Rend. Sem. Mat. Univ. Pol. Torino - Vol. 66, 3 (2008)
Lezioni Lagrangiane 2007–2008

M. Tomassini

GAMES, EVOLUTION, AND SOCIETY

Abstract. Classical game theory has been conceived to model strategic situations in which individuals' decision making depends on the actions chosen by others. The theory has been extremely successful in several economic, social and technological contexts. However, many empirical results tend to indicate that the explanatory power of the theory is not completely satisfactory. Here we describe how the evolutionary interpretation of the theory is able to address some of the theory's deficiencies. In particular, we describe extensions that try to take into account the agents' population structure, based on empirical data about social networks. We find that, at the expense of sacrificing some mathematical rigor, these models give results that are closer to the observed behavior of finite populations of interacting socio-economic agents.

1. Introduction

Many interactions between people or other entities in society such as firms and governments take the form of mutually outcome-dependent decision making. Multi-party decision making often leads to conflicting situations that have profound consequences for the participating entities and for society as a whole. For example, whether to pay taxes or not, how to best exploit a common good, or how to bargain in a market-based economy are all familiar expressions of this kind of socio-economic phenomena. It is thus not surprising that researchers have been searching for models and theories that can describe these situations and, possibly, dictate solutions or at least methods for investigating and ranking the different possibilities that arise. Game theory has been the most successful attempt in this direction so far. It was founded by von Neumann and Morgenstern with the publication of their book in 1944 [27], although early results had been obtained previously by von Neumann himself and by Émile Borel. In this article, which should hopefully be self-contained to a large extent, I shall first describe the basic ideas of game theory, both the classical theory as well as a more recent extension called *evolutionary game theory* which is particularly suited for the description of game-theoretical interactions in large populations. After this introductory material, we shall see how the main ideas can be applied to socio-economically relevant problems, with particular emphasis on the influence of the structure of the participating groups. This will lead us to the idea of a social network and to the usage of evolutionary game theory in a setting that takes these network structures into account. Since the latter case does not yet enjoy a complete and rigorous mathematical formalization, I shall show how numerical simulation can help us to understand the complex phenomena that emerge when many entities interact in a non-linear manner. In the interest of readability, and in the spirit of the Lagrangian Lectures, I shall often sacrifice rigor in favor of a more qualitative presentation which should make the article accessible to a wide audience.

2. The basics of game theory

According to Myerson, a well known game theorist and a recent Nobel prize winner, game theory is the study of decision making under interacting, and often conflicting situations [13, 26]. Such situations are common in everyday life: when we are driving our cars to work in the morning, when we pay our taxes, when two firms bid in an auction for a market, when two countries are deciding whether they should declare war to each other or not, are all playing a game. But of course game theory cannot analyze these situations in all their complexities. For example, when you take a bus and have to decide whether to pay the ticket or take a free ride, assuming that there is no control, game theory conclusion is that on rational grounds you shouldn't buy the ticket, independent of your detailed state of mind and psychological propensities. Rationality is the key word here and is to be understood in the *homo economicus* sense. The concept goes back to von Neumann, Morgenstern and Savage [27, 21], among others, and it is the cornerstone of much economic theory. We shall only briefly recall the main ideas here, the interested reader will find complete accounts in [13] among many other references. The agents that the theory assumes are *intelligent* and *rational* in the following sense:

- an agent knows what are the possible choices in a given situation
- given a choice, an agent is able to associate a value, called *utility*, to the consequence of that choice
- rationality principle: each agent takes the decision that *maximizes* its utility
- the fact that agents are intelligent and rational is common knowledge

The numerical ordering of preferences is a binary relation that must satisfy a number of axioms (see [13]). It is represented in the theory by a real function called utility, or *payoff*, that assigns a numerical value to each possible outcome if everything is known with certainty. In case there is uncertainty, i.e. there is more than one possible consequence to a given action, a rational player maximizes its *expected utility*. The most important thing to remember is that more is better for a player and that any affine transformation leaves the preference ordering invariant. This means that the actual numerical values used don't matter much; the only thing that really matters is their order. In this article we shall use small integer values to represent utilities in games. It goes without saying that such a characterization of an agent has been criticized many times and that other models of players have been proposed where agent's rationality is "bounded" in a way or another. However, although full rationality is clearly unrealistic in practice, it has at least the merit of being exactly formalizable. This allows a complete and rigorous theory to be developed which we shall call here conventionally *standard game theory* to distinguish it from more recent advances. On the other hand, there is a vast literature on how people take their decisions in laboratory experiments in the framework of what is called *behavioral* economics. These experiments more often than not reveal a picture that doesn't match the prescriptions of the theory and this is a subject of heated debate. The interested reader may wish to consult [3, 6] for recent accounts where different interpretations and explanations are offered for the behavior of real players.

The part of game theory that will be made use of here is called *Non-Cooperative Game Theory*, to distinguish it from *Cooperative Game Theory*. The difference is simply that players are assumed to be selfish utility maximizers in the non-cooperative theory, while they can form binding coalitions in the cooperative form. However, cooperation as an individual strategy may emerge in non-cooperative games, as we shall see in the following.

Another useful distinction is between games that will be played only once, so to speak, and those in which the game will be played repeatedly by the same players. The first kind is called *one-shot*, to emphasize the fact that, even if the same players are chosen to play against each other one more time, the new interaction happens as if they had never met before, i.e. they don't have a memory of the past. In contrast, in *repeated* or *iterated* games, the same players may meet again and again and have a representation of the past encounters. This enables agents to develop much more complicated strategies. Both iterated and one-shot game situations are common in society. Iterated games are probably more socially relevant since social ties do not change so often and thus people are led to interact with each other repeatedly. However, it would be incorrect to conclude that iterated games are all that count: the increased mobility and pervasive information structure of modern society allow people to engage often in one-shot games too. Just to give an example, think of auctions taking place on e-bay. For reasons of space, we shall only deal with the easier one-shot scenario here which is sufficient to expose the main ideas simply and concretely. The reader wishing to know more about repeated games may start with ref. [2] and, at a more advanced level, with [13, 26].

2.1. Representations of a game

Before showing more formal definitions, it is useful to describe some simple games in an informal manner to get used to the meaning of symbols. Finite games contemplate a finite number of players N, and a finite set of *strategies* S. In the common case of two-person games, all the possibilities that may arise can be represented with a two-dimensional array which is called the game's *payoff matrix*. Let's consider, for instance, the following child "toy" game sometimes called "matching pennies". In this game, each of two players hides a coin in her hand and both show their side of the coin simultaneously to each other. One player wins if they are both the same, otherwise the other player wins. This gives rise to the following payoff matrix:

		C	
		heads	*tails*
R	*heads*	$1,-1$	$-1,1$
	tails	$-1,1$	$1,-1$

The two players will conventionally be called R and C, where R stands for “row player”, and C stands for “column player”[*]. Each player has two strategies at her disposal: $S = \{heads, tails\}$. A given element of the double matrix represents the utilities of each player when this particular pair of strategies is played, the first number being the utility of the row player. For example, if R plays *heads* and C plays *tails*, then R gets an utility of -1 and C obtains 1. This particular game is called a *zero-sum* game because the payoffs in each cell add up to 0. Many important games are not zero-sum, as we will see in the following. The above representation of a game is called the *normal* or *strategic* form for the two-person game; it ignores all timing issues and treats the players as if they choose their strategies simultaneously. The normal form is the one that we shall use in the rest of the article. However, it should be noted that there is another, more complete representation of a game which is called the *extensive* form. In this representation individuals play one at a time and the sequence of moves is explicitly shown as a tree. For the “matching pennies” game we would have the situation depicted in Fig. 1.

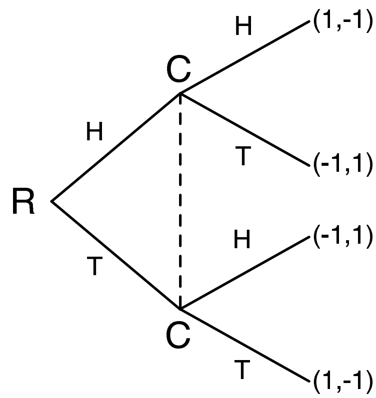

Figure 1: the *Matching Pennies* game in extended form.

This will remind many readers of the game of chess, in which one could in principle draw such a game tree, except that tree sizes grow exponentially and any pair of sensible strategies would give rise to enormous trees, which can only be partially enumerated by the fastest electronic computers. Of course, given this representation in which R moves first, one would say that C will always win by just playing the opposite strategy. But we would reach exactly the opposite conclusion if C were to move first. The point is that, although the representation implies sequentiality, according to the rules of the game, when it is time for the second player to choose her strategy, she shouldn’t know whether the game is at the first node or at the second node, i.e. whether R has played *heads* or *tails*. We formalize this fact by drawing a vertical dashed line

[*]We shall adopt this convention in the following in a tacit manner, without explicitly writing R and C.

between the two nodes of the tree and we say that the *information set* of C, when it is her turn to move, comprises both states, i.e. she is uncertain as to whether R has played *heads* or *tails*. Thus, she could also randomize her choice and choose *heads* or *tails* with probability $1/2$. This is perfectly possible but it's another story that we shall comment upon later. With this interpretation, the two forms of the game are perfectly equivalent.

Let us now go back to the general definition of a game. A finite game Γ with complete information can be formally specified in normal form as follows:

$$\Gamma = (N, S_i, u_i),\ \forall i \in N,$$

where N is the finite set of players, S_i is the finite ensemble of strategies available to player i, and $u_i : \times_{j\in N} S_j \rightarrow \mathbb{R}$ is the utility of player i. Games with incomplete information, also called *Bayesian games* will not be treated here but useful information can be found in [13].

2.2. Nash's solution

Once one has a game Γ defined as above, what does it mean to find a solution of the game? Historically, people have found several ways in which it is possible to come up with an answer. Von Neumann and Morgenstern were the first to give a solution concept for the class of games of the zero-sum type through von Neumann's famous minimax theorem [27]. Other approaches use the concept of *dominance* of a strategy with respect to another and, in some cases, may lead to a solution [13]. However, the work of J. Nash at the end of the forties and beginning of the fifties is the most elegant and commonly used solution concept, although it is not uncontroversial nowadays. Nash's solution is conceptually very simple but it needs the notion of a *mixed strategy* to be defined first.

Mixed Strategies. A *randomized* (or mixed) strategy σ_i for player i is a probability distribution $\Delta(S_i)$ over the set of "pure" strategies S_i. Since the set S_i for each player is finite, any mixed strategy σ_i for player i can be seen as a vector in $\mathbb{R}^{m_i}$, where m_i is the number of strategies available to player i. Any strategy $s_i \in S_i$ must be chosen with nonnegative probability and they must sum up to one:

$$\sum_{s_i \in S_i} \sigma(s_i) = 1.$$

Thus the vector $\sigma_i \in \mathbb{R}^{m_i}$ belongs to the unit *simplex* which, for a number of pure strategies $m_i = 2$ and $m_i = 3$, is represented in Fig. 2. In general the mixed strategy simplex of player i has dimension $m_i - 1$. It is also evident that a pure strategy s_k is a particular case of a mixed strategy in which $\sigma(s_k) = 1$ and $\sigma(s_j) = 0, \forall j \neq k$. In other words, the pure strategies are represented by the corner points of the simplex, i.e. all the points of the type $\mathbf{e}_1 = (1,0,\ldots,0)$, $\ldots, \mathbf{e}_{m_i} = (0,0,\ldots,1)$ where the $\mathbf{e}_i$ are unit vectors in $\mathbb{R}^{m_i}$.

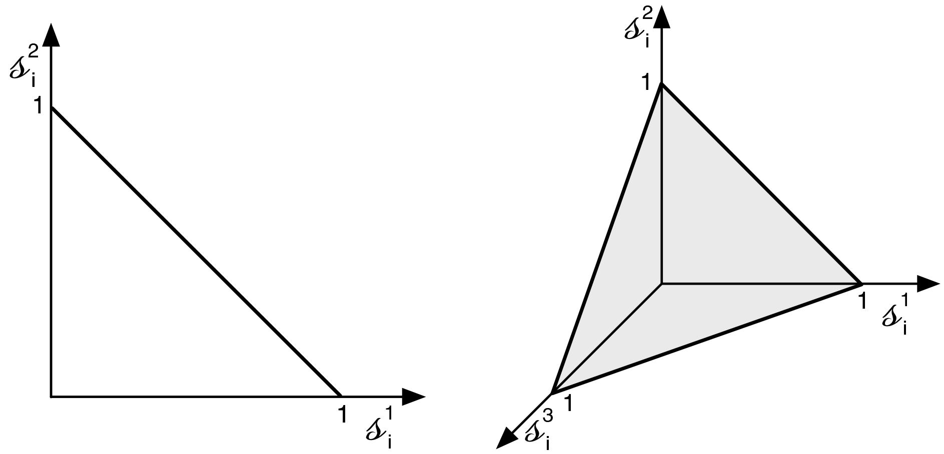

(a) (b)

Figure 2: representation of the unit simplex for two (a) and three strategies (b).

Let us now define the concept of a *randomized strategy profile*. A randomized strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$, $\sigma \in \times_{i \in N} \Delta(S_i)$, is a vector where each component $\sigma_i \in \mathbb{R}^{m_i}$ is a mixed strategy for player $i \in N$ where n is the cardinality of the set N. Let us call $u_i(\sigma)$ the expected payoff that player i would get when the players choose their strategies independently according to the strategy profile σ, and let us denote by (σ_{-i}, τ_i) a randomized strategy profile equal to σ except for the i-th component τ_i, i.e. in which player i deviates from σ by choosing another strategy profile τ_i.

Nash Equilibrium. A *Nash Equilibrium* (NE) is a randomized strategy profile σ^* such that:

$$u_i(\sigma^*) \geq u_i(\sigma_{-i}, \tau_i), \ \ \forall i \in N, \ \ \forall \tau_i \in \Delta(S_i)$$

In other words, a randomized strategy profile σ^* is a Nash equilibrium if and only if no player could increase her expected payoff by unilaterally deviating from this strategy profile. J. Nash was able to show that such an equilibrium always exists for a finite game Γ:

THEOREM 1. *Every finite game Γ in strategic form has at least one equilibrium in mixed strategies.*

Here I report this existence result without proof. For a formal proof of Nash's theorem, which makes use of Kakutani's fixed point theorem, the reader may wish to consult, for example, ref. [13].

Some Simple Games

In order to make these concepts more concrete, now it's time to illustrate them on a few simple two-person games that will be used throughout the rest of the article. These games will all have the general normal form shown in the following Table 1, which is symmetric, and in which each player has the same two possible strategies, conventionally called s1 and s2.

	s1	s2
s1	u_1, u_1	u_2, u_3
s2	u_3, u_2	u_4, u_4

Table 1: a generic payoff bi-matrix for two-person, two-strategies, symmetric games.

Prisoner's Dilemma. Let's start with the celebrated Prisoner's Dilemma (hereafter PD) on which there exist a vast literature (for an entertaining introduction, see [2]). This game has been designed to model many common situations in society which lead to a tension between personal profit and social welfare. It stipulates that, in situations where individuals may either cooperate or defect, they will rationally choose the latter. However, cooperation would be the preferred outcome when global welfare is considered. We shall write C for s1 (C stands for "Cooperate") and D for s2 (D means "Defect"). From the point of view of Table 1 a game is a PD if the payoffs are ordered as follows: $u_3 > u_1 > u_4 > u_2$. To make things more concrete, let us consider the following payoff table 2 in which the actual numerical values used are consistent with the above ordering:

	C	D
C	3,3	0,**4**
D	**4**,0	**1**,**1**

Table 2: the Prisoner's Dilemma game.

The game is usually introduced as an illustration in which there are two persons that are caught by the police after committing a crime and interrogated separately. However, it is just as easy, and perhaps more useful, to consider other naturally occurring examples which essentially take the form of a PD. For example, consider the *public goods* and *free riding* problem. In economics public goods are those that do not belong to a particular entity but are rather the property of all the society's members since nobody can be excluded from enjoying them. For example, the army or the police are public goods; street lighting or public radio and television broadcasts are also public goods. The reader can surely spot several others. Suppose now that, to preserve the air quality, which can be considered a public good, people are encouraged to install a filter in their cars which can reduce emissions significantly. The car owner bears the cost of the conversion but, if a large number of people do the conversion, the benefits in terms of air's quality are considered to be well worth the investment from the single citizen's point of view, as well as for society as a whole, of course. Suppose now that the conversion is advantageous if more than, say, P people do it; otherwise the cost overcomes the benefit for the single contributor. This situation could be represented by the above PD payoff table if we identify the row player with a single contributor which can either convert (C) or not (D). The column player represents the rest of the possible adopters. Column strategy C means that "more than P others install the filter", while D stands for "less than P others adopt the filter". Cast in this way, this public good provision problem is nothing else than a PD. However, strictly speaking, this is just for the illustration: a true public goods problem should be expressed in terms of an *n*-person PD game.

Looking at the PD payoff table, a Nash equilibrium can be found easily. A Nash equilibrium is such that the strategies that define it are best responses to each other. This can be determined by first putting oneself in the shoes of the row player and then in those of the column player. If the column player were to play C then the best response of the row player is D, since he gets a payoff of four by playing D rather than three by playing C. We correspondingly mark the four in bold. If the column player were to play D then the best response of the row player is D since he would obtain a utility of one instead of zero. We thus mark the one. Now, from the point of view of the column player, if the row player were to play C it is best for her to play D, and we mark the four in bold. If, on the other hand, the row player were to play D, her best response is D and we mark the one. We see that the only pair of strategies where both payoffs are marked is (D,D) which is thus the only Nash equilibrium of the game.

Because playing C gives the strictly worst payoff, a rational individual will choose D and hence will not convert to the air-preserving system. Since all players are rational and have complete common knowledge of the game, nobody will convert. Should someone install the filter, perhaps because he is unconditionally cooperative and does not act as a strict profit maximizer, he will bear a cost. However, he will ripe a benefit only if more than P people act similarly. The rest will free-ride on those, i.e. they will benefit from the improved air quality without paying anything. In conclusion, cooperation on public goods problems and many similar situations is not possible if the problem is cast as a PD game and the players are rational utility maximizers. Since we actually observe that cooperation is achieved in many such cases, either in society or in experiments, there must be other considerations playing a role in the issue. For example, state enforcement or taxes may do the trick by changing the utility values and thus the game itself. Another possibility is that people may have other-regarding preferences and are not completely selfish. These issues are a matter of many discussions today but we cannot devote more space to them here. The interested reader is referred to [3, 6, 7] for the state of the art of the subject.

Matching Pennies. This simple game was presented in section 2.1 and its payoff matrix is repeated here for the sake of simplicity. The game does not belong to the symmetrical class defined above, but this does not have any negative implication here. We study it because it is a simple example in which an equilibrium in mixed strategies appears naturally. If we proceed in the same way as we did for the PD we find that there is no pair of strategies that are best responses to each other as can be observed below. Thus, there is no Nash equilibrium in pure strategies. However, according to

	H	T
H	**1**,-1	-1,**1**
T	-1, **1**	**1**, -1

Nash's theorem there must be an equilibrium in mixed strategies. It can be found in the following way. Suppose that the column player plays strategy H with probability p and strategy T with probability $1-p$. In this case the expected payoff of the row player if he plays H is:

$$E_r[H] = p - (1-p) = 2p - 1$$

and if he plays T his expected payoff will be:

$$E_r[T] = -p + 1 - p = 1 - 2p$$

The row player will be indifferent between playing H or T when p is such that $E_r[H] = E_r[T]$, which implies $2p - 1 = 1 - 2p \Longrightarrow p = 1/2$. Assuming now that the row player plays H with probability q and T with probability $1-q$, one similarly concludes that $q = p = 1/2$ at the equilibrium. There is thus a Nash equilibrium in mixed strategies $\sigma^* = (1/2, 1/2)$ in which each strategy is played with probability $1/2$.

Hawks and Doves. This well known game is a metaphor for conflict that appears to capture some important features of social and geopolitical interactions. There are two possible strategies (behaviors): either a player is cautious (dove or D), or it is bold and aggressive (hawk or H). If two doves meet they fly away and each gets a certain payoff. If two hawks meet instead, they fight and get injured obtaining the worst possible payoff each. When a hawk is confronted with a dove the hawk gets the highest payoff of all, while the dove is punished but not as much as a hawk fighting against another hawk. A possible payoff bi-matrix for the game is represented by Table 3: each player would prefer to play hawk given that the other one plays dove,

	H	D
H	$-2, -2$	**2**, **0**
D	**0**, **2**	$1, 1$

Table 3: the Hawks-Doves game.

but if both play hawk they get the worst possible outcome. This game, which is also known under the names of "Chicken" or "Snowdrift", may model many situations in which "parading", "retreating", and "escalating" are common. One striking example of a situation that has been thought to lead to a Hawks-Doves dilemma is the Cuban missile crisis in 1962 [19]. Territorial threats at the border between nations are another case in point as well as bullying in teenage gangs. Other well known applications are found in the animal kingdom during ritualized fights [12].

In contrast to the Prisoner's Dilemma which has a unique Nash equilibrium that corresponds to both players defecting, the strategy pairs (H,D) and (D,H) are both Nash equilibria of the Hawks-Doves game in pure strategies, so the game is *antagonistic*, and there is a third equilibrium in mixed strategies in which strategy H is played with probability p, and strategy D with probability $1-p$, where $0 < p < 1$ depends on the actual utility values (the reader can easily check that $p = 1/3$ for the above payoff matrix by following the same line of reasoning that was used for the Matching Pennies game).

Stag Hunt. The origin of this game is attributed to a story told by the philosopher J. J. Rousseau in his 'A discourse of inequality". The story (with some variation with respect to the original) goes like this. Two hunters go out to hunt a stag which, if caught, could feed many people in the village. However, to successfully hunt a stag they need to coordinate and collaborate, given that a hunter cannot catch a stag alone. On the other hand, a hunter may easily catch a rabbit, say, that happens to pass within reach, without help by the other hunter. If these are the choices that are offered to the hunters, represented in the following payoff matrix by S (hunt the stag) and R (hunt the rabbit) respectively, what would be the rational choice? In the Stag Hunt (SH) mutual coop-

	R	S
R	**1**,**1**	2,0
S	0,2	**3**,**3**

Table 4: the Stag-Hunt game.

eration (S,S) is the best outcome, and a Nash equilibrium, as it can be easily checked by looking at the game's payoff matrix 4. However, there is a second equilibrium in which both players defect by going for a rabbit (R,R) and which is somewhat "inferior" to the previous one, although perfectly equivalent from a Nash theoretical point of view. The (R,R) equilibrium is less satisfactory yet "risk-dominant" since playing it "safe" by choosing strategy R guarantees at least a payoff of 1, while playing S might expose a player to a R response by her opponent, with the ensuing minimum payoff 0. Here the dilemma is represented by the fact that the socially preferable coordinated equilibrium (S,S) might be missed for "fear" that the other player will play R instead. There is a third mixed-strategy Nash equilibrium in the game, but it is commonly dismissed because of its inefficiency and also because it is not evolutionarily stable (see next section).

Games of the SH type are ubiquitous in society. The following example is adapted from [11]. Suppose that there is a group of workers in a firm and that each worker can contribute either a high effort to a common task, or a low effort (but not as low as to be fired). Effort is privately costly but the firm's output and a worker's wage are proportional to the effort contributed by the majority of workers. Clearly, if most workers contribute a low effort, there is no point for a single worker to put in high effort, as he will work extra hours without getting any reward. On the other hand, if he decides to put in low effort, it might perhaps happen that most other workers contribute high effort; in this case the lazy worker would enjoy the general positive result to which he did not contribute. Clearly, the best outcome is for the workers to all make an effort towards the common goal; however, there is a strong incentive to defect. The game is thus a Stag-Hunt and the above payoff table could be interpreted in the following way:

		Minimum effort of the rest of the workers: low	Minimum effort of the rest of the workers: high
worker's effort	low	1,1	2,0
worker's effort	high	0,2	3,3

Although the PD has received much more attention in the literature than the SH, the latter is also very useful as a metaphor of coordinated social behavior for mutual benefit. These aspects are nicely explained in [23].

Pure Coordination Games. We end this section with a description of a class of games that model many commonly occurring situations: the *pure coordination games*. A pure coordination game has the general normal form depicted in Table 5,

	1	2	...	k
1	u_1, u_1	$0,0$	...	$0,0$
2	$0,0$	u_2, u_2	...	$0,0$
...	...	...	...	...
k	$0,0$	...	...	u_k, u_k

Table 5: a general payoff matrix of a two-person pure coordination game.

with $u_i, u_i > 0$, and $u_i, u_j = 0, 0, i \neq j, \forall i, j \in [1, k]$, where k is the number of strategies available.

A simple coordination game is the *driving game*. In some countries people drive on the right side of the road, while in others they drive on the left side. This can be represented by the pure coordination game of Table 6. There are two Nash equilibria in pure

	right	left
right	$\mathbf{1}, \mathbf{1}$	$0,0$
left	$0,0$	$\mathbf{1}, \mathbf{1}$

Table 6: the driving game.

strategies: (right, right) and (left, left) and obviously there is no reason, in principle, to prefer one over the other, i.e. the two equilibria are equivalent. However, while some countries have got accustomed to drive on the left such as the UK, Australia, and Japan, others have done the opposite such as most European countries and the USA.

Such *norms* or *conventions* have stabilized in time and are the product of social evolution. We shall have more to say on this later; for the time being, however, please note that such a norm may not be the product of historical accidents, it can just be "ordered", as Sweden did in 1967 switching from left to right overnight (of course the switch had been prepared long in advance). The reader can easily check that there is a third equilibrium in mixed strategies in the driving game which consists in playing left and right with probability $1/2$ each. I wouldn't advise to play the game in this way on a real road though!

Coordination games type situations are very common and important in society and it is obviously of collective interest to be able to solve this kind of problems successfully. For example, when I'm driving in the UK, it is of the utmost importance (for me and for the car passengers I happen to meet on the road) that I be able to immediately coordinate on the prevailing convention. Bargaining problems like selling and buying goods can also be seen as coordination problems to some extent, although there are conflicting interests. In conclusion, being able to coordinate on the "right" equilibrium is beneficial for all parties. But standard game theory doesn't help here, as the different strategies are simply labels devoid of any particular meaning. Indeed, looking back to the general two-person finite coordination game (Table 5), one sees that all the diagonal entries of the bi-matrix are Nash equilibria. Pure coordination games thus make it apparent that there may be several equivalent Nash equilibria in a game, a phenomenon that we already met in the Hawks-Doves and Stag-Hunt games. How to pick up one of those is called the *equilibrium selection problem*. This problem has plagued game theory for decades and it is not close to be solved. In fact, it might well be that this degeneracy is absolutely natural and unavoidable to some extent. Many devices have been proposed to refine and select an equilibrium such as the notions of saliency, trembles, and perfect and proper equilibria among others, but no solution concept has proven really conclusive. The concept of "saliency" is particularly easy to grasp: it simply appears that some "extra game" consideration, probably of psychological or cultural order, plays an important role in the strategy choice. For example, T. Schelling showed early on that if two New Yorkers are asked to meet in town at a certain time but without specifying the place, most will choose Grand Central Station (two foreigners would have probably coordinated on the Empire State Building instead). This implicit consensus is called a "focal point" [22].

The literature on equilibrium refinement and equilibrium selection is very rich. We cannot go into the details here; the interested reader can consult Myerson's book [13] for an introduction to the main ideas and a summary of the issues.

This section has tried to offer, in a simplified and intuitive way, an introduction to the main concepts of standard game theory and a number of illustrations of how simple games may model many commonly occurring important social interactions. We have seen that an important idea in game theory is that of the Nash equilibrium. The NE theory rests on a number of assumptions such as intelligence and rationality of agents, and common knowledge of the latter that, while mathematically formalizable, may not fully correspond to the way in which actual socio-economic agents take their decisions. The NE theory has thus *normative* value, but its explanatory power is limited, as can be concluded from the innumerable experiments that have been conducted. As I said before, the extent to which game theory can be applied to the real world is still an open question and a matter of debate. Apart from these philosophical questions, we have seen that, technically, often there is more than one equilibrium in a game, a fact that requires to be taken care of. In the next section I shall present another way of approaching interactive decision-making which does offer an answer to some of the open questions but not to all of them, and which also raises new questions of its own.

3. Evolutionary game theory

An idea which is conspicuously absent from what has been said so far is the concept of *dynamics*. In the real-world, however, interactions do not have the "instantaneous" character that the theory seems to attribute to them. Real agents do observe their own behavior and the behavior of other agents, and they probably adapt their strategies over time as they learn more about the game and about the environment. Assuming that such a dynamical view of the players is a reasonable one, where does it lead us? Do the dynamics converge to some stable state or not? And if yes, to what state or states? The fundamental ideas of *Evolutionary Game Theory* (EGT) have been inspired by the population dynamics view of biological evolution and indeed, among the founders of the field one finds mainly theoretical evolutionary biologists. Biological evolution, *in nuce*, works with the following ideas:

- A population of individuals
- A source of variation that provides diversity through, say, recombination and mutation of genetic material
- A selection mechanism that favors fitter variants over others that are less adapted to the current environment

When translated into the language of game theory, these ideas correspond to the following analogy:

- There is a very large population of players each one of which plays a fixed strategy in a given two-person game
- There is no need for the players to be rational in the sense explained in the previous section; players just do what they have been "hard-wired" to do, so to speak
- Members of the population are anonymously matched in pairs at random to play the game and each one receives its corresponding payoff

- There is a selection mechanism that causes the proportion of those strategies that get a larger payoff in the population to increase. Conversely, the share of individuals using less successful strategies tends to decrease

The previous intuitive ideas have been formalized in what is called *Replicator Dynamics* [25, 29] to be explained in the next section.

3.1. Replicator dynamics

In *replicator dynamics* an infinite population is considered for mathematical convenience. A population *state* at time t is a distribution of n pure strategies given by a vector of population frequencies $x(t) = (x_1(t), x_2(t), \ldots, x_n(t))$ where $x_i(t)$ represents the fraction of the population using strategy i. Alternatively, this *polymorphic* population state can also be seen as the whole population playing the corresponding mixed strategy. This equivalence will be useful when examining the relationships between Nash equilibrium and the stable states of the dynamics. As said above, individuals are randomly drawn in pairs from the population and play the given game. Let's call $u(\mathbf{e}_i, x)$ the expected utility (payoff) of the pure strategy i in state x, and let $u(x,x)$ be the population *average payoff* defined as:

$$u(x,x) = \sum_{i=1}^{n} x_i u(\mathbf{e}_i, x).$$

Now suppose that payoffs are analogous to fitness in biology in the sense that they measure the amount of offspring that inherit the same trait, i.e. the "reproduction" capabilities of the corresponding strategies. With this analogy in mind, the rate of frequency change of any strategy in the population will be proportional to the relative difference between its average payoff and the population average payoff. Therefore, we can write:

$$\frac{dx_i}{dt} = \dot{x}_i = x_i \left[u(\mathbf{e}_i, x) - u(x,x)\right] = x_i(u_i - \overline{u}), \tag{1}$$

where I have written $\overline{u}$ for the average population fitness, and I have simplified $u(\mathbf{e}_i, x)$ to u_i. Although I have suppressed a number of details in the interest of simplicity, this system of linear differential equations essentially represents the replicator dynamics. The reader will find a full account in [29]. One readily sees from the above equation that strategies that do better than the average will tend to increase their share in the population, while those that do worse will tend to decline. However, it is also easy to see that no strategy that is present at some time t with a frequency $x_i > 0$ can totally disappear at a some future time. Conversely, and more important, also note that the model is such that no strategy that was not already present at the beginning can appear during evolution since the dynamics does not have an element of innovation such as a concept of mutation. Thus, replicator dynamics highlights the role of selection. We shall see that the concept of variation is also essential to fully characterize an evolving population of strategies.

In the above pseudo-biological interpretation strategies "reproduce" in proportion to their relative payoffs. Indeed, evolutionary game theory has been conceived for biological contexts in which Darwinian selection has a well established meaning. However, it is difficult to translate these ideas in terms of social behavior. Social agents in their lifetime use additional mechanisms such as imitation, trial and error, and all forms of learning. Let's consider imitation, which is probably the most primitive type of learning. The good news are that if we endow agents with the possibility of imitating more successful strategies, the equations that result turn out to be indistinguishable from standard replicator dynamics equations. Replicator dynamics can thus be seen in social terms as the population strategy dynamics induced by the progressive adoption of the more successful strategies by agents using imitation. Although this view is extremely limited with respect to reality, it has at least the advantage of being socially realistic, and mathematically fully formalizable.
Now, given this simple dynamics, the next question is: what are its stable states, and, is there a relationship whatsoever between the latter and the previous static concept of Nash equilibrium? Before discussing these issues we shall present an alternative evolutionary view of game equilibrium: the notion of an evolutionarily stable strategy.

3.2. Evolutionarily stable strategies

The concept of an *evolutionarily stable strategy* (ESS) was introduced by Maynard-Smith and Price in 1973 [9, 12]. It is a very simple and elegant idea and it has a close connections with the NE concept, which establishes a clear link between the standard and the evolutionary approaches. The key idea of an evolutionary stable strategy can be expressed simply: a strategy is an ESS in a population if it cannot be displaced (invaded) by a sufficiently small number of individuals playing an alternative game strategy. The notion can be easily put in mathematical form in the usual large population setting in which pairs of individuals are repeatedly drawn at random and play the given game.
Assume that all the individuals in the population play (mixed or pure) strategy x and call $y \in \Delta(S_i)$ the strategy played by the "mutant" individuals, which are few in number. The payoff to strategy $x \in \Delta(S_i)$ when played against strategy y is $u(x,y)$.
Let $\varepsilon \in (0,1)$ be the share of mutants in the population. Given that pairs of players are drawn from the population with uniform probability to play the game, the probability that a player will play y is ε, while the probability of playing x is $1-\varepsilon$; this is equivalent to playing the mixed strategy $w = \varepsilon\, y + (1-\varepsilon)\, x$. The payoff of the "established" strategy x against w is thus $u(x,w)$ and that of the "mutant" strategy y is $u(y,w)$. Strategy x will be evolutionarily stable if

$$u[x, \varepsilon y + (1-\varepsilon)x] > u[y, \varepsilon y + (1-\varepsilon)x],$$

$\forall y \in \Delta(C_i), y \neq x$, and granted that the share of mutants ε is "sufficiently small".
There is an alternative formulation of evolutionary stability that highlights its relationships with Nash equilibrium. It can be stated through the following two conditions:

(2) $$u(x,x) \geq u(y,x) \quad \forall y,$$

(3) $$u(x,x) = u(y,x) \Rightarrow u(x,y) > u(y,y) \quad \forall y \neq x$$

This can be read as: x is evolutionarily stable if either x is a strict best reply to any y, or it is as good against itself as any other mutant, and x is a better reply to any mutant y than y is to itself.
The first condition is equivalent to the NE equilibrium for the underlying bilateral game. However, the second condition states that, in addition, an ESS x must also be a best reply to any $y \neq x$ than y to itself. In conclusion, the concept of an ESS can be seen as a refinement of NE and $\Delta^{ESS} \subset \Delta^{NE}$, which means that some Nash equilibrium may not be an ESS, a result that promises to alleviate somewhat the equilibrium selection problem.

An example should be useful at this point to illustrate the meaning of what has been said. We shall take the Hawks-Dove game again (sect.2.2), of which we report the payoff matrix below for convenience. Let's assume that the population is entirely

	H	D
H	$-2,-2$	$\mathbf{2},\mathbf{0}$
D	$\mathbf{0},\mathbf{2}$	$1,1$

constituted by individuals playing "dove" (D) at the beginning. Now imagine that for some reason one or a few players switch their strategy from D to "hawk" (H). This can occur because of an error, or a "mutation", if we want to keep the biological analogy. Now, while almost all encounters will be between two doves (each earning a payoff of 1), from time to time a mutated hawk will be selected to play a dove; in this case the hawk's payoff will be 2 and the dove's 0. Of course, hawk-hawk encounters are very unlikely if the fraction ε of hawks is low enough. Thus, the few hawks will be more successful on the average than the doves and, by virtue of proportional reproduction, their number in the population will start to grow. This is because the condition $u(H,D) \leq u(D,D)$ is not satisfied since $2 > 1$. Consequently, outcome (D,D) is not evolutionarily stable.
The same considerations apply to a population of pure hawks. Invasion by a few doves is possible for the same reason: two hawks will earn -2 each, while a mutant dove against a hawk will earn 0. Thus the number of doves in the population will increase. Consequently, outcome (H,H) is also evolutionary unstable. This leads us to the first important result namely, that of the three NE of the game, the two in pure strategies are not ESS which, at least for this game, solves the equilibrium selection problem. The second consideration is the following: if neither a population of pure doves nor one of pure hawks is evolutionarily stable, when will it become stable? If we call p the proportion of doves in the population, the answer is that the population will be evolutionarily stable when p is equal to the probability with which strategy D is played in the corresponding Nash equilibrium.

So, the set of ESS is a proper subset of the set of NE; however, there are simple games in which this subset is empty, i.e. the game in question does not have any ESS. Consider, for example, the following children game, which is commonly called the *rock-scissors-paper* game (see, e.g. [29, 26]). In this game, Rock (R) beats scissors (S), S beats paper (P), and P beats R. We assume that beating the opponent gets a payoff of 1 from him, while nothing happens if both players choose the same action and the utility is 0 in this case, see Table 7. It is easy to see by inspection that there is no

	R	S	P
R	$0,0$	$\mathbf{1},-1$	$-1,\mathbf{1}$
S	$-1,\mathbf{1}$	$0,0$	$\mathbf{1},-1$
P	$\mathbf{1},-1$	$-1,\mathbf{1}$	$0,0$

Table 7: the Rock-Scissors-Paper game.

NE equilibrium in pure strategies, as there is no cell in the table where the highlighted best replies meet. The only NE of the game is $x = (1/3, 1/3, 1/3)$ in which players randomize uniformly among their available strategies. Let's now consider the payoff matrix M corresponding to the row player:

$$M = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$$

Since there is only one NE (x,x) in mixed strategies, we can ask whether it is also an ESS. If we choose, for instance, strategy R as a possible mutant strategy y, that is $y = (1,0,0)$, then simple matrix multiplication will tell us that

$$xMx = yMx = 0,$$

and,

$$xMy = yMy = 0.$$

Given that $xMx = u(x,x)$, $yMx = u(y,x)$, $xMy = u(x,y)$, and $yMy = u(y,y)$, this yields

$$u(x,x) = u(y,x) = 0,$$

$$u(x,y) = u(y,y) = 0.$$

The first equation satisfies the first condition for a strategy x to be an ESS (Eq. 2), but the second equation violates the condition expressed by Eq. 3 and therefore x cannot be an ESS.

The possible non-existence of an ESS can obviously be a problem, but we shall see that the replicator dynamics description offers some alternatives stability concepts that can rescue the evolutionary approach.

4. Replicator dynamics and stability

In this section we can only hint at the main results of stability analysis for the replicator dynamics equations and their relationships with the Nash equilibrium concept. The subject is a vast one, and the interested reader is referred to the books [8, 29, 26] for a full analysis.
In the first place, note that every NE is a fixed point of the replicator dynamics [29]. Now, since we have seen in the previous section that $\Delta^{ESS} \subset \Delta^{NE}$, it follows that all ESS are among the fixed points of the replicator dynamics. However, the replicator dynamics equations may have attractors that are not Nash equilibria. A trivial example are the vertices of the strategy simplex for which we have $\dot{x}_i = 0$ from Eq. 1 with $x_i = (0, \ldots, \mathbf{e}_j, \ldots, 0)$ and $j \in [1, m_i]$. These degenerate monomorphic population strategy profiles are formally fixed points, since the dynamics never actually moves out of them when the whole populations plays initially the same pure strategy. Of course the degeneracy arises because most of these pure strategies are not NE. In addition, it has been shown [29] that there can also be polymorphic fixed points – i.e. populations with at least two different strategies – that do not correspond to NE. In summary, rather than providing a refinement of the set of NE, the replicator dynamics lead to an *extension*.
However, if one requires additional criteria of robustness and stability to be obeyed by the stationary points of the dynamics, then more precise conclusions can be obtained. The dynamical systems concepts of *asymptotic* and *Lyapunov* stability have been used to provide further characterization (see e.g. [29, 8]).

Asymptotic Stability. Let $\dot{x} = f(x)$ represent the state of a dynamical system in a set $D \subset \mathbb{R}$, and let x^* be an equilibrium point, i.e. $f(x^*) = 0$.

- An equilibrium point x^* is said to be Lyapunov stable if, given a neighborhood V_1 of $x^*, \exists\, V_2 : \forall x(.)$, if $x(0) \in V_2, \Longrightarrow x(t) \in V_1, \forall t \geq 0$
- An equilibrium point x^* is said to be asymptotically stable if there exists some neighborhood V of x^* such that for any path $x(.)$, if $x(0) \in V$, then $\lim_{t \to \infty} x(t) = x^*$

Thus, asymptotic stability requires Lyapunov stability, which means that any trajectory that starts sufficiently close to the equilibrium point at $t = 0$ will remain close to the equilibrium as t increases. But to be asymptotically stable, the equilibrium state must also meet the second condition, i.e. that any trajectory that starts sufficiently close to the equilibrium converges to it as t increases without bound.

The concept of asymptotic stability has been found to be useful in reducing the set of replicator dynamics equilibria. These results are summarized in the following two theorems (see [29, 8]).

THEOREM 2. *If x^* is an asymptotically stable state of the RD, then the mixed strategy $\sigma^* = x^*$ is such that the pair (σ^*, σ^*) is a symmetric NE.*

THEOREM 3. *If σ^* is an ESS, then the population state vector $x^* = \sigma^*$ is asymptotically stable (the converse is not necessarily true, and thus the condition is only sufficient).*

These results provide interesting links between the static Nash equilibrium and ESS approaches, and the dynamical view as expressed by the replicator dynamic. Furthermore, they also provide a restriction on the set of asymptotically stable states.
Thanks to the concept of dynamic stability, it is also possible to sometimes shed some light in cases where there is no ESS, such as the rock-scissors-paper game of the previous section, because the existence of an ESS is a sufficient but not necessary condition for asymptotic stability. When one applies dynamic stability considerations to the rock-scissors-paper game, one finds that the unique interior fixed point $(1/3, 1/3, 1/3)$ which corresponds to the unique Nash equilibrium of the game, is not asymptotically stable. However, it is Lyapunov stable, which implies that small perturbations at equilibrium will only have small effects. Therefore, the state has some of the properties of an equilibrium state; Fig. 3 provides a graphical illustration of this behavior.

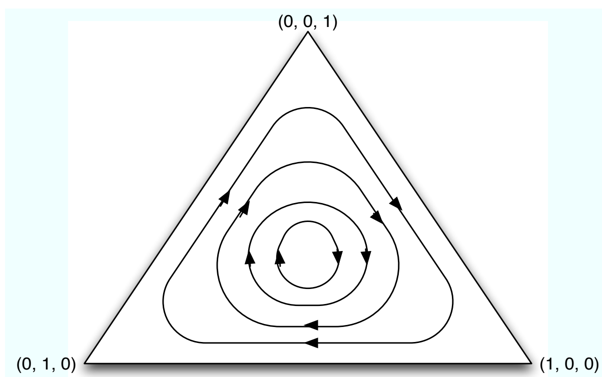

Figure 3: replicator dynamics in the rock-scissors-paper game.

5. Replicator dynamics of some simple games

To wrap-up the discussion of evolutionary game theory, this section presents a qualitative discussion of the replicator dynamics of some of the games of section 2.2. The discussion is inspired by the one appearing in [7]. Let's consider again symmetric 2×2 games as represented in Table 1.
Assume that $P\{s1\} = p, P\{s2\} = 1 - p$. Thus:

$$E[s1] = pu_1 + (1-p)u_2, \;\; E[s2] = pu_3 + (1-p)u_4$$

Let's call δ the net gain (or loss) arising from choosing the first strategy $s1$ over the second $s2$:

$$\delta = E[s1] - E[s2] = pu_1 + (1-p)u_2 - [pu_3 + (1-p)u_4],$$

which is a straight line $\delta = A + Bp$, with

$$A = (u_2 - u_4), \quad B = (u_1 + u_4) - (u_2 + u_3) \tag{4}$$

The replicator dynamics will favour the relatively more successful strategy. We thus study the behavior of $\Delta\delta/\Delta p = f(p)$. Remark that δ is positive when the first strategy is relatively more successful than the second, which means that the first strategy share increases at the expenses of the second strategy and thus p raises. When δ is negative, the second strategy spreads and the first shrinks and thus p declines. When $\delta = 0$ p is stationary.
For the Prisoner's Dilemma, strategy s1 is called C, strategy s2 is called D, and we use the PD payoff numerical values appearing in the game's table of section 2.2. The behavior of δ as a function of the probability p of playing D is shown in Fig. 4. We

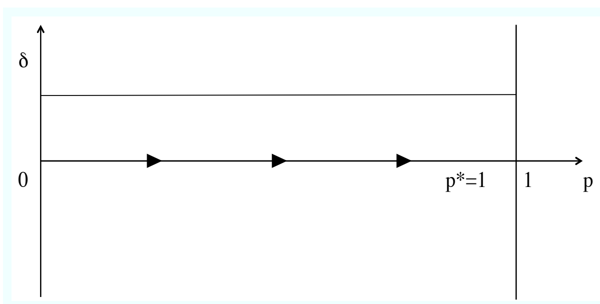

Figure 4: replicator dynamics in the Prisoner's Dilemma.

easily find that $E[C] = 3p + 0 \times (1-p) = 3p$, and $E[D] = 4p + 1 \times (1-p) = 3p + 1$. Therefore, $\delta = E[D] - E[C] = 1$ for all values of p. Thus defection always predominates and, at equilibrium, $p = p^* = 1$. D is the dominant strategy, the unique Nash equilibrium, and the unique evolutionary equilibrium.
For the Hawks-Doves game, using the game's payoff table of section 2.2 and replacing the appropriate numerical values in Eq. 4 we have: $A = u_2 - u_4 = 2 - 1 = 1, \; B = (u_1 + u_4) - (u_2 + u_3) = (-2+1) - (2+0) = -3$, which gives $\delta = 1 - 3p$; As a consequence, $\delta > 0$ when $p < 1/3$, it is negative when $p > 1/3$, and $\delta = 0$ when $p = 1/3$. Therefore, the share of hawks increases with increasing p up to $p = 1/3$. If one starts with a share of hawks in the population $p > 1/3$, then it decreases, as $\delta < 0$ in this case. At equilibrium $\delta = 0$ and the population is composed by $1/3$ hawks and $2/3$ doves, which is also the Nash equilibrium in mixed strategy for this game. In conclusion, the two pure strategies are unstable (they are not ESS in the language of section 3.2) and the only evolutionary stable state is the polymorphic population. This eliminates

two of the three NE of the game. This situation is schematically represented in Fig. 5. A similar analysis for the Stag-Hunt game (see the payoff table in section 2.2) yields

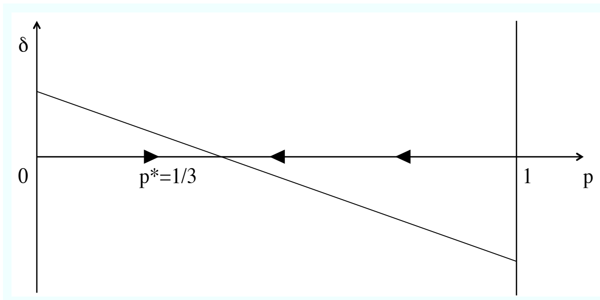

Figure 5: replicator dynamics for the Hawks-Doves game.

$\delta = E[R] - E[S] = 2 - p - (3 - 3p) = -1 + 2p$. This gives $\delta > 0$ when $p > 1/2$, $\delta < 0$ when $p < 1/2$, and $\delta = 0$ when $p = 1/2$. Thus the share of stag hunters (S) increases with decreasing p when $p < 1/2$, and it decreases with increasing p for $p > 1/2$. The equilibrium point $p = 1/2$ is therefore unstable and the evolutionary equilibrium will be a monomorphic population of either stag hunters or rabbit hunters. This is schematically depicted in Fig. 6 in which the direction of the arrows intend to show that the stable states are the pure strategies. One can also say that when the initial population composition falls in the basin of attraction of stag hunting, the population will converge to all stag hunters, while the reverse happens when the population composition is such that it falls into the basin of attraction of hunting rabbits. With respect to NE, we see that the evolutionary approach has been able to eliminate the NE in mixed strategies.

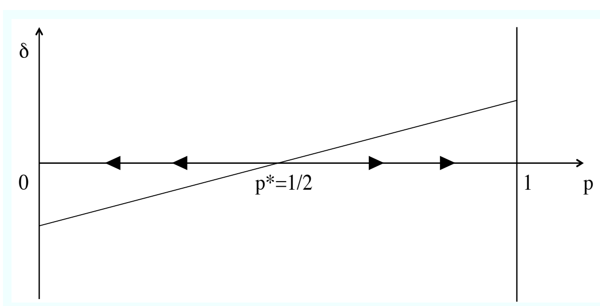

Figure 6: replicator dynamics for the Stag-Hunt game.

We conclude this section with a few general observations concerning the evolutionary approach applied to game theory:

- EGT takes a fresh, dynamical view of game theory

- EGT partially succeds in the equilibrium selection problem
- EGT does away with rationality and common knowledge problems but the price to pay is stereotyped behavior
- EGT does not seem to have more explanatory power than standard game theory does when it comes to observed social phenomena

Beyond the simple symmetric games studied here, evolutionary game theory has been extended to asymmetric games and multipopulation models [29]. There have also been studies of the effect of finite populations and other deviations with respect to the infinite populations of the basic theory [16]. One particularly promising direction is the inclusion of more realistic forms of learning for the agents, such as reinforcement learning and fictitious play. In conjunction with the notion of stochastic stability, these new stochastic dynamics constitute an improvement with respect to the simple imitation learning implicit in the replicator dynamics. We do not have enough space here for a decent treatment of these extensions. For an excellent introduction, the reader is referred to [26]. However, we shall at least challenge at length one of the main assumptions of evolutionary game theory: the mixing population assumption. In the biological jargon, a mixing, or *panmictic* population is one in which any member of the population may meet any other member with equal probability. This assumption is essential for mathematical reasons to insure that the successive random choices of individuals that play the game are completely uncorrelated [29]. In the next section, we shall see that when populations are endowed with a relational or geographical structures there are important consequences on the evolutionary equilibria and their basins of attraction.

6. Evolutionary dynamics in structured populations

Our everyday experience tells us that biological populations, and especially human ones, cannot be considered to approach the mixing behavior to any remarkable extent. On the contrary, social and others constraints originate a large amount of *locality* in the interactions between agents belonging to the population. Recognizing this fact, some early work was done assuming that populations have a two-dimensional structure that can be modeled by grids, which are a kind of regular graph, such as the one depicted in Fig. 7.

This diagram is to be interpreted in the following way: each point represents a player in the population, and a link between two individuals means that there is an interaction between them. Moreover, to avoid awkward border conditions, it is usually assumed that the grid wraps around, i.e. it is is actually a torus. The fact that this kind of spatial structure with only local interactions in a restricted neighborhood favors cooperation in the PD was first remarked by Axelrod in the iterated case [2], in which there are however other reciprocity mechanisms that may lead to cooperative behavior. The idea was used by Nowak and May for the one-shot case in 1992 [15]. The one-shot case is of course harder for cooperation to emerge in the PD since we know that the

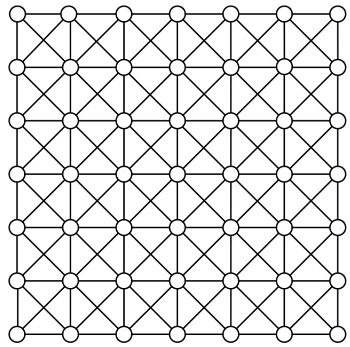

Figure 7: a lattice-structured population. The neighborhood of an individual is constituted only by the cells directly connected to it. In the figure this comprises eight cells.

unique NE and the only evolutionary stable state is defection. However, Nowak and May's and many other successive studies showed that cooperation may emerge and be stable. How can a mere change in the population structure sustain cooperation in the PD? The answer is that, when players' strategies are randomly assigned to grid points, statistical fluctuations may cause some cooperator cell to be surrounded by a sizable number of cooperators. In certain conditions, several of those cooperators only interact with cooperators and thus accumulate a large payoff. Since strategies are changed proportionally to the success of neighbors in the grid, it follows that sufficiently large clusters of cooperators keep their strategies and thus may survive in the long run. Of course, the frontiers of the clusters do not stay still, they are rather "liquid" and change often their shapes; in other words, if the strategy switching rule contains a stochastic element, these clusters never freeze. A typical "quasi-equilibrium" situation is the one shown in Fig. 8.

Evolutionary games on grids have been the object of many investigations. The reader can refer to [17] for a recent overview.

These results were quite remarkable, given the theoretical impossibility for cooperation to survive. However, although regular grids are a first easy to treat approximation to structured societies, they are still far from actually observed social networks. For example, in a real social network there is no reason for all the agents to have the same number of connections; on the contrary, this number will in general vary from one individual to the other. Moreover, grids do not show the *small-world* effect, which simply means that the average distance (in the path-length sense) between two arbitrary nodes is typically small and grows slowly as $O(logN)$ where N is the number of vertices in the network . Researchers have found that many man-made, biological and social networks are small-worlds. Indeed, in recent years large complex networks have been investigated in detail for the first time thanks to the availability of abundant computer-based data and the existence of new techniques inspired by statistical physics

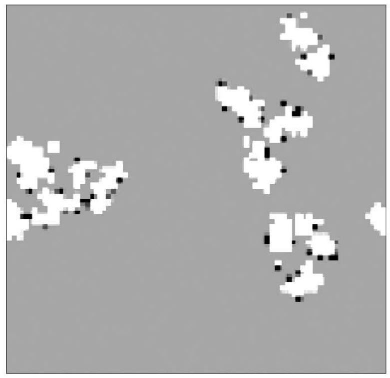

Figure 8: cooperator clusters (white) surviving in a sea of defectors (dark gray). Cells marked in black along the borders of the clusters have just changed strategy.

approaches. The flurry of studies was started by the seminal paper of Watts and Strogatz [28]; for an excellent review of the field, the reader is referred to [14]. In the rest of this section, I shall present the results of simulations of evolutionary games based on network-structured societies. Since these networks can be modeled as graphs, we first need a minimum of notation and concepts from graph theory.

The network of agents is represented by an undirected graph $G(V,E)$, where the set of vertices V represents the agents, while the set of edges (or links) E represents their symmetric interactions. The population size N is the cardinality of V. A neighbor of an agent i is any other agent j such that there is an edge $\{ij\} \in E$. The set of neighbors of i is called V_i and its cardinality is the degree k_i of vertex $i \in V$. The average degree of the network will be called $\bar{k}$. In general, the graph G is to be understood as a *relational* entity. In other words, links between vertices only model a relationship which is in general independent of a distance defined for an underlying metric space. Think, for example, of links between web pages or friendship ties.

We shall use a few well-defined graph types to represent societies. The Erdös-Rény random graph is a model which, in its simplest form, consists of N vertices joined by edges that are placed between pairs of vertices uniformly at random. In other words, each of the possible $N(N-1)/2$ edges is present with probability p and absent with probability $1-p$. The model is often called $G_{N,p}$ to point out the fact that, rigorously speaking, there is no such thing as a random graph, but rather an ensemble $G_{N,p}$ of equiprobable graphs [4]. The random graph will mainly be used for comparison purposes, as it is not a good model for a social network.

The *degree distribution function* (DDF) $p(k)$ of a graph G represents the probability that a randomly chosen node has degree k. Random graphs are characterized by DDF of Poissonian form $p(k) = \bar{k}^k e^{-\bar{k}}/k!$, while social and technological real networks often show long tails to the right, i.e. there are nodes that have an unusually large number of neighbors [14]. In some extreme cases the DDF has a power-law form $p(k) \propto k^{-\gamma}$; the tail is particularly extended and there is no characteristic degree. Graphs enjoying this property are called *scale-free* and they have been found in several fields. For example, the Internet router graph is scale-free, and so are the portions of the web that have been measured, article citations in scientific publications, as well as the Wikipedia oriented graph of connected articles [14]. Scale-free graphs are not a good representation of social networks [1, 14] but they are a much better approximation to them than regular graphs such as grids or random graphs. Finally, we shall use a more realistic model of a social network, and a real measured social network representing the graph of co-authorship coming from the field of evolutionary computation, in which vertices represent authors and there is an edge between two vertices if the corresponding authors have published at least a paper together.

Now the question is, how do agents playing some simple bilateral games behave on such structured societies? To answer this question, we shall use again the three classical social dilemmas as examples namely, the PD, the HD, and the SH games. For the rest of the section, we shall use the following generic payoff table for the three games:

	C	D
C	R,R	S,T
D	T,S	P,P

The reason is that we shall present cooperation results in a unified way not only for a fixed set of numerical payoff values, but rather for a whole portion of the parameter's space, which will allow us to detect general trends much more easily and reliably. In the table, strategies C and D are the usual ones for the PD, while they correspond to hawk and dove for the HD game, and to stag and rabbit for the SH respectively. In this matrix, R stands for the *reward* the two players receive if they both cooperate (C), P is the *punishment* for bilateral defection (D), and T is the *temptation*, i.e. the payoff that a player receives if it defects, while the other cooperates. In this case, the cooperator gets the *sucker's payoff* S. In the three games the condition $2R > T + S$ is imposed so that mutual cooperation is preferred over an equal probability of unilateral cooperation and defection. The relative ordering of the payoffs defines the game at hand. As we already know, for the PD the payoff values are ordered numerically in the following way: $T > R > P > S$. In the HD, the order of P and S is reversed yielding $T > R > S > P$, while for the SH we have: $R > T > P > S$. For a detailed discussion of the three games we direct the reader to sections 2.2, 3.2, and 5.

The dynamics is analogous to the replicator dynamics of section 3.1 with two major differences. First, the pairing of players is restricted to the immediate neighborhood V_i of a player i, instead of being drawn from the whole population. Second, the population size is necessarily finite, and usually relatively small with a few thousand individuals at most. Both changes represent major deviations from the assumptions of the theory. The consequences are that the model is much closer to reality but the previously developed mathematical formalism cannot be applied. This is the reason why most work in this field has been done by means of numerical simulations.

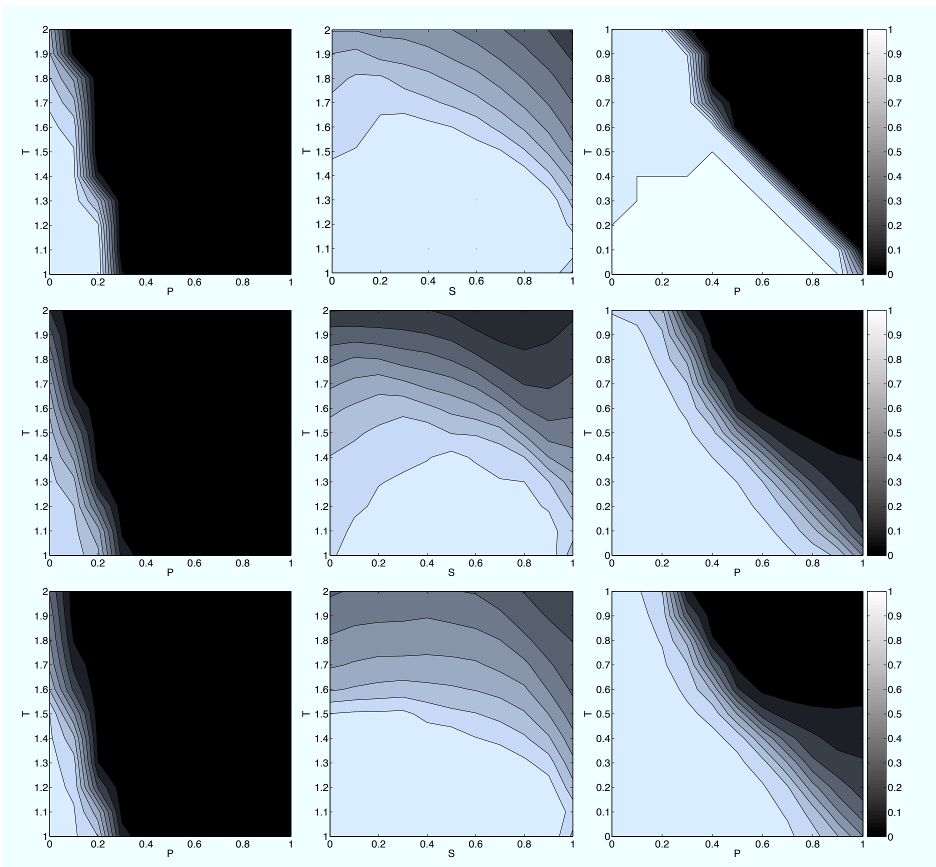

Figure 9: level of cooperation at the end of the simulation. From left to right: PD, HD, SH; from top to bottom: the scale-free case, an artificial social network model, and the real collaboration network. Results are averaged over 50 runs for each game, each network structure, and parameter set. Lighter areas are more cooperative.

The following figure 9 shows the average degree of cooperation for the three games on three different types of networks, for systems having attained a steady-state, after the transient equilibration period is well over. Initially, an equal amount of cooperators and defectors is randomly distributed among the graph vertices. As expected, the region in which cooperation is possible is much more restricted in the PD than for the other two games. Cooperation is more widespread for the HD, as mutual defection is the worst outcome in this game. For the PD and the SH, cooperation is sensitive to the "punishment" level P, for a given T, with the PD being influenced in a higher degree. Concerning the HD, one can see that the S parameter has moderate influence on cooperation for a given T. We also notice that the transition from cooperation to defection is much steeper in the PD and SH cases than for the HD which is easy to understand, given that the only evolutionarily stable state is the mixed one for the HD, while there are two monomorphic stable states for the SH.

Now, interpreting the results in terms of the three different topologies, scale-free networks are the structures that yield the highest cooperation levels for the three games (compare the top row with the second and third rows of Figure 9). This is encouraging but, as we already said, scale-free graphs are not a likely structure for actual social interaction networks. Nevertheless, the results are good for social networks too (second and third rows in the figure). Even for the PD, which is the hardest case for cooperation to emerge, there is a sizable parameter's region in which cooperation may thrive, while the theory for mixing populations would lead to all-defecting states for the whole space. For comparison, a mixing population would be represented by a complete network in graph terms, and the parameter space would be all black meaning full defection (not shown here). The same happens when the network is a random graph.

The HD and the SH are more favorable in terms of cooperation than the PD, still, the cooperation levels that can be reached on social nets are enhanced with respect to the mixing population case also for these games.

What is the mechanism that allows cooperation to emerge in such structured populations? We use the PD as an example since it is the hardest case, as we said above and we refer the reader to [10] for a fuller discussion. The answer is to be found in the clustering structure of any social networks which is also called its *community structure*. A community can be loosely defined as a group of highly connected vertices having few connections with vertices belonging to other communities. There are several algorithms that can partition the network into a set of disjoint or overlapping communities. When this community or cluster partition is determined for our social graphs, the important observation is that, independent of the cooperation level, in most communities either cooperators or defectors predominate.

In Figure 10 a portion of the evolutionary computation collaboration graph is depicted distinguishing between cooperators and defectors for the PD. As noted above, tightly-bound communities are mostly composed of players with the same strategy. Although only a small portion of the whole network is shown for reasons of clarity, I could have chosen many other places as the phenomenon is widespread. Cooperators tend to "protect" themselves by occupying sites with many links toward other cooperators. On the other hand, a cooperator like the central one in the largest defecting community will have a tendency to become a defector since its neighbors are nearly all defectors; but when its highly connected "wealthy" cooperator neighbor on the left of the figure is probabilistically selected to be imitated, then it will certainly become a cooperator again. So, the rare cooperators that are not tightly clustered with other cooperators will tend to oscillate between strategies. In fact, It turns out that in all cases a cooperator is surrounded by a large majority of cooperators, whereas a defector mainly interacts with other defectors. Thus the mechanism for the emergence and maintenance of cooperation is qualitatively similar to the one described for regular grids, with the important differences that now it is the intrinsic cluster structure of the network and

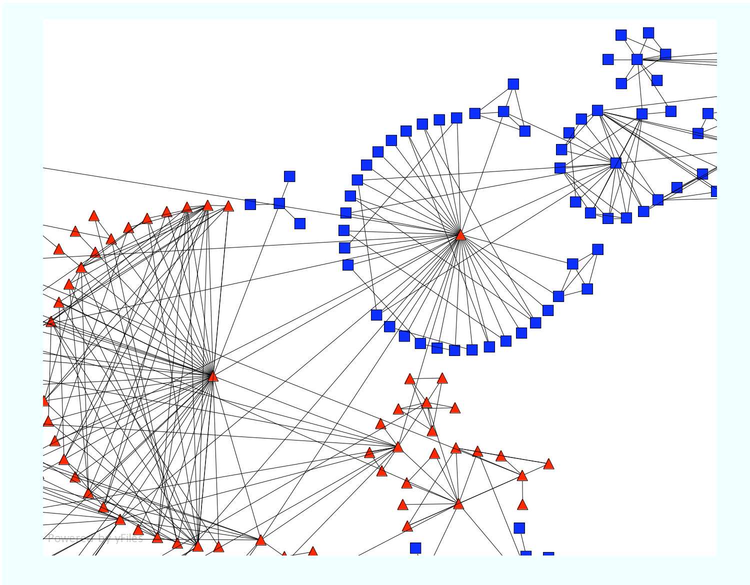

Figure 10: communities: cooperators are represented by triangles and defectors by squares.

its degree inhomogeneity that plays an important role, and both phenomena have been observed in all social networks studied so far.

One question that remains is: given that real social networks do not stay fixed once and for all, what is the effect of the co-evolution of the players' strategy and of the underlying network itself? Indeed, in real life new individuals may join or leave the network at any time, and links may be added or deleted between individuals already in the network. In our context, depending of the kind of game and its payoffs, individuals might contemplate dismissing some interactions and establishing new ones. Research along this line is very recent but there are already promising results showing that the added degrees of freedom of changing the graph topology on a local or global scale using social-like decision-making rules tend to provide another mechanism that promotes cooperation. We don't have space to discuss these interesting extensions here, the reader is referred to [5, 20, 18] for details.

The main positive message of this section is that, when the actual or likely structure of a population of players is taken into account, the results of evolutionary game theory tend to favor cooperation and socially valuable behavior with respect to the unstructured mixing population. We conclude by referring the interested reader to the recent review [24] for a much more detailed description of the state of the art on evolutionary games on graphs.

7. Conclusions and outlook

In this article I have first presented an overview of the fundamentals of game theory, followed by an introduction to its evolutionary interpretation. Evolutionary game theory has marked an important progress by doing away with strict instrumental rationality and introducing simple learning and imitation mechanisms. In spite of its success and its mathematical elegance, this stance is only partially satisfactory however. The standard problem of equilibrium selection is mitigated but it doesn't disappear, and the lack of rationality does not agree with observed individual's decision-making criteria in human societies. Moreover, EGT only applies to populations of players, not to situations in which there are few players and small groups. This makes it eminently suited for the study of the emergence of social behaviors such as cooperation and norm establishment. Many causes have been invoked for justifying the emergence of cooperation and reciprocity among human and even animal populations. Here we have presented one explanation, which is not mutually exclusive with other propositions, that only uses the observable fact that such populations possess a recognizable network structure. Using graph theory and numerical simulations it has been possible to show that the mere fact of acting locally in a widely inhomogeneous network environment makes it easier for the individuals to achieve socially useful global results. Other sources of diversity in the populations are also surely important and should be studied along with the fact that societies are dynamical entities.

Acknowledgements. This research work is funded by the Swiss National Science Foundation under grant number 200020-119719. I gratefully acknowledge this financial support. I also wish to express my gratitude to C. Bach, J. Peña, E. Pestelacci, and H. Volken for reading the manuscript and making useful suggestions.

References

[1] AMARAL L. A. N., SCALA A., BARTHÉLEMY M. AND STANLEY H. E., *Classes of small-world networks*, Proc. Natl. Acad. Sci. USA **97** (2000), 11149–11152.

[2] AXELROD, *The evolution of cooperation*, Basic Books, Inc., New-York 1984.

[3] BINMORE K., *Does game theory work? The bargaining challenge*, MIT Press, Cambridge, MA 2007.

[4] BOLLOBÁS B., *Modern graph theory*, Springer, Berlin, Heidelberg, New York 1998.

[5] EGUÍLUZ V. M., ZIMMERMANN M. G., CELA-CONDE C. J. AND SAN MIGUEL M., *Cooperation and the emergence of role differentiation in the dynamics of social networks*, American J. of Sociology **110** (4) (2005), 977–1008.

[6] GINTIS H., BOWLES S., BOYD R. T. AND FEHR E., *Moral sentiments and material interests*, MIT Press, Cambridge, MA 2005.

[7] HARGREAVES S. P. AND VARUFAKIS Y., *Game theory: a critical text*, Routledge, London, New York 2004.

[8] HOFBAUER J. AND SIGMUND K., *Evolutionary games and population dynamics*, Cambridge University Press, Cambridge 1998.

[9] MAYNARD-SMITH J. AND G. R. PRICE, *The logic of animal conflict*, Nature **246** (1973), 15–18.

[10] LUTHI L., PESTELACCI E. AND TOMASSINI M., *Cooperation and community structure in social networks*, Physica A **387** (2008), 955–966.

[11] MAILATH G. J., *Do people play Nash equilibrium? Lessons from evolutionary game theory*, J. of Economic Literature **36** (1998), 1347–1374.

[12] MAYNARD SMITH J., *Evolution and the theory of games*, Cambridge University Press, Cambridge 1982.

[13] MYERSON R. B., *Game theory: analysis of conflict*, Harvard University Press, Cambridge, MA 1991.

[14] NEWMAN M. E. J., *The structure and function of complex networks*, SIAM Review **45** (2003), 167–256.

[15] NOWAK M. A. AND MAY R. M., *Evolutionary games and spatial chaos*, Nature **359** (1992), 826–829.

[16] NOWAK M. A., SASAKI A., TAYLOR C. AND FUDENBERG D., *Emergence of cooperation and evolutionary stability in finite populations*, Nature **428** (2004), 646–650.

[17] NOWAK M. A. AND SIGMUND K., *Games on grids*, in: "The geometry of ecological interactions: simplifying spatial complexity" (Eds. Dieckmann U., Law V. and Metz J. A. J.), Cambridge University Press, Cambridge, UK 2000, 135–150.

[18] PESTELACCI E., TOMASSINI M. AND LUTHI L., *Evolution of cooperation and coordination in a dynamically networked society*, J. of Biological Theory, to appear.

[19] POUNDSTONE W., *The prisoner's dilemma*, Doubleday, New York 1992.

[20] SANTOS F. C., PACHECO J. M. AND LENAERTS T., *Cooperation prevails when individuals adjust their social ties*, PLOS Comp. Biol. **2** (2006), 1284–1291.

[21] SAVAGE L. J., *The foundations of statistics*, J. Wiley and Sons, New York 1954.

[22] SCHELLING T., *The strategy of conflict*, Harvard University Press, Cambridge, MA 1960.

[23] SKYRMS B., *The stag hunt and the evolution of social structure*, Cambridge University Press, Cambridge, UK 2004.

[24] SZABÓ G. AND FÁTH G., *Evolutionary games on graphs*, Physics Reports **446** (2007), 97–216.

[25] TAYLOR P. AND JONKER L., *Evolutionary stable strategies and game dynamics*, Mathematical Biosciences **40** (1978), 145–156.

[26] VEGA-REDONDO F., *Economics and the theory of games*, Cambridge University Press, Cambridge, UK 2003.

[27] VON NEUMANN J. AND MORGENSTERN O., *Theory of games and economic behavior*, Princeton University Press, Princeton NJ 1944.

[28] WATTS D. J. AND STROGATZ S. H., *Collective dynamics of "small-world" networks*, Nature **393** (1998), 440–442.

[29] WEIBULL J. W., *Evolutionary game theory*, MIT Press, Boston, MA 1995.

AMS Subject Classification: 91A10, 91A22, 91D25, 91D30.

Marco TOMASSINI, HEC, Information Systems Department, University of Lausanne, 1015 Lausanne, SWITZERLAND
e-mail: `marco.tomassini@unil.ch`