

7 Evolutionary Games and Learning

- 7.1 Replicator dynamics and evolutionary stable strategies: the basic biological concepts
 - 7.2 Extensions and generalizations to economics: evolution, rationality and efficiency
 - 7.3 Learning models
 - 7.4 Applications
-

Evolutionary games have been developed in biology and mathematics, notably by J. Maynard Smith and his collaborators and after them by many other researchers, including the Nobel prize in Economics, R. Selten. Evolutionary models have proven very helpful to study the behaviour of animals in various strategic contexts and to understand biological evolution. Now, these models are becoming popular among game theorists and the approach appears to be fruitful in the field of economics and business studies. One feels that the kind of adjustment dynamics studied by evolutionary games may be useful for discussing various types of economic issues where agents react slowly through emulation, imitation, or learning. Of course the kind of modelling suitable for applications in economics must be different from that used in biology. Even if managers, sellers and buyers, or other economic decision makers sometimes have very limited rationality, they can be expected to behave more rationally than birds or rats. It is then the introduction of *specific learning processes* that can bridge the gap between biological games and economic behaviour.

There is a particular reason why economists could now be interested in evolutionary models. Non-cooperative GT, as applied in economics, is facing two difficulties: first, in many economic problems, it is not entirely clear how a NE can be finally reached by the players and, second, when there are many equilibria, with different implications, it is important to understand how a particular equilibrium will eventually be selected. It happens that the dynamic adjustments described by evolutionary models may give interesting answers to both these questions. So, even though economic applications are still rare and some progress is still to be made in order to adapt the modelling, the path seems a very promising one to follow.

The present chapter will mainly develop the framework put forward to study biological games. Some applications to economic problems in the field of industrial organisation or international trade theory will be presented. Section 7.1 introduces to the basic concepts of Evolutionary Game Theory (EGT): the 'Replicator Dynamics' (RD) and 'Evolutionary stable strategies' (ESS) in both symmetrical and asymmetrical evolutionary games. Section 7.2 presents extensions of the basic concepts and discusses their relevance for economics. The relationship between RD, ESS and other equilibrium concepts is also presented in this section. In section 7.3, we present a brief introduction to learning models, both an older approach in the tradition of the Cournot adjustment model, and new developments which are designed to make EGT fit better traditional economic problems. Finally, section 7.4 will be devoted to a few examples of applications of the concepts of evolutionary games to economic issues.

7.1 Replicator Dynamics and evolutionary stable strategies: the basic biological concepts

The most interesting cases for the economist involve different populations of players. However, in order to keep things simple at first, let us introduce the RD and the concept of ESS by considering the case often studied in biology of a population of a single species, also called a homogeneous population, playing a symmetric game. Then, we will see how these concepts are used when there are more than one population.

In the single-population case, one studies the random matching of individuals who have the same set of strategies available and whose payoffs are entirely symmetric. Here a 'strategy' means a special behaviour and each individual is genetically programmed for playing a particular strategy. In biological games, a payoff may be interpreted as the number of offspring. It is also called 'fitness'.

Let us call $X = \{x_1, \dots, x_i, \dots, x_n\}$ the set of pure strategies available to the players, and $U(x, x')$ the payoff when one agent plays x and her (or his) opponent plays x' . In this chapter we will denote explicitly by m the mixed strategies.

7.1.1 The Replicator Dynamics

At one point in time there may be different fractions of the total population programmed to play a particular strategy. In order to study the evolution of these groups, admitting that only the fittest will survive, a type of dynamic adjustment must be imagined. It is common in biological games to use a kind of Malthusian dynamics, also called the RD.

According to the RD, the fraction of the population playing a particular strategy will increase if it performs better in terms of the fitness function than the population average. If a type of players gets less than the average payoff, then its percentage in the population will decrease.

There is a simple large population of agents playing the same symmetric game. In order to keep the presentation as simple as possible, the number of pure strategies is limited to two: x and x' . An extension to a larger number of strategies is straightforward.

Let n_t and n'_t be the number of agents, respectively, playing x and x' at time t and N_t the total population.

Let $s_t(x)$ denote the proportion of agents playing strategy x at time t :

$$s_t(x) = \frac{n_t}{n_t + n'_t} = \frac{n_t}{N_t} \quad (7.1)$$

Agents programmed for playing x have the expected payoff:

$$u_t(x) = s_t(x)u_t(x, x) + s_t(x')u_t(x, x') \quad (7.2)$$

The average payoff in the population is then:

$$\bar{u}_t = s_t(x)u_t(x) + s_t(x')u_t(x') \quad (7.3)$$

Starting from these assumptions, one can define several versions of the RD. The most common version in a continuous time is expressed in the following differential equation:

$$\dot{s}(x) = s(x)[u(x) - \bar{u}] = F(s) \quad (7.4)$$

This equation of replication describes the evolution process of populations programmed for playing the diverse strategies: in this case x and x' . It reflects the basic idea defining the RD: if strategy x is performing better than the average, the agents who play it will see their proportion increase in the total population.

There are different ways to obtain (7.4). We present two ways below, but if the logic of the RD is understood, the reader may prefer to go directly to the numerical examples which follow.

A non-overlapping generations model of the Replicator Dynamics (Van Damme, 1991)

In each period, agents are paired at random to play a symmetric game. Their payoffs correspond to their offspring who replace them in the next period. The number of players choosing each strategy depends on the payoffs in the previous play of the game.

If n_t players play x at t , then $n_t u_t$ players will play x at $t+1$. The expected number of players in period $t+1$ is:

$$n_t u_t(x) + n'_t u_t(x') = N_{t+1}$$

It is equal to $N_t \bar{u}_t$. The proportion of players choosing x in period $t+1$ is

$$s_{t+1}(x) = \frac{n_t u_t(x)}{(n_t + n'_t) \bar{u}_t} = s_t(x) \frac{u_t(x)}{\bar{u}_t} \quad (7.5)$$

In this discrete time model, the evolution of the population can be represented as:

$$s_{t+1} - s_t = s_t(x) \frac{u_t(x) - \bar{u}_t}{\bar{u}_t} \quad (7.6)$$

Considering very short time periods allows us to write (7.6) as:

$$\dot{s} = s(x) \frac{u(x) - \bar{u}}{\bar{u}} \quad (7.7)$$

Finally, a rescaling of time leads to the following equation:

$$\dot{s} = s(x)[u(x) - \bar{u}] \quad (7.8)$$

which has the same solution trajectories as (7.7) and which is (7.4).

An overlapping generations model of the Replicator Dynamics (Binmore, 1992; Samuelson, 1997)

In the previous model, all agents were reproducing themselves at the same time and none of them could survive after reproduction. This assumption may apply to some kinds of animal species, but certainly not all. Moreover, it does not fit well the applications to economics, where we would like to have agents learning over time.

Now we assume that in each period of time of length τ , a fraction τ of the population reproduces itself. Payoffs of the game are again taken as representing offspring and each agent playing strategy x will give birth to $u_t(x)$ offspring at time t .

In period $t + \tau$, the number of agents playing strategy x is given by:

$$n_{t+\tau} = n_t + \tau n_t u_t(x)$$

The total number of agents next period is:

$$N_{t+\tau} = n_t(1 + \tau u_t(x)) + n'_t(1 + \tau u_t(x'))$$

The proportion of agents playing x next period is:

$$s_{t+\tau}(x) = \frac{n_{t+\tau}}{N_{t+\tau}} = \frac{n_t(1 + \tau u_t(x))}{n_t(1 + \tau u_t(x)) + n'_t(1 + \tau u_t(x'))}$$

or:

$$s_{t+\tau}(x) = \frac{s_t(x)(1 + \tau u_t(x))}{s_t(x)(1 + \tau u_t(x)) + s_t(x')(1 + \tau u_t(x'))}$$

Then, the evolution of the population can be represented as:

$$s_{t+\tau}(x) - s_t(x) = s_t(x) \frac{\tau u_t(x) - \tau \bar{u}_t}{1 + \tau \bar{u}_t} \quad (7.9)$$

Taking the limit $\tau \rightarrow 0$ in (7.9) gives:

$$\dot{s} = s(x)(u(x) - \bar{u})$$

which is again (7.4).

Example 1

Consider the game of coordination described by the payoff matrix in Figure 7.1.

The players have two alternative strategies, x_1 and x_2 . Call s the proportion of players programmed for playing strategy x_1 . The player programmed for 1 will get:

$$u_1 = s.3 + (1 - s)2 = s + 2$$

Similarly:

$$u_2 = s.0 + (1 - s)4 = (1 - s).4$$

The average payoff is:

$$\bar{u} = s(s + 2) + (1 - s)(1 - s).4 = 5s^2 - 6s + 4$$

Figure 7.1 A coordination game

	x_1	x_2
x_1	3, 3	2, 0
x_2	0, 2	4, 4

Then the equation of replication $F(s)$ is:

$$F(s) = s[s + 2 - (5s^2 - 6s + 4)]$$

and finally:

$$F(s) = s(1 - s)(5s - 2)$$

The equation of replication can be represented graphically by the phase diagram in Figure 7.2.

It appears that, in such a game, if the proportion of players programmed for playing strategy x is initially greater than $\frac{2}{5}$, the RD will increase it continuously, up to the point where everybody is playing it. Inversely, this percentage will fall to zero if the game starts at a level of s lower than $\frac{2}{5}$.


Notice that the RD has three steady states, that is points for which $F(s) = 0$: for $s = 0$, $s = \frac{2}{5}$ and $s = 1$. Are they all equilibrium points of the game? Clearly, one would not like to consider the point $s = \frac{2}{5}$ as an equilibrium, since a slight departure from it implies a continuous move towards $s = 0$, or $s = 1$.

This example shows how stability is important in the evolutionary approach. Let us now leave this particular example and get a more general view about the equilibrium definition.

Evolutionary equilibrium

In the evolutionary approach, equilibrium means not only a state of rest of the dynamical process (a steady state, or a fixed point of the function describing the dynamics), but also a certain form of stability of this rest point. More precisely, one can propose the following definition.

Figure 7.2 Phase diagram of the RD in the coordination game



Definition 1 (Evolutionary equilibrium)

An evolutionary equilibrium (EE) is any asymptotically stable fixed point of the dynamical process of evolution (see Appendix 1 of this chapter, p. 379, for the definition of asymptotic stability).■

In the above example, a simple way to study stability is to consider the sign of the slope of function $F(s)$. When $dF(s)/ds < 0$, the steady state is stable. At points $s = 0$ and $s = 1$, the slope of the function $F(s)$ is negative; at $s = \frac{2}{5}$, the slope is positive. These signs confirm that the former points are stable steady states, while the latter is not. In other words, $s = 0$ and $s = 1$ are EE of this game.

In Example 1 above the EE corresponds to a 'monomorphic' population, that is a population where everybody will play the same strategy (before any mutation brought by possible 'mutants'). We want to provide another example showing that the EE may also correspond to a polymorphic population.

Example 2

A good example is provided by the Hawk-Dove game. The members of a very large population are fighting for the use of a territory, or any other particular valuable scarce resource. V is the value of that resource for any member of this population. Individuals meet at random and can play one or the other of the two following strategies: behave like a Hawk (aggressively), or behave like a Dove (nicely). The Hawk is always ready to fight, the Dove will always avoid fighting. Every fight has a cost, noted C .

When a Hawk meets a Dove, the latter refuses the fight and leaves the place. Of course, the Hawk gets all the benefit. When two Doves meet, the payoff is shared pacifically, $V/2$ for each. Finally, when two Hawks meet each other, they fight until they get half of the value of the resource minus the cost of fighting, that is $(V - C)/2$. Figure 7.3 summarizes the game.

If one assumes a Malthusian dynamics, one gets the following equation of replication:

$$F(s) = s(1-s)[s(V-C)/2 + (1-s)V/2]$$

Figure 7.3 The Hawk-Dove game

	Hawk	Dove
Hawk	$(V-C)/2, (V-C)/2$	$V, 0$
Dove	$0, V$	$V/2, V/2$

where s is the percentage of the population playing Hawk. To be more precise, give numerical values to the parameters V and C ; for instance $V = 4$ and $C = 16$. Then:

$$F(s) = s(1-s)(2-8s)$$

This equation has three roots: $s = 0$, $s = \frac{1}{4}$ and $s = 1$. The phase diagram is represented in Figure 7.4.


A simple look at the arrows describing the dynamics and the fact that $F(s)$ at point $s = \frac{1}{4}$ is negative shows that this point is an EE. It is easy to check that this is the only EE in this game. So here the population is 'polymorphic' at equilibrium. There is a proportion of $\frac{1}{4}$ playing Hawk and $\frac{3}{4}$ playing Dove. The equilibrium can be interpreted either as a pure strategy equilibrium or as a mixed strategy equilibrium. In the latter case, we have to assume that individuals can be programmed for playing pure or mixed strategies.

The mixed strategy equilibrium is the only EE of this game. It is also the only symmetric NE (while there are also two asymmetric NE: Hawk, Dove and Dove, Hawk). We will see in the next sub-section that it has another stability property since it is an 'evolutionary stable strategy' (ESS) (see definition below).

Remark 1

Of course, one may use many other ways to represent the evolution dynamics. One can, for instance, consider a *discrete* time rather than a *continuous* time dynamics. Unfortunately, this change can have dramatic effects on the results (see sub-section 7.2.2, the result obtained by Dekel and Scotchmer, 1992, showing that in a discrete time model the RD may not eliminate the strictly dominated strategies).■

Figure 7.4 Phase diagram of the RD in the Hawk–Dove game



Remark 2

The s presented in (7.4) links the rate of growth of a strategy to its performance with respect to the average. But one can imagine other relationships. For example, the *instantaneous rate of change* of a strategy, instead of its rate of growth, may be related to the difference between the payoff from playing this strategy and the average payoff. Again, the assumption may make a difference (see, for instance, in Friedman, 1991). Other kinds of dynamics can still be envisaged for describing the evolution process. However, it is reasonable to require that they all verify a condition of compatibility with the fitness function (Friedman, 1991): the fitter strategies must grow compared to those which fit less well. ■

7.1.2 Evolutionary stable strategies

A particular concept of equilibrium, called 'Evolutionary stable strategy' (ESS) has been proposed by Maynard Smith and Price (1973) in order to describe the stable state of the evolutionary process. The basic idea of ESS is to require that the equilibrium can 'resist' mutant invasion. While the RD is studying the dynamical property of a given strategy, the ESS concept focuses on a different issue: the study of possible mutations to unused strategies.

Suppose that the population is originally playing a strategy x , which can be a pure or a mixed strategy, and that a small percentage of 'mutants', say ε , play another strategy x' .

Definition 2 (Evolutionary stable strategy)

An ESS means that the incumbent population gets a higher payoff than the invaders when pairs of players are randomly chosen. Each player has a $(1 - \varepsilon)$ chance to meet someone playing strategy x and a probability ε to meet an invader. So the condition defining the ESS can be written:

$$u[x, (1 - \varepsilon)x + \varepsilon x'] > u[x', (1 - \varepsilon)x + \varepsilon x'] \quad (7.10)$$

where ε is positive and sufficiently small ($0 < \varepsilon < \bar{\varepsilon}$). ■

The above definition is sufficient to define an ESS, but it is common and useful to present an equivalent two-part condition. Notice first, using the linearity of expected utility, that:

$$u[x, (1 - \varepsilon)x + \varepsilon x'] = (1 - \varepsilon)u(x, x) + \varepsilon u(x', x')$$

Then (7.10) can be rewritten:

$$(1 - \varepsilon)u(x, x) + \varepsilon u(x, x') > (1 - \varepsilon)u(x', x) + \varepsilon u(x', x') \quad (7.10')$$

This inequality has to be verified only for values of ε close to 0. Hence for any $x' \neq x$:

$$(i) \quad u(x, x) > u(x', x)$$

or

$$(ii) \quad \text{if } u(x, x) = u(x', x), \text{ then: } u(x, x') > u(x', x') \quad (7.11)$$

In condition (7.11), part (i) makes it clear that an ESS must be a NE. It means that when all players play x , it is unprofitable for any player to deviate and play x' . Part (ii) shows that ESS is a kind of refinement of symmetric NE. Even if strategy x cannot do better than x' against players programmed for x (that is, in a weak NE case), it can still win against strategy x' when opponents play x' .

Example 1

A Prisoner's dilemma game is given in Figure 7.5. If the game is played only once, the DSE (D, D) also corresponds to an ESS. Clearly, with x for the strategy Defect (or 'Aggressive'), and x' for Cooperate (or 'Pacific'), and applying (7.10'), one gets:

$$3(1 - \varepsilon) + 6\varepsilon > 2(1 - \varepsilon) + 5\varepsilon$$

Now what would happen if the game were repeated many times without discounting, and if mutants were playing the so-called 'Tit-for-tat' (TFT) strategy'. It has been shown (see Axelrod and Hamilton, 1981) that the strategy (D, D) cannot resist the invaders.

A player programmed for TFT will eventually lose once against a player programmed for Defection, but thereafter will always also play Defection. On average (recalling that there is no discounting) over a large number of runs, she (or he) will get as much as the players who choose Defection at the first run, specifically a payoff of 3. So:

$$u(\text{TFT}, D) = u(D, D).$$

However, when they meet each other the mutants can perform better than when a player programmed for Defection meets them: respectively, on

Figure 7.5 A Prisoner's dilemma

	Cooperate (C)	Defect (D)
Cooperate (C)	5, 5	2, 6
Defect (D)	6, 2	3, 3

average 5 and 3 (that is because the mutants maintain a strategy of cooperation when the opponent does the same and switch to Defection when the opponent chooses this strategy). So, we also have:

$$u(D, TFT) < u(TFT, TFT).$$

Strategy D does not verify condition (7.11)(ii), and is not in this case an ESS.

Example 2

The payoffs of a coordination game are given in Figure 7.6.

This game has three NE: (x_1, x_1) , (x_2, x_2) , and the mixed strategy equilibrium $(m, m) = (\frac{1}{4}, \frac{3}{4})$. The first two are ESS, but the latter is not.

Let us check for the strategy (x_2, x_2) , which provides a payoff of 1 for each player. The expected payoff of an individual playing strategy 2, when she (or he) has a $(1 - \varepsilon)$ chance of meeting someone playing the same strategy and a ε probability of meeting a mutant is:

$$(1 - \varepsilon)1 + \varepsilon.0 = (1 - \varepsilon)$$

The expected payoff of someone playing strategy 1 and facing the same distribution of opponents is:

$$(1 - \varepsilon)0 + \varepsilon.3 = 3\varepsilon$$

Obviously, with a small value for ε , the former number is higher than the latter and condition (7.10') is met. The strategy (x_2, x_2) is an ESS. It is interesting to note here that the *inefficient* NE (x_2, x_2) may still appear as an equilibrium outcome of an evolutionary process.

By the same reasoning one can check that the strategy (x_1, x_1) is also an ESS. However the mixed strategy (m, m) , although being a NE, is not an ESS. The strategy (m, m) can be invaded by a pure strategy x_1 or x_2 .

Let us consider a population playing (m, m) , and mutants playing x_2 . We have:

$$u(m, m) = u(x_2, m) = \frac{3}{4}$$

Figure 7.6 A coordination game

	x_1	x_2
x_1	3, 3	0, 0
x_2	0, 0	1, 1

So we are typically in the case where condition (ii) of (7.11) must be considered. But precisely here the condition is not met:

$$u(m, x_2) = \frac{3}{4}, \text{ but: } u(x_2, x_2) = 1$$

so:

$$u(m, x_2) < u(x_2, x_2)$$

The mutants are doing as well as the incumbent population when they meet an incumbent but they perform better when they meet each other. Of course, mutants playing strategy x_1 will also invade a population playing (m, m) .

Remark 1

In some games there is no ESS. As Haigh (1975) has shown, the number of ESS is always finite, possibly zero. ■

7.1.3 Neutral stability, evolutionary stable sets and robustness against equilibrium entrants

Neutral stability or weak ESS

In the formal definition of an ESS in 7.1.2, we used a strict inequality, but what happens if the mutant strategy can do as well against itself as the incumbent strategy does? In the latter case, mutants are called 'neutral'. The concept of a 'Neutrally stable strategy' (NSS) then is defined as an ESS with *weak* inequality.

Definition 3 (Neutrally stable strategy)

$x \in X$ is a NSS if for every strategy $x' \neq x$ there exists some $\bar{\varepsilon} \in (0, 1)$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$ the following inequality is satisfied:

$$u[x, \varepsilon x' + (1 - \varepsilon)x] \geq u[x', \varepsilon x' + (1 - \varepsilon)x] \blacksquare$$

This inequality can be rewritten as:

- (i) $u(x, x) > u(x', x)$
- (ii) if $u(x, x) = u(x', x)$ then: $u(x', x) \geq u(x', x')$

Obviously neutral stability is less demanding than evolutionary stability, since the condition guarantees here only that the mutants cannot do better than the incumbent population. Neutral stability is still a refinement of symmetric NE:

$$\text{ESS} \subset \text{NSS} \subset \text{NE}$$

Now, since neutral mutants can get as much as the incumbents, they may stay in the game. Hence, each neutral mutation will increase the share of neutral mutants in the total population. Such a progressive 'invasion' of the population by neutral mutants is called an 'evolutionary drift'. The problem is that at some point the evolutionary drift may destabilize the NSS. A new concept of stability defined on sets of NE (Thomas, 1985) is then useful in order to deal with this risk of destabilizing evolutionary drift.

Evolutionary stable sets

A set of symmetric NE strategies is called 'Evolutionary stable' (ES) if each strategy in the set can do at least as well against any close strategy x' as it can do against itself and if the mutant strategy also belongs to the set in the case of equal payoffs.

Definition 4 (Evolutionary stable set) (Kandori, 1997)

A closed set $X^* \subset X$ is ES in a symmetric two-player game if:

- (i) each element of X^* is a NSS and
- (ii) $x \in X^*$, $u(x, x) = u(y, x)$ and $u(x, y) = u(y, y)$ implies $y \in X^*$ ■

In other words, the evolutionary drift can never lead to an unstable point. At the border, the mutants would do strictly worse than the incumbent population.

Example

In the repeated Prisoner's dilemma, the TFT strategy is an NSS. However one can check that it does not belong to an ES set, since for TFT and C (cooperate) we have:

$$U(\text{TFT}, \text{TFT}) = U(\text{C}, \text{TFT}) \text{ and } U(\text{TFT}, \text{C}) = U(\text{C}, \text{C})$$

There can be an evolutionary drift from TFT to C. Yet, C is not a NE, and hence not a NSS.

This example brings us to the interesting extension suggested by Swinkels (1992a). In economic applications one may want to consider that the agents, including the mutants, are smart enough to avoid trying stupid strategies against incumbent strategies. In the latter example strategy C is indeed very fragile compared to TFT. Even with weak rationality, an economic agent would hesitate before using C against TFT.

Equilibrium entrants and robustness to equilibrium entrants

Swinkels (1992a) argues that conditions on ESS might be too stringent in economic contexts. When players are firms it seems reasonable to restrict attention to mutant strategies that themselves fulfil a stability condition. In this line, Swinkels proposes a new stability concept, that he calls 'Robustness against equilibrium entrants' (REE) permitting us to consider only 'rational' mutations or experimentation and eliminate the 'stupid' ones. Assuming that the population share of mutants is ε , the post-mutation mixed strategy is: $\omega = \varepsilon x' + (1 - \varepsilon)x$. An entrant is called an equilibrium entrant if x' is a best reply to ω . But if x' is a best reply to ω , then a plan of invasion of mutants representing ε per cent of the population becomes self-enforcing. One may want to define situations where this risk does not exist, and this is precisely the objective of the REE.

Definition 5 (Robustness against equilibrium entrants)

A symmetric strategy profile (x, x) is REE if there exists some $\bar{\varepsilon} \in (0, 1)$ such that if $x' \neq x$ and $\varepsilon \in (0, \bar{\varepsilon})$ then x' is not a best reply to $\varepsilon x' + (1 - \varepsilon)x$:

$$x' \notin BR[\varepsilon x' + (1 - \varepsilon)x] \quad (\text{where } BR \text{ is the set of best responses}). \blacksquare$$

It can be shown that if a strategy is REE, it is also a best reply to itself, that is, it must be a NE. One can even go further (Swinkel, 1992a) and show that REE implies properness, thus being a particular refinement of NE (see sub-section 4.2.1 of chapter 4 for a definition of a proper equilibrium). Since REE is less stringent than ESS, we have the following inclusions of sets:

$$\text{ESS} \subset \text{REE} \subset \text{NE}$$

The motivations behind the Swinkels' stability concept are twofold. On the one hand, when all mutations are considered then some games can fail to have an ESS (see Remark 1 in sub-section 7.1.2). On the other hand, if some mutations are unpalatable then the relationship between ESS and Nash refinements become less compelling if it does not persist in the presence of rational mutants. Note that although some games without ESS present strategies that are REE, the Swinkels' notion also fails to exist in some games.

An extension in terms of set-valued concepts is suggested by Swinkels (1992a) under the name: Equilibrium evolutionary stable sets (EES).

Definition 6 (Equilibrium evolutionary stable set)

A set $X^* \subset X$ is EES if it is a minimal closed non-empty set such that:

- (i) X is a sub-set of the set of NE
- (ii) for some $\bar{\varepsilon} \in (0, 1)$, if $\varepsilon \in (0, \bar{\varepsilon})$, $x \in X^*$, $x' \in X$ and $x' \in BR(\varepsilon x' + (1 - \varepsilon)x)$, then:

$$\varepsilon x' + (1 - \varepsilon)x \in X^*.$$

The definition states that an EES set is a minimal closed set of symmetric NE such that the population can never be led out of X by a series of small equilibrium entries.

7.1.4 Asymmetrical evolutionary games

The previous sections dealt with single-population games. But, even if the payoff matrix is symmetric, an evolutionary game may easily show interesting asymmetries between the players. Maynard Smith (1982) raised the possibility that the animals could condition their behaviour on whether they are the 'row' or the 'column' player, that is to say that the strategies are then conditioned by the role played in the game (summarized by the phrases 'row and column'). Of course asymmetries can be worse still, payoffs being different between the row and the column players. In economics or business, the situations involve different kinds of players, in the sense that their available strategies are different and that they do not get identical payoffs from their participation in the game. If we want to model games with sellers and buyers, or incumbent firms and potential entrants, it is necessary to extend the evolutionary approach to multi-population settings. For reasons of simplicity, we will consider only the case of two different populations, but extensions to K populations ($K > 2$) is fairly possible.

Suppose that two large populations interact, the members of the first one being chosen randomly to meet also randomly chosen members of the second population. In order to avoid the difficulty coming from different speeds of adjustment, let us assume that the two populations are of equal size. The natural extension of the evolutionary model is then to consider a RD for each population. We have thus two replication equations:

$$\dot{s}_i^j(x) = s_i^j(x^j)[u_i^j(x^j) - \bar{u}_i^j] \quad i = 1, 2 \quad (7.12)$$

One can study the stationary points of the dynamics and verify the stability condition. A major finding in the multi-population case is that the mixed strategy profiles cannot be asymptotically stable.

One can also apply the ESS concept, although adapted to the multi-population context (see Weibull, 1995, section 5.1), and check whether the strategies are ESS or not. Since the members of a given population are now supposed to meet randomly members of a different population, the definition of an ESS given in (7.2) cannot any longer apply. A new definition of an ESS, allowing for asymmetric mutants, can be defined, as in Swinkels (1992b).

Another way to apply the concept of ESS is to symmetrize the game (see exercises 7.7 and 7.8). 'Symmetrizing' means that the player's role is chosen by Nature before the game starts. *Ex ante*, players do not know which role they will play and must then calculate their expected payoff under the assumption that they have an equal chance to be a row or a column player. An important result proven by Selten (1980) is that an ESS in such a symmetrized game is a *strict* equilibrium. Since a mixed strategy equilibrium by definition is not a strict equilibrium, mixed equilibria cannot be ESS of asymmetric games. On the other hand, one must notice that the 'symmetrization' of the game may work well for some games but may not fill many asymmetric games, the specific roles not being interchangeable (for a more complete discussion on this topic see Binmore and Samuelson, 2001).

Example I

A game of entry may oppose incumbent firms (player 1) and potential entrants (player 2). In order to maintain the assumption of large populations of players, one must imagine here a large number of incumbent firms, being in a position of local monopolies and many potential entrants, choosing randomly the location of their attack. The numerous possibilities of matching and the fact that two particular firms may meet only occasionally means that any reputational effect is negligible. What remains is that each entrant knows she (or he) is facing either an aggressive incumbent or a passive (accommodating) one. Similarly, the incumbent has only a probability on the chances of entry. Let p be the proportion of aggressive incumbents and q the proportion of potential entrants who actually enter.

The payoffs are given in Figure 7.7.

What is the RD for each category of player? For the local monopolies,

Figure 7.7 An entry game

	Enter	Do not enter
Aggressive	-1, -1	8, 0
Passive	3, 3	8, 0

We get:

$$\dot{p} = p[u_1^1 - \bar{u}^1]$$

where: $u_1^1 = q(-1) + (1-q)8 = 8 - 9q$

and since: $u_2^1 = 3q + (1-q)8 = 8 - 5q$

then: $\bar{u}^1 = p(8 - 9q) + (1-p)(8 - 5q) = 8 - 5q - 4pq$

and: $\dot{p} = p[(8 - 9q) - (8 - 5q - 4pq)] = p(1-p)(-4q)$

Therefore if q is positive, $\dot{p}$ is negative; which means that the percentage of incumbents playing aggressive is continuously decreasing.

Concerning the potential entrants, one gets by the same reasoning:

$$\dot{q} = q[u_1^2 - \bar{u}^2]$$


where $u_1^2 = 3 - 4p$ and $\bar{u}^2 = 3q - 4pq$

Then: $\dot{q} = q(1-q)(3 - 4p)$ and $\dot{q}$ is positive if $p < \frac{3}{4}$. In other words, the number of potential entrants who actually enter increases continuously if the proportion of aggressive local monopolies is lower than $\frac{3}{4}$.

The phase diagram in Figure 7.8 represents the dynamics of this example.

It appears clearly on the graph that the only stable point of the RD is the point corresponding to $q = 1$ and $p = 0$. The incumbent firms accommodate and the potential entrants do enter. Notice that the dynamical process has other stationary points: $(p = 0, q = 0)$, $(p = 1, q = 1)$, $(p = 1, q = 0)$ and $(p = \frac{3}{4}, q = 0)$. Among these four steady states, three are NE: two equilibria in pure strategies points: $(p = 0, q = 1)$, $(p = 1, q = 0)$, and the mixed strategy equilibrium $(p = \frac{3}{4}, q = 0)$. Yet only the equilibrium: $(p = 0, q = 1)$ is perfect. It is interesting to notice that it is also the EE of the game (the only asymptotically stable point of the dynamics). We will say more on the relationship between evolutionary equilibria and non-cooperative game equilibria in the next section.

Figure 7.8 Phase diagram of the RD in the entry game



We must emphasize that the couple of strategies ($p = 0, q = 1$) also has the property of being an ESS (see exercise 7.7). In the end, the only SPE of this game is also an EE and an ESS. This correspondence will be investigated further later.

It is intriguing to see how players who do not choose rationally (they are just programmed) end up playing a SPE. A criterion like perfection seems to require a very strong intelligence of the game and a high capacity for computation, to solve for instance the backward induction type of calculation. Without any capacity for computation the players can find the same solution through an evolution process. The only drawback of the evolutionary approach is that the process takes time. In economic or business applications, the external conditions of the game may change rapidly. The evolutionary approach can be applied only for problems not characterized by rapidly changing environments (see Camerer, 1991, for a pessimistic view on this matter).

Example 2

We saw above that in the Hawk-Dove game the only EE was the mixed strategy: $(\frac{1}{4}, \frac{3}{4})$. Yet, what would happen if the game were played by two distinct populations? The mixed strategy is no longer an EE in this case.

Distinguishing two populations in this game means that there are two (large) groups of individuals who do not give the same value to the scarce resource: $V_1 \neq V_2$. Let us say that the scarce resource is a territory, and players 1 are residents, while players 2 are foreigners. So, normally, $V_1 > V_2$. For instance $V_1 = 6, V_2 = 2$, and as in the previous case, the cost of a fight is $C = 16$. Then the payoff matrix is as in Figure 7.9.

Now p is the proportion of residents playing Hawk and q is the proportion of foreigners playing Hawk. The Replication equations are:

$$\dot{p} = p(u_R - \bar{u}_R)$$


and

$$\dot{q} = q(u_I - \bar{u}_I)$$

where R is used for residents and I for invaders. After some calculation, we get:

Figure 7.9 An asymmetric Hawk-Dove game

	Hawk	Dove
Hawk	-5, -7	6, 0
Dove	0, 2	3, 1

Figure 7.10 Phase diagram of the RD in an asymmetric Hawk–Dove game

$$\dot{p} = p(1-p)(3-8q)$$

$$\dot{p} > 0 (< 0) \text{ if } q < \frac{3}{8} (> \frac{3}{8})$$

$$\dot{q} = q(1-q)(1-8p)$$

$$\dot{q} > 0 (< 0) \text{ if } p < \frac{1}{8} (> \frac{1}{8})$$

The dynamics can be represented by the phase diagram in Figure 7.10.

The RD shows two stable points ($p = 0, q = 1$) and ($p = 1, q = 0$). These are the EE of this game. All the players in a given population end up choosing to be Hawks while in the other population people choose to be Doves, depending on the starting point. No combination of Hawks and Doves can be maintained in the same population as an equilibrium of the evolutionary process. One can check that the two EE are also ESS (see exercise 7.8).

7.2 Extensions and generalizations to economics: evolution, rationality and efficiency

7.2.1 Connections between Replicator Dynamics, Evolutionary stable strategies and other equilibrium concepts

In the previous sections we have already underlined some connections between the evolutionary equilibrium concepts and the standard equilibrium concepts of non-cooperative game theory. When we defined an ESS, we stressed that it was a kind of refinement of symmetric NE. We also noticed that the ESS were stable stationary points of the dynamical