore specify what to do in role I, and what in role II, i.e., such a strategy is given by a pair (e_i, f_j) . A mixed strategy is given by an element $z = (z_{ij}) \in S_{nm}$, where z_{ij} denotes the probability to play e_i when in role I and f_j when in role II. To the probability distribution z correspond its *marginals*: $x_i = \sum_j z_{ij}$ and $y_j = \sum_i z_{ij}$. The vectors $x = (x_i)$ and $y = (y_j)$ belong to S_n and S_m , respectively. It is easy to see that for any given $x \in S_n$ and $y \in S_m$ there exists a $z \in S_{nm}$ having x and y as marginals, for instance

$z_{ij} = x_i y_j$ (barring exceptions, there exist many probability distributions with the same marginals).

The payoff for a player using (e_i, f_j) against a player using (e_k, f_l) , with $i, k \in \{1, \dots, n\}$ and $j, l \in \{1, \dots, m\}$, depends on the outcome of the coin-toss and is given by

$$c_{ij,kl} = \frac{1}{2}a_{il} + \frac{1}{2}b_{kj}. \quad (2.14)$$

Let us briefly take for granted that every symmetric game has a symmetric Nash equilibrium. Then it can easily be deduced that every game (A, B) has a Nash equilibrium pair.

Indeed, let us assume that $\bar{z} \in S_{nm}$ is a symmetric Nash equilibrium for the symmetrized game (C, C^T) . This means that

$$z \cdot C\bar{z} \leq \bar{z} \cdot C\bar{z} \quad (2.15)$$

for all $z \in S_{nm}$. Let x and y be the marginals of z , and \bar{x} and \bar{y} the marginals for \bar{z} . Then

$$z \cdot C\bar{z} = \sum_{ijkl} z_{ij} c_{ij,kl} \bar{z}_{kl} \quad (2.16)$$

$$= \frac{1}{2} \sum_{ijkl} z_{ij} a_{il} \bar{z}_{kl} + \frac{1}{2} \sum_{ijkl} z_{ij} b_{kj} \bar{z}_{kl} \quad (2.17)$$

$$= \frac{1}{2} \sum_{il} x_i a_{il} \bar{y}_l + \frac{1}{2} \sum_{jk} y_j b_{kj} \bar{x}_k = \frac{1}{2} x \cdot A\bar{y} + \frac{1}{2} y \cdot B^T \bar{x}. \quad (2.18)$$

Since \bar{z} is a symmetric Nash equilibrium, equation (2.15) implies

$$x \cdot A\bar{y} + \bar{x} \cdot B y \leq \bar{x} \cdot A\bar{y} + \bar{x} \cdot B\bar{y}. \quad (2.19)$$

For $y = \bar{y}$ this yields

$$x \cdot A\bar{y} \leq \bar{x} \cdot A\bar{y}, \quad (2.20)$$

and for $x = \bar{x}$,

$$\bar{x} \cdot B y \leq \bar{x} \cdot B\bar{y}. \quad (2.21)$$

Hence $\bar{x} \in BR(\bar{y})$ and $\bar{y} \in BR(\bar{x})$, i.e., (\bar{x}, \bar{y}) is a Nash equilibrium pair of the game (A, B) .

2.6 POPULATION DYNAMICS MEETS GAME THEORY

We now consider a symmetric game with payoff matrix A and assume that in a large, well-mixed population, a fraction x_i uses strategy e_i , for $i = 1, \dots, n$. The state of the population is thus given by the vector $x \in S_n$. A player with strategy e_i has as expected payoff

$$(Ax)_i = \sum_j a_{ij} x_j. \quad (2.22)$$

Indeed, this player meets with probability x_j a co-player using e_j . The average payoff in the population is given by

$$x \cdot Ax = \sum_i x_i (Ax)_i. \quad (2.23)$$

It should be stressed that we are committing an abuse of notation. The same symbol $x \in S_n$ which denoted in the previous sections the mixed strategy of one specific player now denotes the state of a population consisting of different types, each type playing its pure strategy. (We could also have the players use mixed strategies, but there will be no need to consider this case.)

Now comes an essential step: we shall assume that populations can evolve, in the sense that the frequencies x_i change with time. Thus we let the state $x(t)$ depend on time, and denote by $\dot{x}_i(t)$ the velocity with which x_i changes, i.e., $\dot{x}_i = dx_i/dt$. In keeping with our population dynamical approach, we shall be particularly interested in the (per capita) growth rates \dot{x}_i/x_i of the frequencies of the strategies.

How do the frequencies of strategies evolve? How do they grow and diminish? There are many possibilities for modeling this process. We shall mostly assume that the state of the population evolves according to the *replicator equation*. This equation holds if the growth rate of a strategy's frequency corresponds to the strategy's payoff, or more precisely to the difference between its payoff $(Ax)_i$ and the average payoff $x \cdot Ax$ in the population. Thus we posit

$$\dot{x}_i = x_i [(Ax)_i - x \cdot Ax] \quad (2.24)$$

for $i = 1, \dots, n$. Accordingly, a strategy e_i will spread or dwindle depending on whether it does better or worse than average.

This yields a deterministic model for the state of the population. Indeed, any *ordinary differential equation* $\dot{x} = F(x)$ with a smooth right hand side (such as eq. (2.24)) has a unique solution for each initial condition x , i.e., a function $t \mapsto x(t)$ from an open interval I (containing 0) into R^n such that $x(0) = x$ and such that $\dot{x}(t) = F(x(t))$ holds for all $t \in I$. For all differential equations that we consider in this book, the interval I can always be taken to be the whole real line R .

We may interpret the right hand side of the differential equation as a vector field $x \mapsto F(x)$. It associates to each point x in the domain of definition of F (an open subset B of R^n) the "wind velocity" $F(x) \in R^n$ at that point. The solution then describes the motion of a particle, released at time 0 at x and carried along by the wind. At a point z such that $F(z) = 0$, the velocity is zero. This corresponds to a *rest point*: a particle released at z will not move. We note that multiplying the right hand side $F(x)$ by a positive function $M(x) > 0$ corresponds to a *change in velocity*. The particle will then travel with a different speed, but along the same orbit.

Before we try to explain (in section 2.7) why we are interested in equation (2.24), let us note that $\sum \dot{x}_i = 0$. Furthermore, it is easy to see that the constant function $x_i(t) = 0$ for all t obviously satisfies the i -th equation in (2.24). From this follows that the state space, i.e., the simplex S_n , is invariant: if $x(0) \in S_n$ then $x(t) \in S_n$ for all $t \in R$. The same holds for all sub-simplices of S_n , (which are given by $x_i = 0$ for one or several i), and hence also for the boundary bdS_n of S_n (i.e., the union of

all such sub-simplices), and moreover also for the interior $\text{int } S_n$ of the simplex (the subset satisfying $x_i > 0$ for all i).

2.7 IMITATION DYNAMICS

The replicator equation initially showed up in the context of biological games. The assumption that payoff corresponds to reproductive success, and that individuals breed true, leads almost immediately to this equation. Clearly, for the economic games we are considering here, strategies are unlikely to be inherited, but they can be transmitted through social learning. If we assume that individuals imitate each other, we meet the replicator equation again.

To be more precise, let us assume that from time to time, a randomly chosen individual randomly samples a model from the population and imitates that model with a certain likelihood. Thus the probability that during an interval of length Δt , an individual switches from strategy e_j to e_i is given by $x_i f_{ij} \Delta t$. The corresponding input-output equation is

$$x_i(t + \Delta t) - x_i(t) = \sum_j f_{ij} x_i x_j \Delta t - \sum_j f_{ji} x_i x_j \Delta t, \quad (2.25)$$

which in the limit $\Delta t \rightarrow 0$ yields

$$\dot{x}_i = x_i \sum_j (f_{ij} - f_{ji}) x_j. \quad (2.26)$$

In general, the rates f_{ij} will depend on the state \mathbf{x} . For instance, we can assume that

$$f_{ij} = [(A\mathbf{x})_i - (A\mathbf{x})_j]_+. \quad (2.27)$$

This means that an e_j player comparing himself with an e_i player will adopt the latter's strategy only if it promises a higher payoff: and if this is the case, the switch is more likely if the difference in payoff is higher. In that case, since $f_{ij} - f_{ji} = (A\mathbf{x})_i - (A\mathbf{x})_j$, the input-output equation yields

$$\dot{x}_i = x_i \sum_j [(A\mathbf{x})_i - (A\mathbf{x})_j] x_j = x_i [(A\mathbf{x})_i - \mathbf{x} \cdot A\mathbf{x}], \quad (2.28)$$

which is just the replicator equation (2.24). We would obtain it in a similar way if, instead of the payoff $(A\mathbf{x})_i$, we use a more general "fitness" term measuring the success of a strategy, for instance $(1-s)B + s(A\mathbf{x})_i$, with $0 < s \leq 1$. This is the convex combination of a "baseline fitness" $B = B(\mathbf{x}) > 0$ (the same for all types) and the payoff. The size of s specifies the importance of the game in evaluating the "appeal" of a strategy.

We could also assume that

$$f_{ij} = (1-s)B + s(A\mathbf{x})_i, \quad (2.29)$$

which means that the switching rate depends only on the success of the model (and not on the payoff of the e_j player); or that

$$f_{ij} = (1-s)B - s(A\mathbf{x})_j, \quad (2.30)$$

which means that players are all the more prone to imitate one another the more reason they have to be dissatisfied with their own payoff. The role of the convex combination is to guarantee that (at least for small s) the rate is positive.

Not every imitation mechanism leads to the replicator equation. For instance, we could assume that if two players compare their payoffs, the better will always be imitated by the worse. Thus $f_{ij} = 0$ if $(A\mathbf{x})_i < (A\mathbf{x})_j$, $f_{ij} = 1$ if $(A\mathbf{x})_i > (A\mathbf{x})_j$, and $f_{ij} = \frac{1}{2}$, say, in the case of a tie. This leads to a differential equation with a discontinuous right hand side. The dynamics reduces, incidentally, to that of a replicator equation in every region of the state space defined by a specific ordering of the payoff values $(A\mathbf{x})_i$.

Not all learning is social learning (i.e., learning from others). We can also learn from our own experience, for instance by using mostly those strategies that have brought success so far. Moreover, social learning could disregard the success of a model, for instance, by simply imitating whatever is most frequent.

It is worth emphasizing that imitation (like selection, in genetics) does not produce anything new. If a strategy e_i is absent from the population, it will remain so (i.e., if $x_i(t) = 0$ holds for some time t , it holds for all t). There exist game dynamics that are more innovative. For instance, clever players could adopt the strategy that offers the highest payoff, even if no one in the population is currently using it. Other innovative dynamics arise if we assume a steady rate of switching randomly to other strategies. This can be interpreted as an "exploration rate," and corresponds to a mutation term in genetics.

2.8 BASIC PROPERTIES OF THE REPLICATOR EQUATION

It is easy to see that if we add an arbitrary function $f(\mathbf{x})$ to all payoff terms $(A\mathbf{x})_i$, the replicator equation (2.24) remains unchanged: what is added to the payoff is also added to the average payoff $\mathbf{x} \cdot A\mathbf{x}$, since $\sum x_i = 1$, and cancels out in the difference of the two terms. In particular, this implies that we can add a constant c_j to the j -th column of A (for $j = 1, \dots, n$) without altering the replicator dynamics in S_n . We shall frequently use this to simplify the analysis.

Another useful property is the quotient rule: if $x_j > 0$, then the time-derivative of the quotient satisfies

$$\left(\frac{x_i}{x_j} \right)' = \left(\frac{x_i}{x_j} \right) [(A\mathbf{x})_i - (A\mathbf{x})_j]. \quad (2.31)$$

More generally, if $V = \prod x_i^{p_i}$ then

$$\dot{V} = V \left[\mathbf{p} \cdot A\mathbf{x} - \left(\sum p_i \right) \mathbf{x} \cdot A\mathbf{x} \right]. \quad (2.32)$$

The rest points \mathbf{z} of the replicator equation are those for which all payoff values $(A\mathbf{z})_i$ are equal, for all indices i for which $z_i > 0$. The common value of these payoffs is the average payoff $\mathbf{z} \cdot A\mathbf{z}$. In particular, all vertices e_i of the simplex S_n are rest points. (Obviously, if all players are of the same type, imitation leads to no change.)

The replicator equation admits a rest point in $\text{int}S_n$ if there exists a solution (in $\text{int}S_n$) of the linear equations

$$(Ax)_1 = \dots = (Ax)_n. \quad (2.33)$$

Similarly, all rest points on each face can be obtained by solving a corresponding system of linear equations. Typically, each sub-simplex (and S_n itself) contains one or no rest point in its interior.

One can show that if no rest point in S_n exists in the interior of S_n , then all orbits in $\text{int}S_n$ converge to the boundary, for $t \rightarrow \pm\infty$. In particular, if strategy e_i is *strictly dominated*, i.e., if there exists a $w \in S_n$ such that $(Ax)_i < w \cdot Ax$ holds for all $x \in S_n$, then $x_i(t) \rightarrow 0$ for $t \rightarrow +\infty$. In the converse direction, if there exists an orbit $x(t)$ bounded away from the boundary of S_n (i.e., such that for some $a > 0$ the inequality $x_i(t) > a$ holds for all $t > 0$ and all $i = 1, \dots, n$), then there exists a rest point in $\text{int}S_n$. One just has to note that for $i = 1, \dots, n$,

$$(\log x_i)' = \dot{x}_i/x_i = (Ax(t))_i - x(t) \cdot Ax(t). \quad (2.34)$$

Integrating this from 0 to T , and dividing by T , leads on the left hand side to $[\log x_i(T) - \log x_i(0)]/T$, which converges to 0 for $T \rightarrow +\infty$. The corresponding limit on the right hand side implies that for the accumulation points z_i of the time averages

$$z_i(T) = \frac{1}{T} \int_0^T x_i(t) dt, \quad (2.35)$$

the relations $z_i \geq a > 0$, $\sum z_i = 1$, and

$$\sum a_{1j} z_j = \dots = \sum a_{nj} z_j \quad (2.36)$$

must hold. Thus z is a rest point in $\text{int}S_n$.

2.9 THE CASE OF TWO STRATEGIES

Let us discuss the replicator equation when there are only two types in the population. Since the equation remains unchanged if we subtract the diagonal term in each column, we can assume without restricting generality that the 2×2 matrix A is of the form

$$\begin{pmatrix} 0 & a \\ b & 0 \end{pmatrix}. \quad (2.37)$$

Since $x_2 = 1 - x_1$, it is enough to observe x_1 , which we denote by x . Thus $x_2 = 1 - x$, and

$$\dot{x} = x[(Ax)_1 - x \cdot Ax] = x[(Ax)_1 - (x(Ax)_1 + (1-x)(Ax)_2)], \quad (2.38)$$

and hence

$$\dot{x} = x(1-x)[(Ax)_1 - (Ax)_2]. \quad (2.39)$$




Figure 2.2 Classification of the replicator dynamics for $n=2$: (a) inertia; (b) dominance; (c) bi-stability; (d) stable coexistence. Circles denote rest points. Filled circles correspond to stable rest points.

Since $(Ax)_1 = a(1-x)$ and $(Ax)_2 = bx$, (2.39) reduces to

$$\dot{x} = x(1-x)[a - (a+b)x]. \quad (2.40)$$

We note that

$$a = \lim_{x \rightarrow 0} \frac{\dot{x}}{x}. \quad (2.41)$$

Hence a corresponds to the limit of the per capita growth rate of the missing strategy e_1 . Alternatively,

$$a = \frac{d\dot{x}}{dx}, \quad (2.42)$$

where the derivative is evaluated at $x=0$.

Let us omit the trivial case $a=b=0$: in this case all points of the state space S_2 (i.e., the interval $0 \leq x \leq 1$) are rest points. The right hand side of our differential equation is a product of three factors, the first vanishing at 0 and the second at 1; the third factor has a zero $\hat{x} = \frac{a}{a+b}$ in the open interval $]0, 1[$ if and only if $ab > 0$. Thus we obtain three possible cases, see figure 2.2:

1. There is no fixed point in the interior of the state space. This happens if and only if $ab \leq 0$. In this case, \dot{x} always has the same sign in $]0, 1[$. If this sign is positive (i.e., if $a \geq 0$ and $b \leq 0$, at least one inequality being strict,) this means that $x(t) \rightarrow 1$ for $t \rightarrow +\infty$, for every initial value $x(0)$ with $0 < x(0) < 1$. The strategy e_1 is said to *dominate* strategy e_2 . It is always the best reply, for any value of $x \in]0, 1[$. Conversely, if the sign of \dot{x} is negative, then $x(t) \rightarrow 0$ and e_2 dominates. In each case, the dominating strategy converges towards fixation.

As an example, we consider the Prisoner's Dilemma game from section 1.3. The payoff matrix is transformed into

$$\begin{pmatrix} 0 & -5 \\ 5 & 0 \end{pmatrix} \quad (2.43)$$

and defection dominates.

2. There exists a rest point \hat{x} in $]0, 1[$ (i.e., $ab > 0$), and both a and b are negative. In this case $\dot{x} < 0$ for $x \in]0, \hat{x}[$ and $\dot{x} > 0$ for $x \in]\hat{x}, 1[$. This means that the orbits lead away from \hat{x} : this rest point is unstable. As in the previous case, one strategy will be eliminated: but the outcome, in this *bistable* case, depends on the initial condition. If x is larger than the threshold \hat{x} , it will keep growing; if it is smaller, it will vanish—a positive feedback.

As an example, we can consider the repeated Prisoner's Dilemma from section 1.5. The payoff matrix is transformed into

$$\begin{pmatrix} 0 & -5 \\ -45 & 0 \end{pmatrix} \quad (2.44)$$

and it is best to play TFT if the frequency of TFT-players exceeds 10 percent.

3. There exists a rest point \hat{x} in $]0, 1[$ (i.e., $ab > 0$), and both a and b are positive. In this case $\dot{x} > 0$ for $x \in]0, \hat{x}[$ and $\dot{x} < 0$ for $x \in]\hat{x}, 1[$. This negative feedback means that $x(t)$ converges towards \hat{x} , for $t \rightarrow +\infty$: the rest point \hat{x} is a stable attractor. No strategy eliminates the other: rather, their frequencies converge towards a *stable coexistence*.

This example can be found in the Snowdrift game from section 1.4. The payoff matrix is transformed into

$$\begin{pmatrix} 0 & 10 \\ 15 & 0 \end{pmatrix} \quad (2.45)$$

and the fixed point corresponds to 40 percent helping and 60 percent shirking.

These three cases (dominance, bi-stability and stable coexistence) will be revisited in the next section. But first, we relate the replicator dynamics to the Nash equilibrium concept.

2.10 NASH EQUILIBRIA AND SATURATED REST POINTS

Let us consider a symmetric $n \times n$ game (A, A^T) with a symmetric Nash equilibrium z . This means that

$$x \cdot Az \leq z \cdot Az \quad (2.46)$$

for all $x \in S_n$. With $x = e_i$, this implies

$$(Az)_i \leq z \cdot Az \quad (2.47)$$

for $i = 1, \dots, n$. Equality must hold for all i such that $z_i > 0$, as we have seen in section 2.2. Hence z is a rest point of the replicator dynamics. Moreover, it is a

saturated rest point: this means by definition that if $z_i = 0$, then

$$(Az)_i - z \cdot Az \leq 0. \quad (2.48)$$

Conversely, every saturated rest point is a Nash equilibrium. The two concepts are equivalent.

Every rest point in $\text{int}S_n$ is trivially saturated; but on the boundary, there may be rest points that are not saturated, as we shall presently see. In that case, there exist strategies not present in the population z , that would do better than average (and better, in fact, than every type that is present). Rest points and Nash equilibria have in common that there exists a c such that $(Az)_i = c$ whenever $z_i > 0$; the additional requirement, for a Nash equilibrium, is that $(Az)_i \leq c$ whenever $z_i = 0$.

Hence every symmetric Nash equilibrium is a rest point, but the converse does not hold. Let us discuss this for the examples from the previous section. It is clear that the fixed points $\hat{x} \in]0, 1[$ are Nash equilibria. In case (1), the dominant strategy is a Nash equilibrium, and the other is not. In case (2), both pure strategies are Nash equilibria. In case (3), none of them is a Nash equilibrium. If you play a bi-stable game, you are well advised to choose the same strategy as your co-player; but in the case of stable coexistence, you should choose the opposite strategy. In both cases, however, the two of you might have different ideas about who plays what.

In the bi-stable case, which of the two pure equilibria, e_1 or e_2 , should be chosen? The first idea is: the one with the higher payoff (if it exists). This is said to be the *Pareto-optimal* outcome. In the example given in section 1.7, this is clearly the TFT strategy. The definition of Pareto-optimality depends on the actual payoff values, and is not specified by the replicator dynamics: after adding constants to every column of the payoff matrix, a different strategy may be Pareto-optimal.

The Pareto-optimal solution is not always convincing. Consider for instance the payoff matrix

$$\begin{pmatrix} 2 & -1000 \\ 0 & 1 \end{pmatrix} \quad (2.49)$$

Clearly, e_1 is Pareto-optimal. But will you play it against an unknown adversary? That player might be a fool, and choose e_2 . In that case, you would lose much. Obviously, e_2 is the safer Nash equilibrium. (And on second thought, your co-player may not be a fool, but just suspect that you might be one; or suspect that you might suspect, etc., . . .)

In a bi-stable game

$$\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \quad (2.50)$$

(with $\alpha > \gamma$ and $\delta > \beta$), the strategy e_1 is said to be *risk-dominant* if it provides the higher payoff against a co-player who is as likely to play e_1 as e_2 . This means that $(1/2)(\alpha + \beta) > (1/2)(\gamma + \delta)$, or

$$\gamma - \alpha < \beta - \delta. \quad (2.51)$$

This condition is invariant with respect to adding constants to each column, and implies for the normalized matrix (2.37) that $a > b$, i.e., (since both values are negative) $\hat{x} < 1/2$. Hence the risk-dominant equilibrium, in a bistable 2×2 game, is the one with the larger basin of attraction.

A handful of results about Nash equilibria and rest points of the replicator dynamics are known as *folk theorem of evolutionary game theory*. For instance, any limit, for $t \rightarrow +\infty$, of a solution $\mathbf{x}(t)$ starting in $\text{int}S_n$ is a Nash equilibrium; and any stable rest point is a Nash equilibrium. (A rest point \mathbf{z} is said to be stable if for any neighborhood U of \mathbf{z} there exists a neighborhood V of \mathbf{z} such that if $\mathbf{x}(0) \in V$ then $\mathbf{x}(t) \in U$ for all $t \geq 0$.) Both results are obvious consequences of the fact that if \mathbf{z} is not Nash, there exists an i and an ϵ such that $(A\mathbf{x})_i - \mathbf{x} \cdot A\mathbf{x} > \epsilon$ for all \mathbf{x} close to \mathbf{z} . In the other direction, if \mathbf{z} is a strict Nash equilibrium, then \mathbf{z} is an asymptotically stable rest point (i.e., not only stable, but in addition *attracting* in the sense that for some neighborhood U of \mathbf{z} , $\mathbf{x}(0) \in U$ implies $\mathbf{x}(t) \rightarrow \mathbf{z}$ for $t \rightarrow +\infty$). The converse statements are generally not valid.

2.11 EXISTENCE OF NASH EQUILIBRIA

In order to prove the existence of a symmetric Nash equilibrium for the symmetric game with $n \times n$ matrix A , i.e., the existence of a saturated rest point for the corresponding replicator equation (2.24), we perturb that equation by adding a small constant term $\epsilon > 0$ to each component of the right hand side. Of course, the relation $\sum \dot{x}_i = 0$ will no longer hold. We compensate this by subtracting the term $n\epsilon$ from each growth rate $(A\mathbf{x})_i - \mathbf{x} \cdot A\mathbf{x}$. Thus we consider

$$\dot{x}_i = x_i[(A\mathbf{x})_i - \mathbf{x} \cdot A\mathbf{x} - n\epsilon] + \epsilon. \quad (2.52)$$

Clearly, $\sum \dot{x}_i = 0$ is satisfied again. On the other hand, if $x_i = 0$, then $\dot{x}_i = \epsilon > 0$. This influx term changes the vector field of the replicator equation: at the boundary of S_n , (which is invariant for the unperturbed replicator equation), the vector field of the perturbed equation points towards the interior.

We shall see presently that (2.52) admits at least one rest point in $\text{int}S_n$, which we denote by \mathbf{z}_ϵ . It satisfies

$$(A\mathbf{z}_\epsilon)_i - \mathbf{z}_\epsilon \cdot A\mathbf{z}_\epsilon = \epsilon \left(n - \frac{1}{(z_\epsilon)_i} \right). \quad (2.53)$$

Let ϵ tend to 0, and let \mathbf{z} be an accumulation point of the \mathbf{z}_ϵ in S_n . The limit on the left hand side exists, it is $(A\mathbf{z})_i - \mathbf{z} \cdot A\mathbf{z}$. Hence the right hand side also has a limit for $\epsilon \rightarrow 0$. This limit is 0 if $z_i > 0$, and it is ≤ 0 if $z_i = 0$. This implies that \mathbf{z} is a saturated rest point of the (unperturbed) replicator equation (2.24), and hence corresponds to a Nash equilibrium.

All that remains to be shown is the existence of a rest point for equation (2.52). Readers who know Brouwer's fixed point theorem will need no proof. All others can find it in the next two sections.




Figure 2.3 (a) This is not a simplicial decomposition. (b) The barycentric decomposition is a simplicial decomposition.

2.12 SPERNER'S LEMMA

Let us consider an $n - 1$ -dimensional simplex S , i.e., the closed convex hull of n points $\mathbf{y}_1, \dots, \mathbf{y}_n$ such that the vectors $\mathbf{y}_i - \mathbf{y}_n$, $i = 1, \dots, n - 1$, are linearly independent. Any non-trivial subset of the vertices $\mathbf{y}_1, \dots, \mathbf{y}_n$ spans a *sub-simplex* of S . The boundary of S is the union of the n *full* (i.e., $n - 2$ -dimensional) faces. A *simplicial decomposition* of S consists of finitely many $n - 1$ -dimensional simplices whose union is S and whose interiors are pairwise disjoint. We furthermore require that if two such (closed) sub-simplices are not disjoint, they must share a face: i.e., if the intersection contains a vertex of one sub-simplex, that point is also a vertex of the other, see figure 2.3.

An example is the barycentric subdivision. (The *barycenter* of the simplex S is $(\mathbf{y}_1 + \dots + \mathbf{y}_n)/n$). We begin with the barycenters of all 1-dimensional sub-simplices, i.e., the midpoints of the edges. They divide the edges of S into 1-dimensional sub-simplices. We then introduce the barycenters of the 2-dimensional sub-simplices of S , and consider the 2-dimensional sub-simplex obtained as a convex hull of such a barycenter and a 1-dimensional sub-simplex on the boundary of the corresponding face; and so on into higher dimensions.

Now suppose that we are given a *coloring* of the vertices of the simplicial decomposition by n colors, in the sense that we associate to each vertex an $i \in \{1, \dots, n\}$. We require that the vertices \mathbf{y}_i of S are colored by the colors i , and that for any sub-simplex of S , only the colors of the vertices spanning that sub-simplex are used. We say that a sub-simplex is I -colored if $I \subset \{1, \dots, n\}$ is the list of all colors actually occurring at the vertices of that sub-simplex.

Sperner's lemma states that there always exists an odd number of $\{1, \dots, n\}$ -colored sub-simplices. (In particular, we need the full set of colors for at least one sub-simplex of S .)

The proof goes by induction. For $n = 2$ (i.e., for the segment S_2) the statement is obvious. Suppose it is proved up to $n - 1$. We can apply this to the boundary face of S , which is opposed to \mathbf{y}_n : its simplicial decomposition contains an odd number of sub-simplices which are $\{1, \dots, n - 1\}$ -colored.




Figure 2.4 A coloring of a simplicial decomposition of S_3 , and the graph described in section 2.12.

We now construct a graph having as vertices the barycenters of the sub-simplices of S . We join two such barycenters by an edge if and only if the corresponding sub-simplices share a $\{1, \dots, n-1\}$ -colored face, see figure 2.4. We add one further vertex o lying outside of S , and connect it with the barycenters of those sub-simplices having a $\{1, \dots, n-1\}$ -colored face on the boundary of S . We see immediately that o is connected to an odd number of barycenters, which belong to sub-simplices having a full face belonging to the face of S opposite of y_n .

If a sub-simplex is $\{1, \dots, n\}$ -colored, it has exactly one $\{1, \dots, n-1\}$ -colored face. Hence its barycenter lies on exactly one edge of the graph; it is an end-point of the graph. As to the other barycenters, they either lie on two edges, or on none at all. Indeed, if a sub-simplex that is not $\{1, \dots, n\}$ -colored has an $\{1, \dots, n-1\}$ -colored face, then the opposite vertex must have one of the colors $1, \dots, n-1$, and hence that sub-simplex has exactly one additional $\{1, \dots, n-1\}$ -colored face.

We note that it is possible that the graph has closed loops. But since an odd number of edges issues from o , there must be an odd number of end-points of the graph, and

hence an odd number of fully colored sub-simplices of S . Thus Sperner's lemma holds.

2.13 A FIXED-POINT THEOREM

We now show that a smooth vector field on the plane $\sum x_i = 1$ satisfying

$$\sum \dot{x}_i = 0 \quad (2.54)$$

and

$$x_i = 0 \Rightarrow \dot{x}_i > 0 \quad (2.55)$$

has a fixed point in $\text{int} S_n$. We proceed indirectly and assume that it has no fixed point. To each point $\mathbf{x} \in S_n$ we can associate the "color" $i := \min\{j : \dot{x}_j < 0\}$. This is possible because $\dot{\mathbf{x}} \neq \mathbf{0}$ and $\sum_j \dot{x}_j = 0$. We note that condition (2.55) implies that on each sub-simplex of S_n , only the colors of the vertices spanning that face are used.

This induces a coloring for any simplicial decomposition. Each such decomposition must have an odd number of fully-colored sub-simplices. Now consider a sequence of simplicial decompositions whose width (the size of the largest sub-simplex) converges to 0. (For instance, we can start with the barycentric subdivision of S , and then iterate this ad lib.)

This yields a sequence of fully-colored sub-simplices: by compactness, the sub-sequence converges to a point $\mathbf{x} \in S_n$. For each i , this point is a limit of i -colored vertices, and hence must satisfy $\dot{x}_i \leq 0$. Since $\sum \dot{x}_i = 0$ this implies $\dot{\mathbf{x}} = \mathbf{0}$, a contradiction.

Hence the vector field (2.52) must have some fixed point in S_n . This closes the gap in the proof that each replicator equation admits a saturated fixed point.

2.14 ROCK-SCISSORS-PAPER

Whereas there exist only four possible types of replicator dynamics for $n = 2$, there exist about a hundred of them for $n = 3$ (and for $n > 3$ a full classification seems presently out of sight). A particularly interesting example occurs if the three strategies dominate each other in a cyclic fashion, i.e., if e_1 dominates e_2 , in the absence of e_3 ; and similarly if e_2 dominates e_3 ; and e_3 , in turn, dominates e_1 . Such a cycle occurs in the game of Rock-Scissors-Paper. If we assume that the winner receives one dollar from the loser, the payoff matrix is

$$\begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}. \quad (2.56)$$

This is a zero-sum game: one player receives what the other player loses. Hence the average payoff in the population, $\mathbf{x} \cdot A\mathbf{x}$, is zero. There exist only four rest points, one in the center, $\mathbf{m} = (1/3, 1/3, 1/3) \in \text{int} S_3$, and the other three at the vertices e_i . The only Nash equilibrium is \mathbf{m} .




Figure 2.5 The replicator dynamics of the Rock-Scissors-Paper game with payoff matrix (2.58): (a) $a = 1$; (b) $a > 1$; (c) $0 < a < 1$.

Let us consider the function $V := x_1 x_2 x_3$, which is positive in the interior of S_3 (with its maximum at \mathbf{m}) and vanishes on the boundary. Using (2.32) we see that $t \rightarrow V(\mathbf{x}(t))$ satisfies

$$\dot{V} = V(x_2 - x_3 + x_3 - x_1 + x_1 - x_2) = 0. \quad (2.57)$$

Hence V is a *constant of motion*: all orbits $t \rightarrow \mathbf{x}(t)$ of the replicator equation remain on constant level sets of V , see figure 2.5. This implies that all orbits in $\text{int } S_n$ are closed orbits surrounding \mathbf{m} . The invariant set consisting of the three vertices \mathbf{e}_i and the orbits connecting them along the edges of S_3 is said to form a *heteroclinic set*. Any two points on it can be connected by “shadowing the dynamics.” This means to travel along the orbits of that set and, at appropriate times that can be arbitrarily rare, to make an arbitrarily small step. In the present case, it means for instance to flow along an edge towards \mathbf{e}_1 , and then step onto the edge leading away from \mathbf{e}_1 . This step can be arbitrarily small: travelers just have to wait until they are sufficiently close to the “junction” \mathbf{e}_1 .

Now let us consider the game with matrix

$$\begin{pmatrix} 0 & a & -1 \\ -1 & 0 & a \\ a & -1 & 0 \end{pmatrix}. \quad (2.58)$$

For $a > 0$, it has the same structure of cyclic dominance. For $a \neq 1$ the game is no longer a zero sum game, but it has the same rest points. The point \mathbf{m} is a Nash equilibrium and the boundary of S_3 is a heteroclinic set, as before. But now,

$$\mathbf{x} \cdot \mathbf{A}\mathbf{x} = (a - 1)(x_1 x_2 + x_2 x_3 + x_3 x_1) \quad (2.59)$$

and hence

$$\dot{V} = V(a - 1)[1 - 3(x_1 x_2 + x_2 x_3 + x_3 x_1)] \quad (2.60)$$

$$= \frac{V(a - 1)}{2} [(x_1 - x_2)^2 + (x_2 - x_3)^2 + (x_3 - x_1)^2]. \quad (2.61)$$

This expression vanishes on the boundary of S_3 and at \mathbf{m} . It has the sign of $a - 1$ everywhere else on S_3 . If $a > 1$, this means that all orbits cross the constant-level sets of V in the uphill direction, and hence converge to \mathbf{m} . This implies that ultimately,

all three types will be present in the population in equal frequencies: the rest point \mathbf{m} is asymptotically stable, see figure 2.5b. But for $a < 1$, the orbits flow downhill, towards the boundary of S_3 . The Nash equilibrium \mathbf{m} corresponds to an unstable rest point, and the heteroclinic cycle on the boundary attracts all other orbits, see figure 2.5c.

Let us follow the state $\mathbf{x}(t)$ of the population, for $a < 1$. If the state is very close to a vertex, for instance \mathbf{e}_1 , it is close to a rest point and hence almost at rest. For a long time, the state does not seem to change. Then, it picks up speed and moves towards the vicinity of the vertex \mathbf{e}_3 , where it slows down and remains for a much longer time, etc. This looks like a recurrent form of “punctuated equilibrium”: long periods of quasi-rest followed by abrupt upheavals.

2.15 STOCHASTIC PROCESSES AND FIXATION PROBABILITIES

So far, we have considered the limiting case of infinitely large populations. If we assume a population of finite size, we can no longer rely on deterministic models. In finite populations, random fluctuations, due for instance to sampling effects, have to be taken into account. Instead of ordinary differential equations, we must use stochastic processes.

Let us assume, in the simplest case, that a population of finite size M consists of two types of players only, \mathbf{e}_1 and \mathbf{e}_2 . From time to time, one of the players updates strategy, by imitating a model chosen from the population. The state of the population is given by the number i of individuals of type \mathbf{e}_1 (while the number of players of type \mathbf{e}_2 is $M - i$). Let p_{ij} be the probability that the transition leads from i to j . The matrix P is tri-diagonal, i.e., $p_{ij} = 0$ if $|j - i| > 1$. The states 0 and M are absorbing: if all individuals are of the same type, imitation can not introduce the other type. We write $p_{i,i+1} = b_i$ and $p_{i,i-1} = d_i$ (because these transition probabilities correspond, in another interpretation, to *birth* and *death* rates).

We denote by p_i the probability that a population in state i will eventually reach state M , i.e., consist entirely of type \mathbf{e}_1 . This state M is *absorbing*, since once reached it will not be left. The probability p_i that, starting in state i , such a *fixation* of the type \mathbf{e}_1 occurs, must satisfy

$$p_i = d_i p_{i-1} + (1 - b_i - d_i) p_i + b_i p_{i+1} \quad (2.62)$$

for $i = 1, \dots, M - 1$. Indeed, in the first updating event, the number of \mathbf{e}_1 's will either increase or decrease by 1, or remain unchanged (when a player imitates someone of his own kind); and after this first step, fixation must occur. Moreover, we have $p_0 = 0$ and $p_M = 1$. Setting $y_i := p_i - p_{i-1}$, equation (2.62) can be written as

$$y_{i+1} = \frac{d_i}{b_i} y_i. \quad (2.63)$$

Since $p_1 = y_1$ and $\sum_{i=0}^k y_i = p_k - p_0 = p_k$, we obtain

$$1 = p_M = \sum_{i=1}^M y_i = p_1 \left(1 + \frac{d_1}{b_1} + \dots + \frac{d_1 d_2}{b_1 b_2} \dots \frac{d_{M-1}}{b_{M-1}} \right), \quad (2.64)$$

so that

$$p_i = \frac{1 + \sum_{j=1}^{i-1} \prod_{k=1}^j d_k/b_k}{1 + \sum_{j=1}^{M-1} \prod_{k=1}^j d_k/b_k} \quad (2.65)$$

for $i = 1, \dots, M$. In particular, we denote by $\rho_{1,2}$ the fixation probability p_1 , i.e., the probability that a single individual of type e_1 in a population consisting otherwise of type e_2 will eventually be imitated by everyone. It is given by

$$\rho_{1,2} = \frac{1}{1 + \sum_{j=1}^{M-1} \prod_{k=1}^j d_k/b_k}. \quad (2.66)$$

So far, we have not specified the imitation mechanism. In this chapter, we shall consider only the so-called Moran process, developed in the context of population genetics. For this reason, we shall adopt the corresponding terminology, and assume that each individual has a certain "fitness," which in our context means some measure of success, such that individuals with a higher fitness are more likely to be imitated (see section 2.7). The Moran process consists in drawing one individual at random (each has the same probability $1/M$ of being chosen) and endowing it with the type of a "model player" who is selected from the population with a probability proportional to that model's success.

Thus let us assume, as a first example, that individuals of type e_1 have fitness r , while those of type e_2 have a fitness normalized to be equal to 1. We then obtain for the death rate

$$d_i = \left(\frac{i}{M}\right) \left(\frac{M-i}{ri + M-i}\right), \quad (2.67)$$

where the first fraction is the probability that the updating individual is of type e_1 , and the second that the selected model is of type e_2 . Similarly, for the birth rate,

$$b_i = \left(\frac{M-i}{M}\right) \left(\frac{ri}{ri + M-i}\right). \quad (2.68)$$

Hence $d_i/b_i = 1/r$ and

$$\rho_{1,2} = \frac{1 - r^{-1}}{1 - r^{-M}}. \quad (2.69)$$

If $r \rightarrow 1$, we obtain as limiting value $\rho_{1,2} = 1/M$, which is reassuring. This is the fixation probability of a *neutral* type, i.e., the probability that a single individual of type e_1 , doing exactly as well as the resident e_2 individuals, will eventually be copied by the entire population.

2.16 GAMES IN FINITE POPULATIONS

Now suppose that in a population of size M , individuals are engaged in pairwise games, and strategies are determined by type e_1 or e_2 . If the payoff matrix is

$$\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix}, \quad (2.70)$$

then the expected payoff depends on the state i of the population. For a player of type e_1 , it is given by

$$F_i = \alpha \frac{i-1}{M-1} + \beta \frac{M-i}{M-1}, \quad (2.71)$$

and for a player of type e_2 by

$$G_i = \gamma \frac{i}{M-1} + \delta \frac{M-i-1}{M-1}. \quad (2.72)$$

(Players do not play against themselves.) As in section 2.7, we assume that the fitness, i.e., the likelihood to be imitated, is given as a convex combination of the payoff and a "baseline fitness," the same for all, which we normalize to 1. Hence the fitness of an e_1 individual, if the population is in state i , is given by

$$f_i = (1-s) + sF_i \quad (2.73)$$

and that of an e_2 individual is

$$g_i = (1-s) + sG_i. \quad (2.74)$$

Here the parameter $s \in [0, 1]$ measures the "strength of selection," i.e., the importance of the game for overall success. If $s = 0$ the game is irrelevant. In the limiting case of an infinitely large population, the Moran process leads to the switching rate given by equation (2.29) and hence to the replicator dynamics.

The birth and death rates are

$$b_i = \left(\frac{M-i}{M}\right) \left(\frac{if_i}{if_i + (M-i)g_i}\right) \quad (2.75)$$

and

$$d_i = \left(\frac{i}{M}\right) \left(\frac{(M-i)g_i}{if_i + (M-i)g_i}\right), \quad (2.76)$$

so that

$$\frac{d_i}{b_i} = \frac{g_i}{f_i} = \frac{1-s(1-G_i)}{1-s(1-F_i)}. \quad (2.77)$$

The fixation probability (2.66) is therefore given by

$$\rho_{1,2} = \left[1 + \sum_{j=1}^{M-1} \prod_{i=1}^j \frac{1-s+sG_i}{1-s+sF_i} \right]^{-1}. \quad (2.78)$$

2.17 LIMITING CASES

For small s , expression (2.77) can be approximated, up to first order, by

$$\frac{d_i}{b_i} = 1 - s(F_i - G_i). \quad (2.79)$$

Now by equations (2.71) and (2.72),

$$H_i := F_i - G_i = \frac{1}{M-1}[\bar{e} + \bar{f}i] \quad (2.80)$$

with $\bar{e} = -\alpha + \beta M - \delta M + \delta$ and $\bar{f} = \alpha - \beta - \gamma + \delta$. Hence up to first order in s , the fixation probability of type e_1 is, according to equation (2.78), given by

$$\rho_{1,2} = \left[1 + \sum_{k=1}^{M-1} \prod_{i=1}^k (1 - sH_i) \right]^{-1}. \quad (2.81)$$

It is easy to see that

$$\sum_{k=1}^{M-1} \prod_{i=1}^k (1 - sH_i) = M - 1 - s \sum_{i=1}^{M-1} (M-i)H_i, \quad (2.82)$$

and that

$$\sum_{i=1}^{M-1} (M-i)(\bar{e} + \bar{f}i) = M(M-1)\bar{e} + (M\bar{f} - \bar{e}) \sum_{i=1}^{M-1} i - \bar{f} \sum_{i=1}^{M-1} i^2. \quad (2.83)$$

The first sum on the right hand side is $M(M-1)/2$ and the second sum is $M(M-1)(2M-1)/6$. This yields altogether

$$M(M-1)(M\bar{f} + \bar{f} + 3\bar{e})/6 = M(M-1)(eM - f)/6, \quad (2.84)$$

with $e = \alpha + 2\beta - \gamma - 2\delta$ and $f = 2\alpha + \beta + \gamma - 4\delta$. Up to first order in s , equation (2.81) yields

$$\rho_{1,2} = \left[1 - \frac{s}{6}(eM - f) \right]^{-1} / M. \quad (2.85)$$

We say that strategy e_1 is *advantageous* if its fixation probability is higher than that of a neutral mutant, i.e., if $\rho_{1,2} > 1/M$. This condition reads $eM > f$, i.e.,

$$\alpha(M-2) + \beta(2M-1) > \gamma(M+1) + \delta(2M-4). \quad (2.86)$$

For the limit $M \rightarrow \infty$ we obtain

$$\alpha + 2\beta > \gamma + 2\delta, \quad (2.87)$$

or, with the normalization from matrix (2.37), $b < 2a$. This inequality always holds if e_1 dominates e_2 , i.e., if $b \leq 0$ and $a \geq 0$ (one inequality being strict). The dominant strategy is always advantageous. In the case of stable coexistence, i.e., $a > 0$ and $b > 0$, it means that $\hat{x} > 1/3$, where $\hat{x} = \frac{a}{a+b}$ is the Nash equilibrium. Thus if $1/3 < \hat{x} < 2/3$, both strategies are advantageous. Finally, in the case of a bi-stable game, i.e., if $a < 0$ and $b < 0$, inequality (2.86) means that $\hat{x} < 1/3$, where \hat{x} is the unstable Nash equilibrium in $]0, 1[$. This means that for the replicator equation, the basin of attraction of e_1 is more than twice as large as that of e_2 . In particular, if e_1 is advantageous, it is risk-dominant. If \hat{x} lies between $1/3$ and $2/3$, none of the two strategies is advantageous.

For the examples of a Repeated Prisoner's Dilemma game (section 1.7) or a Snowdrift game (section 1.4), we see from inequality (2.86) that cooperation is advantageous for $M > 4$ resp. $M > 20$.

The fixation probability $\rho_{2,1}$ of e_2 is obtained similarly to that of e_1 (by replacing e with $-2\alpha - \beta + 2\gamma + \delta$ and f with $-4\alpha + \beta + \gamma + 2\delta$). The condition $\rho_{1,2} > \rho_{2,1}$ means

$$(M-2)(\alpha - \delta) > M(\gamma - \beta). \quad (2.88)$$

In the limit of large M , this reduces to the condition $\alpha - \delta > \gamma - \beta$. This is just the condition $a > b$ that e_1 is risk-dominant.

For any value of $\pi \in [0, 1]$, the vector $(\pi, 0, \dots, 0, 1 - \pi) \in S_M$ is a stationary distribution of the imitation process. Let us assume that with some probability $\mu > 0$, players can change their strategy without imitating another player, just by random trial. In that case, the resulting Markov chain is recurrent. It describes the interplay between innovation and imitation. Let us assume that μ is so small that we can separate the time scales of the two processes. This means that most of the time, the population is in the homogeneous state 0 or M . Occasionally, a single individual tries the other strategy. Then, the imitation process starts anew, leading either to the extinction of the new type or to its fixation. In this "adiabatic" case, the resulting process can be approximated by a Markov chain with two states, 1 and 2, (which correspond to homogeneous populations consisting of type e_1 or e_2). This Markov chain is given by the matrix

$$\begin{pmatrix} 1 - \rho_{2,1} & \rho_{2,1} \\ \rho_{1,2} & 1 - \rho_{1,2} \end{pmatrix} \quad (2.89)$$

whose unique stationary distribution, the left eigenvector

$$\left(\frac{\rho_{1,2}}{\rho_{1,2} + \rho_{2,1}}, \frac{\rho_{2,1}}{\rho_{1,2} + \rho_{2,1}} \right) \quad (2.90)$$

describes the prevalence of the two types, for large time spans. In particular, for the bi-stable case and large M , strategy e_1 is risk-dominant if and only if the stochastic process spends more time in the corresponding homogeneous state.

The same "adiabatic" argument holds also for n types e_i . If the "innovation rate" μ is sufficiently small, the population will always consist of one or at most two types only. If in a homogeneous population, a single individual switches to a different type, then the imitation process will have caused the fixation or the elimination of that type before the next innovation occurs. If we assume that these innovations are random explorations, i.e., that every non-resident type has the same chance $1/(n-1)$ to occur, we obtain an $n \times n$ Markov chain P with transition probabilities p_{ij} given by

$$p_{ij} = \rho_{j,i}/(n-1) \quad (2.91)$$

for $i \neq j$. Here, $\rho_{j,i}$ is the fixation probability of j in i , i.e., the probability that a single individual of type j in a population consisting otherwise of type i will eventually be copied by the entire population.

2.18 REFERENCES

The first, and already classic book on evolutionary game theory is by Maynard Smith (1982). Among the textbooks which have appeared since, Weibull (1995), Hofbauer and Sigmund (1998), and Cressman (2003) are closest in spirit to the approach presented here, but see also Colman (1995), Fudenberg and Levine (1998), Gintis (2000), and Vincent and Brown (2005). Hofbauer and Sigmund (2003) and Sandholm (2009) consider more general types of game dynamics. The replicator equation was introduced by Taylor and Jonker (1978), see also Hofbauer, Schuster, and Sigmund (1979) and Zeeman (1980); the name was suggested in Schuster and Sigmund (1983). Derivations via social learning were proposed by Helbing (1992) and Schlag (1997). Bomze (1983) gives a classification of replicator dynamics with three strategies. Hofbauer (2000) presents an account of the relations between the concept of Nash equilibrium pairs, existence proofs, and diverse game dynamics. Our treatment of finite population games closely follows Nowak (2006a), see also Nowak et al. (2004), and Taylor et al. (2004). A different, less "stochastic" but related approach is used in Kandori, Mailath, and Rob (1993) and in Peyton Young and Foster (1995).

Chapter Three

Direct Reciprocity: The Role of Repetition

3.1 HELP

As Darwin wrote, "The small strength and speed of man, his want of natural weapons, etc., are more than counterbalanced . . . by his social qualities which lead him *to give and receive* aid from his fellow-men" (italics added). In its simplest form, to help means to confer a benefit b to another individual, at a cost c to oneself. This can be viewed as an atom of social interaction.

In the *Donation game*, two players have to decide simultaneously (more precisely, in ignorance of the co-player's decision) whether to give help to their co-player or not. The two strategies e_1 and e_2 will be denoted by C (for *cooperate*) and D (for *defect*), respectively. This yields the following payoff matrix:

$$\begin{pmatrix} b-c & -c \\ b & 0 \end{pmatrix}. \quad (3.1)$$

If not otherwise stated, we will assume $b > c > 0$. The second strategy D dominates the first. This is an example of a Prisoner's Dilemma game, as described in section 1.3, i.e., a symmetric 2×2 game whose payoff matrix

$$\begin{pmatrix} R & S \\ T & P \end{pmatrix} \quad (3.2)$$

satisfies

$$T > R > P > S. \quad (3.3)$$

The Prisoner's Dilemma game encapsulates the tug-of-war between the common interest (R is larger than P) and the selfish interest (D dominates C). Selfishness ought to win in this conflict. Indeed, the game has a unique Nash equilibrium, namely defection; and imitation of successful individuals leads inexorably to the demise of cooperation, see section 2.10.

It can be interesting to compare this Donation game with the Snowdrift game (see section 1.4). Both players can receive a benefit b each, if they come up with a fee $c < b$. They have to decide simultaneously whether to pay the fee or not, knowing that if both decide to pay, they will share the cost. The payoff matrix is

$$\begin{pmatrix} b - \frac{c}{2} & b - c \\ b & 0 \end{pmatrix}. \quad (3.4)$$

Obviously, it is best to do the opposite of what the other player does. If your co-player is willing to pay the fee, you yourself can safely skip it. But if your co-player