Proceedings of Symposia in Applied Mathematics
Volume **69**, 2011

Introduction to Evolutionary Game Theory

Karl Sigmund

Abstract. This chapter begins with some basic terminology, introducing elementary game theoretic notions such as payoff, strategy, best reply, Nash equilibrium pairs etc. Players who use strategies which are in Nash equilibrium have no incentive to deviate unilaterally. Next, a population viewpoint is introduced. Players meet randomly, interact according to their strategies, and obtain a payoff. This payoff determines how the frequencies in the strategies will evolve. Successful strategies spread, either (in the biological context) through inheritance or (in the cultural context) through social learning. The simplest description of such an evolution is based on the replicator equation. The basic properties of replicator dynamics are analyzed, and some low-dimensional examples such as the Rock-Scissors-Paper game are discussed. The relation between Nash equilibria and rest points of the replicator equation are investigated, which leads to a short proof of the existence of Nash equilibria. We then study mixed strategies and evolutionarily stable strategies. This introductory chapter continues with a brief discussion of other game dynamics, such as the best reply dynamics, and ends with the simplest extension of replicator dynamics to asymmetric games.

1. Predictions and Decisions

Predictions can be difficult to make, especially, as Niels Bohr quipped, if they concern the future. Reliable forecasts about the weather or about some social development may seem to offer comparable challenges, at first sight. But there is a fundamental difference: a weather forecast does not influence the weather, whereas a forecast on economy can influence the economic outcome. Humans will react if they learn about the predictions, and they can anticipate that others will react, too.

When the economist Oskar Morgenstern, in the early 'thirties, became aware of the problem, he felt that he had uncovered an 'impossibility theorem' of a similarly fundamental nature as the incompleteness theorem of his friend, the logician Kurt Gödel. Morgenstern was all the more concerned about it as he was director of the Vienna-based Institut für Konjunkturforschung, the Institute for Business Cycles Research, whose main task was actually to deliver predictions on the Austrian economy. Oscar Morgenstern expained his predicament in many lectures and publications, using as his favorite example the pursuit of Sherlock Holmes by the

---

2000 *Mathematics Subject Classification.* Primary 91A22.

The author wants to thank Ross Cressman for his helpful comments.

infamous Professor Moriarty [**24**]. These two equally formidable adversaries would never arrive at a conclusive solution in mutually outguessing each other.

We can describe the fundamental nature of the problem by using some of the mathematical notation which later was introduced through game theory. Let us suppose that player I has to choose between n options, or *strategies*, which we denote by $\mathbf{e}_1$,..., $\mathbf{e}_n$, and player II between m strategies $\mathbf{f}_1$,..., $\mathbf{f}_m$. If I chooses $\mathbf{e}_i$ and II chooses $\mathbf{f}_j$, 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' preference ranking.

The two players could engage in the game 'Odd or Even?' and decide that the loser pays one dollar to the winner. At a given signal, each player holds up one or two fingers. If the resulting sum is odd, player I wins. If the sum is even, player II wins. Each player then has to opt for *even* and *odd*, which correspond to $\mathbf{e}_1$ and $\mathbf{e}_2$ for player I and $\mathbf{f}_1$ and $\mathbf{f}_2$ for player II, and the payoff matrix is

$$\begin{pmatrix} (-1,1) & (1,-1) \\ (1,-1) & (-1,1) \end{pmatrix} \tag{1.1}$$

If the two players graduate to the slightly more sophisticated Rock-Scissors-Paper game, they would each have to opt between three strategies, numbered in that order, and the payoff matrix would be

$$\begin{pmatrix} (0,0) & (1,-1) & (-1,1) \\ (-1,1) & (0,0) & (1,-1) \\ (1,-1) & (-1,1) & (0,0) \end{pmatrix}. \tag{1.2}$$

If both Rock-Scissors-Paper players opt for the same move, the game is a tie and both obtain payoff zero. If the outcome is $(0,0)$ or $(-1,1)$, then player I (who chooses the row of the payoff matrix) would have done better to choose another strategy; if the outcome is $(1,-1)$ or $(0,0)$, then it is player II, the column player, who would have done better to switch. If a prediction is made public, then at least one of the players would have an incentive to deviate. The other player would anticipate this, and deviate accordingly, and both would be launched into a vicious circle of mutual outguessing.

A few years, however, after Morgenstern had started to broadcast his impossibility result, the topologist Cech pointed out to him that John von Neumann had found, in an earlier paper on parlor games, a way to avoid Morgenstern's dead end [**42**]. It consists in randomizing, i.e. letting chance decide. Clearly, if players opt with equal probability for each of their alternatives, none has an incentive to deviate. Admittedly, this would lead to the expected payoff 0, somewhat of an anti-climax. But John von Neumann's minimax theorem holds for a much larger class of games. Most importantly, it led, in the 'forties, to a collaboration of John von Neumann with Oscar Morgenstern which gave birth to game theory [**43**]. A few years later, John Nash introduced an equilibrium notion valid in an even more general context, which became the cornerstone of game theory [**33**].

2. Mixed strategies and best replies

Suppose that player I opts to play strategy $\mathbf{e}_i$ with probability x_i. This *mixed* strategy is thus given by a stochastic vector $\mathbf{x} = (x_1, ..., x_n)$ (with $x_i \geq 0$ and

$x_1 + ... + x_n = 1$). We denote the set of all such mixed strategies by Δ_n: this is a simplex in R^n, spanned by the unit vectors $\mathbf{e}_i$ of the standard base, which are said to be the *pure* strategies, and correspond to the original set of alternatives.

Similarly, a mixed strategy for player II is an element $\mathbf{y}$ of the unit simplex Δ_m spanned by the unit vectors $\mathbf{f}_j$. If player I uses the pure strategy $\mathbf{e}_i$ and player II uses strategy $\mathbf{y}$, then the payoff for player I (or more precisely, its expected value) is

$$(A\mathbf{y})_i = \sum_{j=1}^{m} a_{ij} y_j. \tag{2.1}$$

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

$$\mathbf{x} \cdot A\mathbf{y} = \sum_i x_i (A\mathbf{y})_i = \sum_{i,j} a_{ij} x_i y_j, \tag{2.2}$$

and the payoff for player II, similarly, is

$$\mathbf{x} \cdot B\mathbf{y} = \sum_{i,j} b_{ij} x_i y_j. \tag{2.3}$$

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

$$BR(\mathbf{y}) = \arg\max_{\mathbf{x}} \mathbf{x} \cdot A\mathbf{y}, \tag{2.4}$$

i.e. the set of all $\mathbf{x} \in \Delta_n$ such that $\mathbf{z} \cdot A\mathbf{y} \leq \mathbf{x} \cdot A\mathbf{y}$ holds for all $\mathbf{z} \in \Delta_n$. Player I has no incentive to deviate from $\mathbf{x}$ and chose another strategy $\mathbf{z}$ instead.

Since the function $\mathbf{z} \mapsto \mathbf{z} \cdot A\mathbf{y}$ is continuous and Δ_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 $\mathbf{e}_i$ for which $x_i > 0$. Indeed, for all i,

$$(A\mathbf{y})_i = \mathbf{e}_i \cdot A\mathbf{y} \leq \mathbf{x} \cdot A\mathbf{y}. \tag{2.5}$$

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 $i = 1, ..., n$ then leads to a contradiction. It follows that the set $BR(\mathbf{y})$ is a face of the simplex Δ_n. It is spanned by the pure strategies which are best replies to $\mathbf{y}$.

If player I has found a best reply to the strategy $\mathbf{y}$ of player II, then player I has no incentive 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 either 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

$$\mathbf{z} \cdot A\mathbf{y} \leq \mathbf{x} \cdot A\mathbf{y} \tag{2.6}$$

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

$$\mathbf{x} \cdot B\mathbf{w} \leq \mathbf{x} \cdot B\mathbf{y} \tag{2.7}$$

holds for all $\mathbf{w} \in \Delta_m$. 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 there always exist Nash equilibrium pairs, for any game (A, B). 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 (see, e.g., [**30**]). But it would not hold if we had not allowed for mixed strategies, as is shown by the Rock-Scissors-Paper game. In that case, the mixed strategy which consists in choosing, with equal probability 1/3, among the three alternative moves, clearly leads to an equilibrium pair. No player has a reason to deviate. On the other hand, if player I uses any other strategy (x_1, x_2, x_3) against the $(1/3, 1/3, 1/3)$ of player II, player I would still have an expected payoff of 0. However, the *other* player II would then have an incentive to deviate, presenting I with an incentive to deviate in turn, and so on.

In this example, $(\mathbf{x}, \mathbf{y})$ with $\mathbf{x} = \mathbf{y} = (1/3, 1/3, 1/3)$ is the unique Nash equilibrium pair. We have seen that as long as player II chooses the equilibrium strategy $\mathbf{y}$, player I has no reason to deviate from the equilibrium strategy $\mathbf{x}$, but that on the other hand, player I has no reason not to deviate, either. This would be different if $(\mathbf{x}, \mathbf{y})$ were a *strict* Nash equilibrium pair, i.e. if

$$\mathbf{z} \cdot A\mathbf{y} < \mathbf{x} \cdot A\mathbf{y} \tag{2.8}$$

holds for all $\mathbf{z} \neq \mathbf{x}$, and

$$\mathbf{x} \cdot B\mathbf{w} < \mathbf{x} \cdot B\mathbf{y} \tag{2.9}$$

holds for all $\mathbf{w} \neq \mathbf{y}$. In this case, i.e. when both best-reply sets are singletons, and hence correspond to pure strategies, each player will be penalized for unilaterally deviating from the equilibrium.

Whereas every game admits a Nash equilibrium pair, some need not admit a strict Nash equilibrium pair, as our previous examples show.

3. Excurse to zero sum

Historically, game theory focused first on zero-sum games, for which $a_{ij} = -b_{ij}$ for all i, j, i.e., $A = -B$ (the gain of player I is the loss of player II). This condition clearly holds for a large set of parlor games. But it certainly restricts the range of applications. For most types of social and economic interactions, the assumption that the interests of the two players are always diametrically opposite does not hold. Even in military confrontations, there often exist outcomes both parties want to avoid. Most interactions are of mixed motive type, and contain elements of cooperation as well as competition. Interestingly, John von Neumann did not greatly appreciate the solution concept proposed by the undergraduate student John Nash. A short interview ended when John von Neumann remarked, somewhat dismissively, 'Oh, just a fixed point theorem' [**32**]. We shall see that the existence proof for Nash equilibrium pairs does indeed reduce to a fixed point theorem, and a rather simple one at that. Nevertheless, it yields a very powerful result, as can be seen by applying it to the special case of zero sum games, where it leads to a three-liner proof of the celebrated maximin theorem, a proof which is considerably simpler than John von Neumann's original, brute-force demonstration.

It is easy to see that $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is a Nash equilibrium pair of a zero-sum game iff

$$\mathbf{x} \cdot A\bar{\mathbf{y}} \leq \bar{\mathbf{x}} \cdot A\bar{\mathbf{y}} \leq \bar{\mathbf{x}} \cdot A\mathbf{y} \tag{3.1}$$

for all $\mathbf{x} \in \Delta_n, \mathbf{y} \in \Delta_m$. Suppose that player II correctly guesses that player I plays $\mathbf{x}$. Then player II will use a strategy which is a best reply, i.e., minimizes player I's payoff, which will reduce to $g(\mathbf{x}) := \min_{\mathbf{y}} \mathbf{x} \cdot A\mathbf{y}$. A player I who expects to be anticipated, then, ought to maximize $g(\mathbf{x})$. Any strategy $\hat{\mathbf{x}}$ yielding this maximum is said to be a maximin strategy for player I. Such a maximin strategy is defined by $\hat{\mathbf{x}} := \arg\max_{\mathbf{x}} g(\mathbf{x})$, and it guarantees player I a security level

$$w_u := \max_{\mathbf{x}} \min_{\mathbf{y}} \mathbf{x} \cdot A\mathbf{y}. \tag{3.2}$$

Similarly, we can expect player II to maximize the own security level, i.e., since $A = -B$, to play a minimax strategy $\hat{\mathbf{y}}$ such that player I has a payoff bounded from above by

$$w_o := \min_{\mathbf{y}} \max_{\mathbf{x}} \mathbf{x} \cdot A\mathbf{y}. \tag{3.3}$$

The pair $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is said to be a *maximin pair*. It satisfies

$$\min_{\mathbf{y}} \hat{\mathbf{x}} \cdot A\mathbf{y} = w_u, \qquad \max_{\mathbf{x}} \mathbf{x} \cdot A\hat{\mathbf{y}} = w_o, \tag{3.4}$$

and leads to a payoff which clearly satisfies

$$w_u \le \hat{\mathbf{x}} \cdot A\hat{\mathbf{y}} \le w_o. \tag{3.5}$$

If $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is a Nash equilibrium pair for a zero sum game, then it is a maximin pair. Indeed, by (3.1),

$$\max_{\mathbf{x}} \mathbf{x} \cdot A\bar{\mathbf{y}} \le \bar{\mathbf{x}} \cdot A\bar{\mathbf{y}} \le \min_{\mathbf{y}} \bar{\mathbf{x}} \cdot A\mathbf{y}. \tag{3.6}$$

Now by (3.3), w_o is less than the left hand side of the previous inequality and by (3.2) w_u larger than the right hand side. Since $w_u \le w_o$ by (3.5), we must actually have equality everywhere. But $w_u = \min_{\mathbf{y}} \bar{\mathbf{x}} \cdot A\mathbf{y}$ means that $\bar{\mathbf{x}}$ is a maximin solution, and $\max_{\mathbf{x}} \mathbf{x} \cdot A\bar{\mathbf{y}} = w_o$ that $\bar{\mathbf{y}}$ is a minimax solution.

For zero-sum games, the existence of a Nash equilibrium pair thus implies the existence of a maximin pair. The previous argument implies $w_u = w_o$, i.e.,

$$\min_{\mathbf{y}} \max_{\mathbf{x}} \mathbf{x} \cdot A\mathbf{y} = \max_{\mathbf{x}} \min_{\mathbf{y}} \mathbf{x} \cdot A\mathbf{y}. \tag{3.7}$$

Conversely, it is easy to see that if $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a maximin pair of a zero sum game, then it is a Nash equilibrium pair.

4. Concerns about the Nash solution

Let us note that if $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ and $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ are two Nash equilibrium pairs for a zero sum game, then so are $(\hat{\mathbf{x}}, \bar{\mathbf{y}})$ and $(\bar{\mathbf{x}}, \hat{\mathbf{y}})$. Indeed,

$$\hat{\mathbf{x}} \cdot A\hat{\mathbf{y}} \le \hat{\mathbf{x}} \cdot A\bar{\mathbf{y}} \le \bar{\mathbf{x}} \cdot A\bar{\mathbf{y}} \le \bar{\mathbf{x}} \cdot A\hat{\mathbf{y}} \le \hat{\mathbf{x}} \cdot A\hat{\mathbf{y}}, \tag{4.1}$$

hence equality holds everywhere, and therefore for all $\mathbf{x}$ and $\mathbf{y}$:

$$\mathbf{x} \cdot A\bar{\mathbf{y}} \le \hat{\mathbf{x}} \cdot A\bar{\mathbf{y}} \le \hat{\mathbf{x}} \cdot A\mathbf{y} \tag{4.2}$$

so that $(\hat{\mathbf{x}}, \bar{\mathbf{y}})$ is a Nash equilibrium pair.

The same need not hold for general (non zero-sum) games. Consider for instance

$$\begin{pmatrix} (1,1) & (-1,-1) \\ (-1,-1) & (1,1) \end{pmatrix} \tag{4.3}$$

It is easy to see that $(\mathbf{e}_1, \mathbf{f}_1)$ and $(\mathbf{e}_2, \mathbf{f}_2)$ are two Nash equilibrium pairs. But $(\mathbf{e}_1, \mathbf{f}_2)$ or $(\mathbf{e}_2, \mathbf{f}_1)$ are not. How should the two players coordinate their choice? The problem becomes even more acute for a coordination game given by

$$\begin{pmatrix} (2,2) & (-100,0) \\ (0,-100) & (1,1) \end{pmatrix}. \tag{4.4}$$

Again, $(\mathbf{e}_1, \mathbf{f}_1)$ and $(\mathbf{e}_2, \mathbf{f}_2)$ are two Nash equilibrium pairs. The former has the advantage of yielding a higher payoff to both players: it is said to be *Pareto-optimal*. But the second is less risky, and therefore said to be *risk-dominant*. Indeed, it can be very costly to go for the Pareto-optimum if the other player fails to do so. It may actually be best to decide against using the Pareto-optimum right away. In any case, if the game is not zero-sum, Nash equilibrium pairs may not offer much help for decision makers.

Moreover, even if there exists a unique Nash equilibrium pair, it can lead to frustration, as in the following example:

$$\begin{pmatrix} (10,10) & (-5,15) \\ (15,-5) & (0,0) \end{pmatrix}. \tag{4.5}$$

In this case, $\mathbf{e_2}$ is the best reply to every (pure or mixed) strategy of player II, and similarly $\mathbf{f}_2$ is always the best reply for player II. Hence $(\mathbf{e}_2, \mathbf{f}_2)$ is the unique Nash equilibrium pair, and it is strict. This game is an example of a *Prisoner's Dilemma game*. The payoff matrix may occur, for instance, if two players are asked to choose, independently and anonymously, whether or not to provide a gift of 15 dollars to the co-player, at a cost of 5 dollars to themselves. It the two players cooperate by both opting for their first strategy, they will end up with 10 dollars each. But each has an incentive to deviate. It is only when both opt for their second solution and defect, that they cannot do better by choosing to deviate. But then, they end up with zero payoff. Let us remark that this dilemma cannot be solved by appealing to non-monetary motivations. It holds whenever the payoff values reflect each players' preference ordering, which may well include a concern for the other.

5. Population Games

So far, we have considered games between two specific players trying to guess each other's strategy and find a best reply. This belongs to the realm of classical game theory, and leads to interesting mathematical and economic developments. Starting with the 'sixties and 'seventies, both theory and applications were considerably stimulated by problems in evolutionary biology, such as sex-ratio theory or the investigation of fighting behavior [**12, 27**]. It required a radical shift in perspective and the introduction of thinking in terms of populations [**29**]. It provided a welcome tool for the analysis of frequency dependent selection and, later, of learning processes.

Let us therefore 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). The outcome of each encounter yields payoff values which are no longer measured on utility scales reflecting the individual preferences of the players, but in the one currency that counts in Darwinian evolution, namely fitness, i.e., average reproductive success. If we assume that strategies can be passed on to the offspring, whether through inheritance or through learning, then we can assume that more successful strategies spread.

In order to analyze this set-up, it is convenient to assume, in a first approach, that all individuals in the population are indistinguishable, except in their way of interacting, i.e. that the players differ only by their strategy. This applies well to games where both players are on an equal footing. Admittedly, there are many examples of social interactions which display an inherent asymmetry between the two players: for instance, between buyers and sellers, or between parents and offspring. We will turn to such interactions later.

Thus 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 $\mathbf{f}_j = \mathbf{e}_j$ for all j; and if a player plays strategy $\mathbf{e}_i$ against someone using strategy $\mathbf{e}_j$ (which is the former $\mathbf{f}_j$), then that player receives the same payoff, whether labeled I or II. Hence $a_{ij} = b_{ji}$, the payoff for a $\mathbf{e}_i$-player against a $\mathbf{e}_j$-players does not depend on who is labelled I and who is II, or in other words $B = A^T$. Thus a symmetric game is specified by the pair (A, A^T), and therefore is defined by a single, square payoff matrix A. All examples encountered so far are symmetric, with the exception of 'Even or Odd'. A zero-sum game which is symmetric must satisfy $A^T = -A$ and hence corresponds to a skew-symmetric payoff matrix.

It is easy to see that the symmetric game given by

$$\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}, \tag{5.1}$$

where success depends on doing the opposite of the co-player, admits $(\mathbf{e}_1, \mathbf{e}_2)$ and $(\mathbf{e}_2, \mathbf{e}_1)$ as asymmetric Nash equilibrium pairs. These are plainly irrelevant as solutions of the game, since it is impossible to distinguish players I and II. Of interest are only symmetric Nash equilibrium pairs, i.e. pairs of strategies $(\mathbf{x}, \mathbf{y})$ with $\mathbf{x} = \mathbf{y}$. A *symmetric Nash equilibrium*, thus, is specified by *one* strategy $\mathbf{x}$ having the property that it is a best reply to itself (i.e. $\mathbf{x} \in BR(\mathbf{x})$). In other words, we must have

$$\mathbf{z} \cdot A\mathbf{x} \leq \mathbf{x} \cdot A\mathbf{x} \tag{5.2}$$

for all $\mathbf{z} \in \Delta_n$. A symmetric *strict* Nash equilibrium is accordingly given by the condition

$$\mathbf{z} \cdot A\mathbf{x} < \mathbf{x} \cdot A\mathbf{x} \tag{5.3}$$

for all $\mathbf{z} \neq \mathbf{x}$.

We shall soon prove that every symmetric game admits a symmetric Nash equilibrium. But first, we consider a biological toy model which played an essential role in the emergence of evolutionary game theory [**27**]. It is due to two eminent theoretical biologists, John Maynard Smith and George Price, who tried to explain the evolution of ritual fighting in animal contests. It had often been observed that in conflicts within a species, animals did not escalate the fight, but kept to certain stereotyped behavior, such as posturing, glaring, roaring or engaging in a pushing match. Signals of surrender (such as offering the unprotected throat) stopped the fight as reliably as a towel thrown into the boxing ring. Interestingly, thus, animal fights seem to be restrained by certain rules, without even needing a referee. Such restraint is obviously all for the good of the species, but Darwinian thinking does not accept this as an argument for its emergence. An animal ignoring these 'gloved fist'-type of rules, and killing its rivals, should be able to spread its genes, and the readiness to escalate a conflict should grow, even if this implies, in the long run, suicide for the species.

Maynard Smith and Price imagined, in their thought experiment, a population consisting of two phenotypes (or strategies). Strategy $\mathbf{e}_1$ is a behavioral program to escalate the conflict until death or injury settles the outcome. Strategy $\mathbf{e}_2$ is a behavioral program to flee as soon as the opponent starts getting rough. The former strategy is called 'Hawk', the latter 'Dove'. Winning the conflict yields an expected payoff G, and losing an escalated fight costs $C > G$ (with G and C measured on the scale of Darwinian fitness). If we assume that whenever two 'Hawks' meet, or two 'Doves', both are equally likely to win the contest, then their expected payoff is $G/2 - C/2$ resp. $G/2$. The payoff matrix thus is

$$\left(\begin{array}{cc} \frac{G-C}{2} & G \\ 0 & \frac{G}{2} \end{array} \right). \tag{5.4}$$

Clearly, neither 'Hawk' nor 'Dove' is a Nash equilibrium. In terms of evolutionary biology, a homogeneous 'Dove' population could be invaded by a minority of 'Hawks' who win all their contests hands down; but similarly, a homogeneous 'Hawk' population could be invaded by a minority of 'Doves', whose payoff 0 is larger than the negative payoff $(G - C)/2$ experienced by the 'Hawks' tearing at each other. It is better to experience no change in reproductive success, rather than a reduction. In this sense neither 'Hawk' nor 'Dove' is an evolutionarily stable strategy. On the basis of this extremely simplified model, we must expect evolution to lead to a mixed population.

6. Population dynamics

Let us consider a symmetric game with payoff matrix A and assume that in a large, well-mixed population, a fraction x_i uses strategy $\mathbf{e}_i$, for $i = 1, ..., n$. The state of the population is thus given by the vector $\mathbf{x} \in \Delta_n$. A player with strategy $\mathbf{e}_i$ has as expected payoff

$$(A\mathbf{x})_i = \sum_j a_{ij} x_j. \tag{6.1}$$

Indeed, this player meets with probability x_j a co-player using $\mathbf{e}_j$. The average payoff in the population is given by

$$\mathbf{x} \cdot A\mathbf{x} = \sum_i x_i (A\mathbf{x})_i. \tag{6.2}$$

It should be stressed that we are committing an abuse of notation. The same symbol $\mathbf{x} \in \Delta_n$ which denoted in the previous sections the mixed strategy of one specific player (cf. (2.1) and (2.2)) 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 will consider this case only later.)

Now comes an essential step: we shall assume that populations can evolve, in the sense that the relative frequencies x_i change with time. Thus we let the state $\mathbf{x}(t)$ depend on time, and denote by $\dot{x}_i(t)$ the velocity with which x_i changes. The assumption of differentiability implies an infinitely large population, or the interpretation of x_i as an expected value, rather than a bona fide frequency. Both ways of thinking are familiar to mathematical ecologists. 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? There are many possibilities for modeling this process. We shall at first assume that the state of the population evolves according to the *replicator equation* (see [**40, 16, 46**] and, for the name, [**37**]). 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 $(A\mathbf{x})_i$ and the average payoff $\mathbf{x} \cdot A\mathbf{x}$ in the population. Thus we posit

$$\dot{x}_i = x_i[(A\mathbf{x}_i) - \mathbf{x} \cdot A\mathbf{x}] \tag{6.3}$$

for $i = 1, ..., n$. Accordingly, a strategy $\mathbf{e}_i$ 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. Before we try to motivate the replicator equation, 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 component of equation (6.3). Hence the hyperplanes $\sum x_i = 1$ and $x_i = 0$ are invariant. From this follows that the state space, i.e. the simplex Δ_n, is invariant: if $\mathbf{x}(0) \in \Delta_n$ then $\mathbf{x}(t) \in \Delta_n$ for all $t \in R$. The same holds for all sub-simplices of Δ_n (which are given by $x_i = 0$ for one or several i), and hence also for the boundary $bd\Delta_n$ of Δ_n (i.e. the union of all such sub-simplices), and moreover also for the interior $int\Delta_n$ of the simplex (the subset satisfying $x_i > 0$ for all i). From now on we only consider the restriction of (6.3) to the state simplex Δ_n.

7. Basic properties of the replicator equation

It is easy to see that if we add an arbitrary function $b(\mathbf{x})$ to all payoff terms $(A\mathbf{x})_i$, the replicator equation (6.3) 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 $j = 1, ..., n$) without altering the replicator dynamics on Δ_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

$$(\frac{x_i}{x_j})^{\cdot} = (\frac{x_i}{x_j})[(A\mathbf{x})_i - (A\mathbf{x})_j]. \tag{7.1}$$

Thus the relative proportions of two strategies change according to their payoff ranking. More generally, if $V = \prod x_i^{p_i}$ then

$$\dot{V} = V[\mathbf{p} \cdot A\mathbf{x} - (\sum p_i)\mathbf{x} \cdot A\mathbf{x}]. \tag{7.2}$$

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 $\mathbf{e}_i$ of the simplex Δ_n are rest points. (Obviously, if all players are of the same type, mere copying leads to no change.) The replicator equation admits a rest point in $int\Delta_n$ if there exists a solution (in $int\Delta_n$) of the linear equations

$$(A\mathbf{x})_1 = ... = (A\mathbf{x})_n. \tag{7.3}$$

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

One can show that if no rest point exists in the interior of Δ_n, then all orbits in $int\Delta_n$ converge to the boundary, for $t \to \pm\infty$. In particular, if strategy $\mathbf{e}_i$ is *strictly dominated*, i.e., if there exists a $\mathbf{w} \in \Delta_n$ such that $(A\mathbf{x})_i < \mathbf{w} \cdot A\mathbf{x}$ holds for all $\mathbf{x} \in \Delta_n$, then $x_i(t) \to 0$ for $t \to +\infty$ [**21**]. In the converse direction, if there exists an orbit $\mathbf{x}(t)$ bounded away from the boundary of Δ_n (i.e. such that for some $a > 0$ the inequality $x_i(t) > a$ holds for all $t > 0$ and all $i = 1, ..., n$), then there exists a rest point in $int\Delta_n$ [**18**]. One just has to note that for $i = 1, ..., n$,

$$(\log x_i)^{\cdot} = \dot{x}_i / x_i = (A\mathbf{x}(t))_i - \mathbf{x}(t) \cdot A\mathbf{x}(t). \tag{7.4}$$

Integrating for $t \in [0, T]$, 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 $T \to +\infty$. The corresponding limit on the right hand side implies that for the accumulation points z_i of the time averages

$$z_i(T) = \frac{1}{T} \int_0^T x_i(t) dt, \tag{7.5}$$

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

$$\sum a_{1j} z_j = ... = \sum a_{nj} z_j \tag{7.6}$$

must hold. Using (7.3), we see that $\mathbf{z}$ is a rest point in $int\Delta_n$.

8. The Lotka-Volterra connection

There is an intimate connection between Lotka-Volterra equations, which are the staple fare of mathematical population ecology, and the replicator equation [**18**]. More precisely, there exists a diffeomorphism from $\Delta_n^- = \{\mathbf{x} \in \Delta_n : x_n > 0\}$ onto R_+^{n-1} mapping the orbits of the replicator equation (6.3) onto the orbits of the Lotka-Volterra equation

$$\dot{y}_i = y_i (r_i + \sum_{j=1}^{n-1} d_{ij} y_j), \tag{8.1}$$

where $r_i = a_{in} - a_{nn}$ and $d_{ij} = a_{ij} - a_{nj}$. Indeed, let us define $y_n \equiv 1$ and consider the transformation $\mathbf{y} \to \mathbf{x}$ given by

$$x_i = \frac{y_i}{\sum_{j=1}^n y_j} \qquad i = 1, \ldots, n \tag{8.2}$$

which maps $\{\mathbf{y} \in R_+^n : y_n = 1\}$ onto Δ_n^-. The inverse $\mathbf{x} \to \mathbf{y}$ is given by

$$y_i = \frac{y_i}{y_n} = \frac{x_i}{x_n} \qquad i = 1, \ldots, n\ . \tag{8.3}$$

Now let us consider the replicator equation in n variables given by (6.3). We shall assume that the last row of the $n \times n$ matrix $A = (a_{ij})$ consists of zeros: since we can add constants to columns, this is no restriction of generality. By the quotient rule (7.1)

$$\dot{y}_i = (\frac{x_i}{x_n})[(A\mathbf{x})_i - (A\mathbf{x})_n]. \tag{8.4}$$

Since $(A\mathbf{x})_n = 0$, this implies

$$\dot{y}_i = y_i (\sum_{j=1}^n a_{ij} x_j) = y_i (\sum_{j=1}^n a_{ij} y_j) x_n\ . \tag{8.5}$$

By a change in velocity, we can remove the term $x_n > 0$. Since $y_n = 1$, this yields

$$\dot{y}_i = y_i(a_{in} + \sum_{j=1}^{n-1} a_{ij}y_j) \tag{8.6}$$

or (with $r_i = a_{in}$) equation (8.1).

The converse direction from (8.1) to (6.3) is analogous.

Results about Lotka-Volterra equations can therefore be carried over to the replicator equation and vice versa. Some properties are simpler to prove (or more natural to formulate) for one equation and some for the other.

For instance, it is easy to prove for the Lotka-Volterra equation that the interior of R^n_+ contains α- or ω- limit points if and only if it admits an interior rest point. Indeed, let $L : \mathbf{x} \to \mathbf{y}$ be defined by

$$y_i = r_i + \sum_{j=1}^{n} a_{ij}x_j \qquad i = 1, \dots, n\ . \tag{8.7}$$

If (8.1) admits no interior rest point, the set $K = L(\text{int } R^n_+)$ is disjoint from $\mathbf{0}$. A well known theorem from convex analysis implies that there exists a hyperplane H through $\mathbf{0}$ which is disjoint from the convex set K. Thus there exists a vector $\mathbf{c}= (c_1, \dots, c_n) \neq \mathbf{0}$ orthogonal to H (i.e. $\mathbf{c} \cdot \mathbf{x} = 0$ for all $\mathbf{x} \in H$) such that $\mathbf{c} \cdot \mathbf{y}$ is positive for all $\mathbf{y} \in K$. Setting

$$V(\mathbf{x}) = \prod x_i^{c_i}, \tag{8.8}$$

we see that V is defined on int R^n_+. If $\mathbf{x}(t)$ is a solution of (8.1) in int R^n_+, then the time derivative of $t \to V(\mathbf{x}(t))$ satisfies

$$\dot{V} = V \sum c_i \frac{\dot{x}_i}{x_i} = V \sum c_i y_i = V \mathbf{c} \cdot \mathbf{y} > 0\ . \tag{8.9}$$

Thus V is increasing along each orbit. But then no point $\mathbf{z} \in$ int R^n_+ may belong to an ω-limit: indeed, by Lyapunov's theorem, the derivative $\dot{V}$ would have to vanish there. This contradiction completes the proof.

In particular, if $int \Delta_n$ contains a periodic orbit of the replicator equation (6.3), it must also contain a rest point.

9. Two-dimensional examples

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

$$\begin{pmatrix} 0 & a \\ b & 0 \end{pmatrix}. \tag{9.1}$$

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

$$\dot{x} = x[(A\mathbf{x})_1 - \mathbf{x} \cdot A\mathbf{x}] = x[(A\mathbf{x})_1 - (x(A\mathbf{x})_1 + (1-x)(A\mathbf{x})_2)], \tag{9.2}$$

and hence

$$\dot{x} = x(1-x)[(A\mathbf{x})_1 - (A\mathbf{x})_2], \tag{9.3}$$

which reduces to

$$\dot{x} = x(1-x)[a-(a+b)x]. \tag{9.4}$$

We note that

$$a = \lim_{x \to 0} \frac{\dot{x}}{x}. \tag{9.5}$$

Hence a corresponds to the limit of the per capita growth rate of the missing strategy $\mathbf{e}_1$. Let us omit the trivial case $a = b = 0$: in this case all points of the state space Δ_2 (i.e. the interval $0 \le x \le 1$) 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 $]0,1[$ if and only if $ab > 0$. Thus we obtain three possible cases.

(1) There is no rest point in the interior of the state space. This happens if and only if $ab \le 0$. In this case, $\dot{x}$ has always the same sign in $]0,1[$. If this sign is positive (i.e. if $a \ge 0$ and $b \le 0$, at least one inequality being strict), this means that $x(t) \to 1$ for $t \to +\infty$, for every initial value $x(0)$ with $0 < x(0) < 1$. The strategy $\mathbf{e}_1$ is said to *dominate* strategy $\mathbf{e}_2$. It is always the best reply, for any value of $x \in]0,1[$. Conversely, if the sign of $\dot{x}$ is negative, then $x(t) \to 0$ and $\mathbf{e}_2$ dominates. In each case, the dominating strategy converges towards fixation.

As an example, we consider the Prisoner's Dilemma Game from (4.5). The two strategies $\mathbf{e}_1$ and $\mathbf{e}_2$ are usually interpreted as 'cooperation' (by providing a benefit to the co-player) and 'defection' (by refusing to provide a benefit). The payoff matrix is transformed, by adding appropriate constants to each column, into

$$\begin{pmatrix} 0 & -5 \\ 5 & 0 \end{pmatrix} \tag{9.6}$$

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 coordination game (4.3). The payoff matrix is transformed into

$$\begin{pmatrix} 0 & -2 \\ -2 & 0 \end{pmatrix} \tag{9.7}$$

and it is best to play $\mathbf{e}_1$ if the frequency of $\mathbf{e}_1$-players exceeds 50 percent. Bistability also occurs if the Prisoner's Dilemma game given by (4.5) is repeated sufficiently often. Let us assume that the number of rounds is a random variable with mean value m, for instance, and let us consider only two strategies of particular interest. One, which will be denoted by $\mathbf{e}_1$, is the Tit For Tat strategy which consists in cooperating in the first round and from then on imitating what the co-player did in the previous round. The other strategy, denoted as $\mathbf{e}_2$, consists in always defecting. The expected payoff values are given by the matrix

$$\begin{pmatrix} 10m & -5 \\ 15 & 0 \end{pmatrix} \tag{9.8}$$

which can be transformed into

$$\left(\begin{array}{cc} 0 & -5 \\ 15-10m & 0 \end{array}\right). \tag{9.9}$$

If $m > 3/2$, it is best to do what the co-player does. Loosely speaking, one should go with the trend. The outcome, in such a population, would be the establishment of a single norm of behavior (either always defect, or play Tit For Tat). Which norm emerges depends on the initial condition.

(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 $t \to +\infty$: the rest point $\hat{x}$ is a stable attractor. No strategy eliminates the other: rather, their frequencies converge towards a *stable coexistence.*

This situation can be found in the Hawk-Dove-game, for example. The payoff matrix (5.4)is transformed into

$$\left(\begin{array}{cc} 0 & \frac{G}{2} \\ \frac{C-G}{2} & 0 \end{array}\right) \tag{9.10}$$

and the rest point corresponds to $x = G/C$. The higher the cost of injury, i.e., C, the lower the frequency of escalation. Another well-known example is the so-called snowdrift game. Suppose that two players are promised 40 dollars each if they contribute 30 dollars to the experimenter. They have to decide independently whether to come up with such a fee or not. If both contribute, they can split the cost equally, and pay only 15 dollars. If $\mathbf{e}_1$ is the decision to contribute, and $\mathbf{e}_2$ not to contribute, the payoff matrix is

$$\left(\begin{array}{cc} 25 & 10 \\ 40 & 0 \end{array}\right) \tag{9.11}$$

which can be normalized to

$$\left(\begin{array}{cc} 0 & 10 \\ 15 & 0 \end{array}\right). \tag{9.12}$$

In this case, it is best to do the opposite of what the co-player is doing, i.e., to swim against the stream.

10. Rock-Scissors-Paper

Turning now to $n = 3$, we meet a particularly interesting example if the three strategies dominate each other in a cyclic fashion, i.e., if $\mathbf{e}_1$ dominates $\mathbf{e}_2$, in the absence of $\mathbf{e}_3$, and similarly $\mathbf{e}_2$ dominates $\mathbf{e}_3$, and $\mathbf{e}_3$, in turn, dominates $\mathbf{e}_1$. Such a cycle occurs in the game of Rock-Scissors-Paper shown in (1.2). It 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, $\mathbf{m} = (1/3, 1/3, 1/3) \in int\Delta_3$, and the other three at the vertices $\mathbf{e}_i$.

Let us consider the function $V := x_1 x_2 x_3$, which is positive in the interior of Δ_3 (with maximum at $\mathbf{m}$) and vanishes on the boundary. Using (7.2), we see that $t \to V(\mathbf{x}(t))$ satisfies

$$\dot{V} = V(x_2 - x_3 + x_3 - x_1 + x_1 - x_2) = 0. \tag{10.1}$$

Hence V is a *constant of motion*: all orbits $t \to \mathbf{x}(t)$ of the replicator equation remain on constant level sets of V. This implies that all orbits in $int\Delta_n$ 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 Δ_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 which can be arbitrarily rare, to make an arbitrarily small step. In the present case, it means for instance to flow along an edge from $\mathbf{e}_2$ towards $\mathbf{e}_1$, and then stepping onto the edge leading away from $\mathbf{e}_1$ and toward $\mathbf{e}_3$. This step can be arbitrarily small: travellers just have to wait until they are sufficiently close to the 'junction' $\mathbf{e}_1$.

Now let us consider the generalized Rock-Scissors-Paper game with matrix

$$\begin{pmatrix} 0 & a & -b \\ -b & 0 & a \\ a & -b & 0 \end{pmatrix}. \tag{10.2}$$

with $a, b > 0$, which is no longer zero-sum, if $a \neq b$. It has the same structure of cyclic dominance and the same rest points. The point $\mathbf{m}$ is a Nash equilibrium and the boundary of Δ_3 is a heteroclinic set, as before. But now,

$$\mathbf{x} \cdot A\mathbf{x} = (a-b)(x_1x_2 + x_2x_3 + x_3x_1), \tag{10.3}$$

and hence

$$\dot{V} = V(a-b)[1 - 3(x_1x_2 + x_2x_3 + x_3x_1)], \tag{10.4}$$

which implies

$$\dot{V} = \frac{V(a-b)}{2}[(x_1-x_2)^2 + (x_2-x_3)^2 + (x_3-x_1)^2]. \tag{10.5}$$

This expression vanishes on the boundary of Δ_3 and at $\mathbf{m}$. It has the sign of $a - b$ everywhere else on Δ_3. If $a > b$, this means that all orbits cross the constant-level sets of V in the uphill direction, and hence converge to $\mathbf{m}$. For $a > b$, the function $V(\mathbf{x})$ is a strict *Lyapunov function*: indeed $\dot{V}(\mathbf{x}) \geq 0$ for all $\mathbf{x}$, and equality holds only when $\mathbf{x}$ is a rest point. This implies that ultimately, all three types will be present in the population in equal frequencies: the rest point $\mathbf{m}$ is asymptotically stable. But for $a < b$, the orbits flow downhill towards the boundary of Δ_3. The Nash equilibrium $\mathbf{m}$ corresponds to an unstable rest point, and the heteroclinic cycle on the boundary attracts all other orbits.

Let us follow the state $\mathbf{x}(t)$ of the population, for $a < b$. 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.

The same holds if all the a's and b's, in (10.2), are distinct positive numbers. There exists a unique rest point $\mathbf{m}$ in the interior of Δ_3 which, depending on the sign of $\det A$ (which is the same as that of $\mathbf{m} \cdot A\mathbf{m}$) is either globally stable, i.e., attracts all orbits in $int\Delta_3$, or is surrounded by periodic orbits, or is repelling. In the latter case, all orbits converge to the heteroclinic cycle formed by the boundary of Δ_n.

Interestingly, several biological examples for Rock-Scissors-Paper cycles have been found. We only mention two examples: (A) Among the lizard species *Uta stansburiana*, three inheritable types of male mating behavior are $\mathbf{e}_1$: attach yourself to a female and guard her closely, $\mathbf{e}_2$: attach yourself to several females and guard them (but inevitably, less closely); and $\mathbf{e}_3$: attach yourself to no female, but roam around and attempt sneaky matings whenever you encounter an unguarded female [**39**]. (B) Among the bacteria *E. coli*, three strains occur in the lab through recurrent mutations, namely $\mathbf{e}_1$: the usual, so-called *wild type*; $\mathbf{e}_2$: a mutant producing colicin, a toxic substance, together with a protein conferring auto-immunity; and $\mathbf{e}_3$: a mutant producing the immunity-conferring protein, but not the poison [**23**]. In case (A), selection leads to the stable coexistence of all three types, and in case (B) to the survival of one type only.

There exist about 100 distinct phase portraits of the replicator equation for $n = 3$, up to re-labeling the vertices [**1**]. Of these, about a dozen are generic. Interestingly, none admits a limit cycle [**19**]. For $n > 3$, limit cycles and chaotic attractors can occur. A classification seems presently out of reach.

11. Nash equilibria and saturated rest points

Let us consider a symmetric $n \times n$-game with payoff matrix A, and $\mathbf{z}$ a symmetric Nash equilibrium. With $\mathbf{x} = \mathbf{e}_i$, condition (5.2) implies

$$(A\mathbf{z})_i \leq \mathbf{z} \cdot A\mathbf{z} \tag{11.1}$$

for $i = 1, ..., n$. Equality must hold for all i such that $z_i > 0$. Hence $\mathbf{z}$ 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

$$(A\mathbf{z})_i - \mathbf{z} \cdot A\mathbf{z} \leq 0. \tag{11.2}$$

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

Every rest point in $int\Delta_n$ is trivially saturated; but on the boundary, there may be rest points which are not saturated, as we shall presently see. In that case, there exist strategies which are not present in the population $\mathbf{z}$, but which 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 $(A\mathbf{z})_i = c$ whenever $z_i > 0$; the additional requirement, for a Nash equilibrium, is that $(A\mathbf{z})_i \leq c$ 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 rest points in the interior of the simplex are Nash equilibria. In case $n = 2$ and dominance, the strategy that is dominant is a Nash equilibrium, and the other is not. In case $n = 2$ with bi-stability, both pure strategies are Nash equilibria. Generically (and in contrast to the example (9.7)), one of the pure strategies fares better than the other in a population where both are equally frequent. This is the so-called risk-dominant equilibrium. It has the larger basin of attraction. In the case $n = 2$ leading to stable co-existence, none of the pure strategies is a Nash equilibrium. If you play a bistable game, you should 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 case $n = 3$ with the Rock-Scissors-Paper structure, the interior rest point $\mathbf{m}$ is the unique Nash equilibrium. Each of the vertex rest points can be invaded.

A handful of results about Nash equilibria and rest points of the replicator dynamics are known as *folk theorem of evolutionary game theory* [**5**]. For instance, any limit, for $t \to +\infty$, of a solution $\mathbf{x}(t)$ starting in $int\Delta_n$ 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, then 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 $\mathbf{x}(t) \to \mathbf{z}$ for $t \to +\infty$). The converse statements are generally not valid.

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 (6.3), 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

$$\dot{x}_i = x_i[(A\mathbf{x})_i - \mathbf{x} \cdot A\mathbf{x} - n\epsilon] + \epsilon. \tag{11.3}$$

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 Δ_n (which is invariant for the unperturbed replicator equation), the vector field of the perturbed equation points towards the interior.

Brouwer's fixed point theorem implies that (11.3) admits at least one rest point in $int\Delta_n$, which we denote by $\mathbf{z}_\epsilon$. It satisfies

$$(A\mathbf{z}_\epsilon)_i - \mathbf{z}_\epsilon \cdot A\mathbf{z}_\epsilon = \epsilon(n - \frac{1}{(\mathbf{z}_\epsilon)_i}). \tag{11.4}$$

Let ϵ tend to 0, and let $\mathbf{z}$ be an accumulation point of the $\mathbf{z}_\epsilon$ in Δ_n. The limit on the left hand side exists, and is given by $(A\mathbf{z})_i - \mathbf{z} \cdot A\mathbf{z}$. Hence the right hand side also has a limit for $\epsilon \to 0$. 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 (6.3), and hence corresponds to a Nash equilibrium (see also [**15, 38**]).

12. Mixed strategies and evolutionary stability

Let us now consider the case when individuals can also use mixed strategies, for instance escalate a conflict with a certain probability. Thus let us assume that there exist N types, each using a (pure or mixed) strategy $\mathbf{p}(i) \in \Delta_n$ (we need not assume $n = N$). The average payoff for a $\mathbf{p}(i)$-player against a $\mathbf{p}(j)$-player is given by $u_{ij} = \mathbf{p}(i) \cdot A\mathbf{p}(j)$, and if $\mathbf{x} \in \Delta_N$ describes the frequencies of the types in the population, then the average strategy within the population is $\mathbf{p}(\mathbf{x}) = \sum x_i \mathbf{p}(i)$. The induced replicator dynamics on Δ_N, namely $\dot{x}_i = x_i[(U\mathbf{x})_i - \mathbf{x} \cdot U\mathbf{x}]$ can be written as

$$\dot{x}_i = x_i[(\mathbf{p}(i) - \mathbf{p}(\mathbf{x})) \cdot A\mathbf{p}(\mathbf{x})]. \tag{12.1}$$

This dynamics on Δ_N induces a dynamics $t \to \mathbf{p}(\mathbf{x}(t))$ of the average strategy on Δ_n.

Let us now turn to the concept of an evolutionarily stable strategy, or ESS. If all members of the population use such a strategy $\hat{\mathbf{p}} \in \Delta_n$, then no mutant minority using another strategy $\mathbf{p}$ can invade (cf. [**29, 25**]). Thus a strategy $\hat{\mathbf{p}} \in \Delta_n$ is said to be *evolutionarily stable* if for every $\mathbf{p} \in \Delta_n$ with $\mathbf{p} \neq \hat{\mathbf{p}}$, the induced replicator equation describing the dynamics of the population consisting of these two types only (the 'resident' using $\hat{\mathbf{p}}$ and the 'invader' using $\mathbf{p}$) leads to the elimination of the invader. By (9.4) this equation reads (if x is the frequency of the invader):

$$\dot{x} = x(1-x)[x(\mathbf{p}\cdot A\mathbf{p} - \hat{\mathbf{p}}\cdot A\mathbf{p}) - (1-x)(\hat{\mathbf{p}}\cdot A\hat{\mathbf{p}} - \mathbf{p}\cdot A\hat{\mathbf{p}})] \tag{12.2}$$

and hence the rest point $x = 0$ is asymptotically stable iff the following conditions are satisfied:

(a) (equilibrium condition)

$$\mathbf{p}\cdot A\hat{\mathbf{p}} \le \hat{\mathbf{p}}\cdot A\hat{\mathbf{p}} \tag{12.3}$$

holds for all $\mathbf{p} \in \Delta_n$;

(b) (stability condition)

$$\text{if} \quad \mathbf{p}\cdot A\hat{\mathbf{p}} = \hat{\mathbf{p}}\cdot A\hat{\mathbf{p}} \quad \text{then} \quad \mathbf{p}\cdot A\mathbf{p} < \hat{\mathbf{p}}\cdot A\mathbf{p}. \tag{12.4}$$

The first condition means that $\hat{\mathbf{p}}$ is a Nash equilibrium: no invader does better than the resident, against the resident. The second condition states that if the invader does as well as the resident against the resident, then it does less well than the resident against the invader. Based on (7.2), it can be shown that the strategy $\hat{\mathbf{p}}$ is an ESS iff $\prod_i x_i^{\hat{p}_i}$ is a strict local Lyapunov function for the replicator equation, or equivalently iff

$$\hat{\mathbf{p}}\cdot A\mathbf{p} > \mathbf{p}\cdot A\mathbf{p} \tag{12.5}$$

for all $\mathbf{p} \neq \hat{\mathbf{p}}$ in some neighborhood of $\hat{\mathbf{p}}$ [**16, 18**]. If $\hat{\mathbf{p}} \in int\Delta_n$, then Δ_n itself is such a neighborhood.

In particular, an ESS corresponds to an asymptotically stable rest point of (6.3). The converse does not hold in general [**46**]. But the strategy $\hat{\mathbf{p}} \in \Delta_n$ is an ESS iff it is *strongly stable* in the following sense: whenever it belongs to the convex hull of $\mathbf{p}(1), ..., \mathbf{p}(N) \in \Delta_n$, the strategy $\mathbf{p}(\mathbf{x}(t))$ converges to $\hat{\mathbf{p}}$, under (12.1), for all $\mathbf{x} \in \Delta_N$ for which $\mathbf{p}(\mathbf{x})$ is sufficiently close to $\hat{\mathbf{p}}$ [**4**].

The relation between evolutionary and dynamic stability is particularly simple for the class of *partnership games*. These are defined by payoff matrices $A = A^T$. In this case the interests of both players coincide. For spartnership games, $\hat{\mathbf{p}}$ is an ESS iff it is asymptotically stable for (6.3). This in turn holds iff it is a strict local maximum of the average payoff $\mathbf{x}\cdot A\mathbf{x}$ [**18**]. Replicator equations for partnership games occur prominently in population genetics. They describe the effect of selection on the frequencies x_i of alleles i on a single genetic locus, for $i \in \{1, ..., n\}$. In this case, the a_{ij} correspond to the survival probabilities of individuals with genotype (i, j) (i.e., having inherited the alleles i and j from their parents).

13. Generalizations of the replicator dynamics

We have assumed so far that the average payoff for a player using strategy i is given by a linear function $(A\mathbf{x})_i$ of the state of the population. This makes sense if the interactions are pairwise, with co-players chosen randomly within the population. But many interesting examples lead to non-linear payoff functions $a_i(\mathbf{x})$, for instance if the interactions occur in groups with more than two members. This is the case, for instance, in the sex-ratio game, where the success of a strategy (i.e., an individual sex ratio) depends on the aggregate sex ratio in the population. Nonlinear payoff functions $a_i(\mathbf{x})$ lead to the replicator equation

$$\dot{x}_i = x_i(a_i(\mathbf{x}) - \bar{a}) \tag{13.1}$$

on Δ_n, where $\bar{a} = \sum_i x_i a_i(\mathbf{x})$ is again the average payoff within the population. Many of the previous results can be extended in a straightforward way. For instance, the dynamics is unchanged under addition of a function b to all payoff functions a_i. Equation (13.1) always admits a saturated rest point, and a straight extension of the folk theorem is still valid. The notion of an ESS has to be replaced by a localized version.

Initially, the replicator dynamics was intended to model the transmission of behavioral programs through inheritance. The simplest inheritance mechanisms lead in a straightforward way to (6.3), but more complex cases of Mendelian inheritance through one or several genetic loci yield more complex dynamics [**13, 7, 45, 17**]. The replicator equation (6.3) can also be used to model imitation processes [**14, 2, 36, 34**]. A rather general approach to modeling imitation processes leads to

$$\dot{x}_i = x_i[f(a_i(\mathbf{x})) - \sum x_j f(a_j(\mathbf{x}))] \tag{13.2}$$

for some strictly increasing function f of the payoff, and even more generally to the *imitation dynamics* given by

$$\dot{x}_i = x_i g_i(\mathbf{x}) \tag{13.3}$$

where the functions g_i satisfy $\sum x_i g_i(\mathbf{x}) = 0$ on Δ_n. The simplex Δ_n and its faces are invariant. Such an equation is said to be *payoff monotonic* if

$$g_i(\mathbf{x}) > g_j(\mathbf{x}) \Leftrightarrow a_i(\mathbf{x}) > a_j(\mathbf{x}), \tag{13.4}$$

where the a_i correspond to the payoff for strategy i. For payoff monotonic equations (13.3), the *folk theorem* holds again [**31, 8**]: Nash equilibria are rest points, strict Nash equilibria are asymptotically stable, and rest points that are stable or ω-limits of interior orbits are Nash equilibria.

The dynamics (13.3) can be reduced (through a change in velocity) to a replicator equation (13.1) if it has the following property:

$$\mathbf{y} \cdot \mathbf{g}(\mathbf{x}) > \mathbf{z} \cdot \mathbf{g}(\mathbf{x}) \Longleftrightarrow \mathbf{y} \cdot \mathbf{a}(\mathbf{x}) > \mathbf{z} \cdot \mathbf{a}(\mathbf{x}) \tag{13.5}$$

for all $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \Delta_n$.

14. Best reply dynamics

It is worth emphasizing that imitation (like selection, in genetics) does not produce anything new. If a strategy $\mathbf{e}_i$ 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). An equation such as (13.1) or more generally (13.3) does not allow the introduction of new strategies. There exist game dynamics which are more innovative. For instance, clever players could adopt the strategy which offers the highest payoff, even if no one in the population is currently using it. We describe this dynamics presently. 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 [**35**].

The best-reply dynamics assumes more sophistication than mere learning by copying others. Let us assume that in a large population, a small fraction of the players revise their strategy, choosing best replies $BR(\mathbf{x})$ to the current mean population strategy $\mathbf{x}$. This approach, which postulates that players are intelligent enough to know the current population state and to respond optimally, yields the *best reply dynamics*

$$\dot{\mathbf{x}} \in BR(\mathbf{x}) - \mathbf{x}. \tag{14.1}$$

Since best replies are in general not unique, this is a differential *inclusion* rather than a differential equation [**26**]. For continuous payoff functions $a_i(\mathbf{x})$, the set of best replies $BR(\mathbf{x})$ is a non-empty convex, compact subset of Δ_n which is upper semi-continuous in $\mathbf{x}$. Hence solutions exist, they are Lipschitz functions $\mathbf{x}(t)$ satisfying (14.1) for almost all $t \geq 0$. If $BR(\mathbf{x})$ is a uniquely defined (and hence pure) strategy $\mathbf{b}$, the solution of (14.1) is given by

$$\mathbf{x}(t) = (1 - e^{-t})\mathbf{b} + e^{-t}\mathbf{x} \tag{14.2}$$

for small $t \geq 0$, which describes a linear orbit pointing straight towards the best response. This can lead to a state where $\mathbf{b}$ is no longer the unique best reply. But for each $\mathbf{x}$ there always exists a $\mathbf{b} \in BR(\mathbf{x})$ which, among all best replies to $\mathbf{x}$, is a best reply against itself (i.e. a Nash equilibrium of the game restricted to the simplex $BR(\mathbf{x})$) [**20**]. In this case $\mathbf{b} \in BR((1-\epsilon)\mathbf{x} + \epsilon\mathbf{b})$ holds for small $\epsilon \geq 0$, if the game is linear. An iteration of this construction yields at least one piecewise linear solution of (14.1) starting at $\mathbf{x}$ and defined for all $t > 0$. One can show that for generic linear games, essentially all solutions can be constructed in this way. For the resulting (multi-valued) semi-dynamical system, the simplex Δ_n is only forward invariant and $bd\Delta_n$ need no longer be invariant: the frequency of strategies which are initially missing can grow, in contrast to the imitation dynamics. In this sense, the best reply dynamics is an *innovative* dynamics.

For $n = 2$, the phase portraits of (14.1) differ only in details from that of the replicator dynamics. If $\mathbf{e}_1$ is dominated by $\mathbf{e_2}$, there are only two orbits: the rest point $\mathbf{e}_2$, and the semi-orbit through $\mathbf{e}_1$ which converges to $\mathbf{e}_2$. In the bistable situation with interior Nash equilibrium $\mathbf{p}$, there are infinitely many solutions starting at $\mathbf{p}$ besides the constant one, staying there for some time and then converging monotonically to either $\mathbf{e}_1$ or $\mathbf{e}_2$. In the case of stable coexistence with interior Nash equilibrium $\mathbf{p}$, the solution starting at some point $\mathbf{x}$ between $\mathbf{p}$ and $\mathbf{e}_1$ converges toward $\mathbf{e}_2$ until it hits $\mathbf{p}$, in finite time, and then remains there forever.

For $n = 3$, the differences to the replicator dynamics become more pronounced. In particular, for the generalized Rock-Scissors-Paper game given by (10.2), all orbits converge to the Nash equilibrium $\mathbf{p}$ whenever $\det A > 0$ (just as with the replicator dynamics); but for $\det A < 0$, all orbits (except possibly $\mathbf{p}$) converge to a limit cycle, the so-called *Shapley triangle* spanned by the three points A_i (given by the intersections of the lines $(A\mathbf{x})_2 = (A\mathbf{x})_3$ etc. in Δ_3). In fact, the piecewise linear function $V(\mathbf{x}) := |\max_i (A\mathbf{x})_i|$ is a Lyapunov function for (14.1). In this case, the orbits of the replicator equation (6.3) converge to the boundary of Δ_n; but interestingly, the time averages

$$\mathbf{z}(T) := \frac{1}{T}\int_0^T \mathbf{x}(t)dt \tag{14.3}$$

have the Shapley triangle as the set of accumulation points, for $T \to +\infty$. Similar parallels between the best reply dynamics and the behavior of time-averages of the replicator equation are quite frequent [**9, 10**].

15. A brief look at asymmetric games

So far, we have considered evolutionary games in the symmetric case only. Thus players are indistinguishable (except by their strategies), and the game is described by a single $n \times n$ payoff matrix A. In the first section, however, we had started out with two players I and II having strategies $\mathbf{e}_i$ and $\mathbf{f}_j$ respectively (with $1 \leq i \leq n$ and $1 \leq j \leq m$), and a game was defined by two $n \times m$ payoff matrices A and B There is an obvious way to turn the 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 $(\mathbf{e}_i, \mathbf{f}_j)$. A mixed strategy is given by an element $\mathbf{z} = (z_{ij}) \in \Delta_{nm}$, where z_{ij} denotes the probability to play $\mathbf{e}_i$ when in role I and $\mathbf{f}_j$ when in role II. To the probability distribution $\mathbf{z}$ correspond its *marginals*: $x_i = \sum_j z_{ij}$ and $y_j = \sum_i z_{ij}$. The vectors $\mathbf{x} = (x_i)$ and $\mathbf{y} = (y_j)$ belong to Δ_n and Δ_m, respectively.

The expected payoff for a player using $(\mathbf{e}_i, \mathbf{f}_j)$ against a player using $(\mathbf{e}_k, \mathbf{f}_l)$, with $i, k \in \{1, ..., n\}$ and $j, l \in \{1, ..., m\}$, is given by

$$c_{ij,kl} = \frac{1}{2} a_{il} + \frac{1}{2} b_{kj}. \tag{15.1}$$

Since every symmetric game has a symmetric Nash equilibrium, it follows immediately that every game (A, B) has a Nash equilibrium pair.

Let us now turn to population games. Players meet randomly and engage in a game (A, B), with chance deciding who is in role I and who in role II. For simplicity, we assume that there are only two strategies for each role. The payoff matrix is

$$\begin{pmatrix} (A, a) & (B, b) \\ (C, c) & (D, d) \end{pmatrix}. \tag{15.2}$$

The strategies for the resulting symmetric game will be denoted by $\mathbf{G}_1 = \mathbf{e}_1\mathbf{f}_1$, $\mathbf{G}_2 = \mathbf{e}_2\mathbf{f}_1$, $\mathbf{G}_3 = \mathbf{e}_2\mathbf{f}_2$ and $\mathbf{G}_4 = \mathbf{e}_1\mathbf{f}_2$. The payoff for a player using $\mathbf{G}_i$ against a player using $\mathbf{G}_j$ is given, up to the factor $1/2$ which we shall henceforth omit, by the (i, j)-entry of the matrix

$$M = \begin{pmatrix} A+a & A+c & B+c & B+a \\ C+a & C+c & D+c & D+a \\ C+b & C+d & D+d & D+b \\ A+b & A+d & B+d & B+b \end{pmatrix}. \tag{15.3}$$

This corresponds to (15.1). For instance, a $\mathbf{G}_1$-player meeting a $\mathbf{G}_3$-opponent is with probability $1/2$ in role I, plays $\mathbf{e}_1$ against the co-player's $\mathbf{f}_2$, and obtains B. With probability $1/2$, the $\mathbf{G}_1$-player is in role II, plays $\mathbf{f}_1$ against the co-players' $\mathbf{e}_2$, and obtains c.

The replicator dynamics

$$\dot{x}_i = x_i[(M\mathbf{x})_i - \mathbf{x} \cdot M\mathbf{x}] \tag{15.4}$$

describes the evolution of the state $\mathbf{x} = (x_1, x_2, x_3, x_4) \in \Delta_4$. Since the dynamics is unaffected if each m_{ij} is replaced by $m_{ij} - m_{1j}$ (for $i, j \in \{1, 2, 3, 4\}$), we can use the matrix

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ R & R & S & S \\ R+r & R+s & S+s & S+r \\ r & s & s & r \end{pmatrix} \tag{15.5}$$

with $R := C - A$, $r := b - a$, $S := D - B$ and $s := d - c$. We shall denote this matrix again by M. It has the property that

$$m_{1j} + m_{3j} = m_{2j} + m_{4j} \tag{15.6}$$

for $j = 1, 2, 3, 4$. Hence

$$(M\mathbf{x})_1 + (M\mathbf{x})_3 = (M\mathbf{x})_2 + (M\mathbf{x})_4 \tag{15.7}$$

holds for all $\mathbf{x}$. From this and (7.2) follows that the function $V = x_1x_3/x_2x_4$ satisfies

$$\dot{V} = V[(M\mathbf{x})_1 + (M\mathbf{x})_3 - (M\mathbf{x})_2 - (M\mathbf{x})_4] = 0 \tag{15.8}$$

in the interior of Δ_4, and hence that V is an invariant of motion for the replicator dynamics: its value remains unchanged along every orbit.

Therefore, the interior of the state simplex Δ_4 is foliated by the surfaces

$$W_K := \{\mathbf{x} \in \Delta_4 : x_1x_3 = Kx_2x_4\}, \tag{15.9}$$

with $0 < K < \infty$. These are saddle-like surfaces which are spanned by the quadrangle of edges $\mathbf{G}_1\mathbf{G}_2$, $\mathbf{G}_2\mathbf{G}_3$, $\mathbf{G}_3\mathbf{G}_4$ and $\mathbf{G}_4\mathbf{G}_1$ joining the vertices of the simplex Δ_4.

The orientation of the flow on the edges can easily be obtained from the previous matrix. For instance, if $R = 0$, then the edge $\mathbf{G}_1\mathbf{G}_2$ consists of rest points. If $R > 0$, the flow along the edge points from $\mathbf{G}_1$ towards $\mathbf{G}_2$ (which means that in the absence of the strategies $\mathbf{G}_3$ and $\mathbf{G}_4$, the strategy $\mathbf{G}_2$ dominates $\mathbf{G}_1$), and conversely, if $R < 0$, the flow points from $\mathbf{G}_2$ to $\mathbf{G}_1$.

Generically, the parameters R, S, r and s are non-zero. This corresponds to 16 orientations of the quadrangle $\mathbf{G}_1\mathbf{G}_2\mathbf{G}_3\mathbf{G}_4$, which by symmetry can be reduced to 4. Since $(M\mathbf{x})_1$ trivially vanishes, the rest points in the interior of the simplex Δ_4 must satisfy $(M\mathbf{x})_i = 0$ for $i = 2, 3, 4$. This implies for $S \neq R$

$$x_1 + x_2 = \frac{S}{S - R}, \tag{15.10}$$

and for $s \neq r$

$$x_1 + x_4 = \frac{s}{s - r}. \tag{15.11}$$

Such solutions lie in the simplex if and only if $RS < 0$ and $rs < 0$. If this is the case, one obtains a line of rest points which intersects each W_K in exactly one point. These points can be written as

$$x_i = m_i + \xi \tag{15.12}$$

for $i = 1, 3$ and

$$x_i = m_i - \xi \tag{15.13}$$

for $i = 2, 4$, with ξ as parameter and

$$\mathbf{m} = \frac{1}{(S - R)(s - r)}(Ss, -Sr, Rr, -Rs) \in W_1. \tag{15.14}$$

Of particular interest is the so-called Wright-manifold W_1, where the strategies, in the two roles, are independent of each other. (On W_1, the probability that a randomly chosen individual uses strategy $\mathbf{e}_1\mathbf{f}_1$ is the product of the probabilities $x := x_1 + x_4$ and $y := x_1 + x_2$ of choosing $\mathbf{e}_1$ when in role I, resp. $\mathbf{f}_1$ when in role II. Indeed, $x_1 = (x_1 + x_4)(x_1 + x_2)$). It then follows that

$$\dot{x} = x(1-x)(s-(s-r)y), \tag{15.15}$$

and

$$\dot{y} = y(1-y)(S-(S-R)x). \tag{15.16}$$

If $rR > 0$, each interior rest point is a saddle point within the corresponding manifold W_K, and the system is bistable: depending of the inital condition, orbits converge either to G_1 or to G_3, if $r < 0$, and either to G_2 or to G_4, if $r > 0$. If $rR < 0$, each rest point has (in addition of the eigenvalue 1) a pair of complex conjugate eigenvalues. Within the corresponding manifold W_K, the eigenvalues spiral around this rest point. Depending on whether K is larger or smaller than 1, they either converge to the rest point (which must be a spiral sink), or else toward the heteroclinic cycle defined by the quadrangle of the edges forming the boundary of W_K. For $K = 1$, the orbits are periodic.

16. Applications

In this lecture course, the authors aim to stress the variety of plausible dynamics which describe adaptive mechanisms underlying game theory. The replicator equation and the best reply dynamics describe just two out of many dynamics. For applications of evolutionary game theory, it does not suffice to specify the strategies and the payoff values. One also has to be explicit about the transmission mechanisms describing how strategies spread within a population.

We end this introductory part with some signposts to the literature using evolutionary games to model specific social interactions. The first applications, and indeed the motivation, of evolutionary game theory are found in evolutionary biology, where by now thousands of papers have proved the fruitfulness of this approach, see [**6**]. In fact, questions of sex-ratio, and more generally of sex-allocation, even pre-date any explicit formulation in terms of evolutionary game theory. It was R.F. Fisher, a pioneer in both population genetics and mathematical statistics, who used frequency-dependent selection to explain the prevalence of a 1:1 sex ratio, and W.D. Hamilton who extended this type of thinking to make sense of other, odd sex ratios [**12**]. We have seen how Price and Maynard Smith coined their concept of evolutionary stability to explain the prevalence of ritual fighting in intraspecific animal contests. The subtleties of such contests are still a favorite topic among the students of animal behavior. More muted, but certainly not less widespread conflicts arise on the issues of mate choice, parental investment, and parent-offspring conflicts. Social foraging is another field where the success of a given behavior (scrounging, for instance) depends on its prevalence; so are dispersal and habitat selection. Communication (alarm calls, threat displays, sexual advertisement, gossip), with all its opportunities for deceit, is replete with game theoretical problems concerning bluff and honest signaling. Predators and their prey, or parasites and their hosts, offer examples of games between two populations, with the success of a trait depending on the state of the other population. Some strategic interactions are surprisingly sophisticated, considering the lowly level of the players: for instance, bacteria can engage in quorum sensing as cue for conditional behavior.

Quite a few biological games turned out to have the same structure as games that had been studied by economists, usually under another name [**3**]: the biologists' 'Hawk-Dove' game, for example, has the same structure as the economists' 'Chicken'-game. Evolutionary game theory has found a large number of applications in economic interactions [**44, 22, 41, 8, 11**].

One zone of convergence for studies of animal behavior and human societies is that of cooperation. Indeed, the theory of evolution and economic theory have each their own paradigm of selfishness, encapsulated in the slogans of the 'selfish gene' and the 'homo economicus'. Both paradigms conflict with wide-spread evidence of social, 'other-regarding' behavior. In ant and bee societies, the relatedness of individuals is so close that their genetic interests overlap and their communities can be viewed as 'super-organisms'. But in human societies, close cooperation can also occur between individuals who are unrelated. In many cases, such cooperation is based on reciprocation. Positive and negative incentives, and in particular the threat of sanctions offer additional reasons for the prevalence of cooperation [**38**]. This may lead to two or more stable equilibria, corresponding to behavioral norms. If everyone adopts a given norm, no player has an incentive to deviate. But which of these norms eventually emerges depends, among other things, on the history of the population.

Animal behavior and experimental economics fuse in this area. Experimental economics, has greatly flourished in the last few years. It often reduces to the investigation of very simple games which can be analyzed by means of evolutionary dynamics. These and other games display the limitations of 'rational' behavior in humans, and have assisted in the emergence of new fields, such as behavioral game theory and neuro-economics.

References

1. I.M. Bomze, *Non-cooperative two-person games in biology: a classification*, Int. J. Game Theory **15** (1986), 31-57.
2. T. Börgers and R. Sarin, *Learning through reinforcement and replicator dynamics*, J. Economic Theory **77** (1997), 1-14.
3. A.M. Colman, Game Theory and its Applications in the Social and Biological Sciences, Oxford: Butterworth-Heinemann (1995).
4. R. Cressman, The Stability concept of Evolutionary Game Theory, Springer, Berlin (1992).
5. R. Cressman, Evolutionary Dynamics and Extensive Form Games, MIT Press (2003)
6. L.A. Dugatkin and H.K. Reeve (eds.), Game Theory and Animal Behavior, Oxford UP (1998).
7. I. Eshel, *Evolutionarily stable strategies and viability selection in Mendelian populations*, Theor. Population Biology **22** (1982), 204-217.
8. D. Friedman, *Evolutionary games in economics*, Econometrica **59** (1991), 637-66.
9. A. Gaunersdorfer, *Time averages for heteroclinic attractors*, SIAM J. Appl. Math **52** (1992), 1476-89.
10. A. Gaunersdorfer and J. Hofbauer, *Fictitious play, Shapley polygons and the replicator equation*, Games and Economic Behavior **11** (1995), 279-303.
11. H. Gintis, Game Theory Evolving, Princeton UP (2000).
12. W.D. Hamilton, *Extraordinary sex ratios*, Science **156**, 477-488.
13. P. Hammerstein and R. Selten, *Game theory and evolutionary biology*, in R.J. Aumann, S. Hart (eds.), Handbook of Game Theory II, Amsterdam, North-Holland (1994), 931-993.
14. D. Helbing, *Interrelations between stochastic equations for systems with pair interactions*, Physica A **181** (1992), 29-52.

15. J. Hofbauer, *From Nash and Brown to Maynard Smith: equilibria, dynamics and ESS*, Selection **1** (2000), 81-88.
16. J. Hofbauer, P. Schuster, and K. Sigmund, *A note on evolutionarily stable strategies and game dynamics*, J. Theor. Biology **81** (1979), 609-612.
17. J. Hofbauer, P. Schuster, and K. Sigmund, *Game dynamics for Mendelian populations*, Biol. Cybernetics **43** (1982), 51-57.
18. J. Hofbauer and K. Sigmund, The Theory of Evolution and Dynamical Systems, Cambridge UP (1988).
19. J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge UP (1998).
20. J. Hofbauer and K. Sigmund, *Evolutionary game dynamics*, Bulletin of the American Mathematical Society **40**, (2003) 479-519.
21. J. Hofbauer and J. W. Weibull, *Evolutionary selection against dominated strategies*, J. Economic Theory **71** (1996), 558-573.
22. M. Kandori: *Evolutionary Game Theory in Economics*, in D. M. Kreps and K. F. Wallis (eds.), Advances in Economics and Econometrics: Theory and Applications, I, Cambridge UP (1997).
23. B. Kerr, M.A. Riley, M.W. Feldman, and B.J.M. Bohannan, *Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors*, Nature **418** (2002), 171-174.
24. R. Leonard, Von Neumann, Morgenstern and the Creation of Game Theory: from Chess to Social Science, 1900-1960, Cambridge, Cambridge UP (2010).
25. S. Lessard, *Evolutionary stability: one concept, several meanings*, Theor. Population Biology **37** (1990), 159-70.
26. A. Matsui, *Best response dynamics and socially stable strategies*, J. Econ. Theory **57** (1992), 343-362.
27. J. Maynard Smith and G. Price, *The logic of animal conflict*, Nature **246** (1973), 15-18.
28. J. Maynard Smith, *Will a sexual population converge to an ESS?*, American Naturalist **177** (1981), 1015-1018.
29. J. Maynard Smith, Evolution and the Theory of Games, Cambridge UP (1982)
30. R. Myerson, Game Theory: Analysis of Conflict, Cambridge, Mass., Harvard University Press (1997)
31. J. Nachbar, *"Evolutionary" selection dynamics in games: convergence and limit properties*, Int. J. Game Theory **19** (1990), 59-89.
32. S. Nasar, A Beautiful Mind: A Biography of John Forbes Nash, Jr., Winner of the Nobel Prize in Economics, New York, Simon and Schuster (1994).
33. J. Nash, *Non-cooperative games*, Ann. Math. **54** (1951), 287-295.
34. M.A. Nowak, Evolutionary Dynamics, Cambridge MA, Harvard UP (2006).
35. W.H. Sandholm, Population Games and Evolutionary Dynamics, Cambridge, MA, MIT Press (2010).
36. K.H. Schlag, *Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits*, J. Econ. Theory **78** (1997), 130-156.
37. P. Schuster and K. Sigmund, *Replicator Dynamics*, J. Theor. Biology **100** (1983), 533-538.
38. K. Sigmund, The Calculus of Selfishness, Princeton, Princeton UP (2010).
39. B. Sinervo and C.M. Lively, *The rock-paper-scissors game and the evolution of alternative male strategies*, Nature **380** (1996), 240-243.
40. P.D. Taylor and L. Jonker, *Evolutionarily stable strategies and game dynamics*, Math. Biosciences **40** (1978), 145-156.
41. F. Vega-Redondo, Evolution, Games, and Economic Theory, Oxford UP (1996).
42. J. von Neumann, *Zur Theorie der Gesellschaftsspiele*, Mathematische Annalen **100** (1928), 295-320.
43. J. von Neumann and O. Morgenstern *Theory of Games and Economic Behavior*, Princeton UP (1944).
44. J. Weibull, Evolutionary Game Dynamics, MIT Press, Cambridge, Mass. (1995).
45. F. Weissing: *Evolutionary stability and dynamic stability in a class of evolutionary normal form games*, in R. Selten (ed.) *Game Equilibrium Models I*, Berlin, Springer (1991), 29-97.
46. E.C. Zeeman, *Population dynamics from game theory*, in *Global Theory of Dynamical Systems*, Springer Lecture Notes in Mathematics **819** (1980).

Faculty of Mathematics, University of Vienna, A-1090 Vienna, Austria
and International Institute for Applied Systems Analysis, A-2361 Laxenburg, Austria
E-mail address: `karl.sigmund@univie.ac.at`