Evolutionary Games: What is the algorist's perspective?*

Spyros Kontogiannis[1,2] and Paul Spirakis[1]

[1] Computer Technology Institute, Riga Feraiou 61, 26221 Patras, Greece. {kontog,spirakis}@cti.gr
[2] Department of Computer Science, University of Ioannina, 45110 Ioannina, Greece.

Abstract. Evolutionary Game Theory is the study of strategic interactions among large populations of agents who base their decisions on simple, myopic rules. A major goal of the theory is to determine broad classes of decision procedures which both provide plausible descriptions of selfish behaviour and include appealing forms of aggregate behaviour. For example, properties such as the correlation between strategies' growth rates and payoffs, the connection between stationary states and Nash equilibria and global guarantees of convergence to equilibrium, are widely studied in the literature. In this paper we discuss some computational aspects of the theory, which we prefer to view more as *Game Theoretic Aspects of Evolution* than *Evolutionary Game Theory*, since the term "evolution" actually refers to strategic adaptation of individuals' behaviour through a dynamic process and not the traditional evolution of populations. We consider this dynamic process as a *self-organization procedure*, which under certain conditions leads to some kind of stability and assures robustness against invasion.

1 Introduction

Classical game theory deals with a rational individual who is engaged in a given interaction (a "game") with other individuals (its co-players) and has to adopt a strategy (for selecting among a set of allowable actions) that maximizes its own payoff. Of course each player's payoff is dependent on the other players' strategies for choosing their own actions.

Evolutionary game theory deals with the entire (typically large) population of players, where all players are programmed to adopt some strategy. Strategies with high payoff (given the current state of the population) are expected to spread within the population (by learning, copying successful strategies, inheriting strategies, or even by infection). The frequencies of the strategies in the population thus change according to their payoffs, which in turn depend on all the players' strategies and thus their frequencies in the population. The subject of evolutionary game theory is exactly the study of the dynamics of this feedback loop. A very good presentation of evolutionary game dynamics is in [3]. For a more thorough study the reader is referred to [1].

* This work was partially supported by the EU within the Future and Emerging Technologies Programme under contract IST-2001-331116 (FLAGS) and within the 6th Framework Programme under contract 001907 (DELIS).

Numerous paradigms for modeling individual choice in a large population have been proposed in the literature. For example, if each agent chooses its own strategy so as to optimize its own payoff ("one against all others" scenario) given the current population state (ie, other agents' strategies), then the aggregate behaviour is described by the best-response dynamics [4]. If each time an arbitrary user changes its strategy for any other strategy of a strictly better (but not necessarily the best) payoff, then the aggregate behaviour is described by the better-response dynamics or Nash dynamics [6]. In case that pairs of players are chosen at random and these players engage in a bimatrix game ("one against one" scenario) whose payoff matrix determines according to some rule the gains of the strategies adopted by these two players, then we refer to imitation dynamics, the most popular version of which is the replicator dynamics[8].

The proposed dynamical systems for describing evolutionary games, despite their appealing and highly intuitive definitions, have some strong weaknesses. For example, the best/better-response dynamics admit multiple solution trajectories from a single initial point, whose terminal (rest) points vary significantly in their aggregate performance. On the other hand, the replicator dynamics admit trajectories, whose rest points are not necessarily Nash equilibria (eg, when all the users of a population choose exactly the same strategy, even if this strategy is strictly dominated, there is nothing to imitate – this is a general weakness of the imitation dynamics).

A typical way out of such situations is to introduce small amounts of noise to the underlying models. For example, we may add (small) payoff perturbations to the optimization dynamic models, or we may add occasional arbitrary behaviour to the model of replication. Of course, such modifications create new difficulties: for the optimization models the rest points may now be slightly away from the Nash equilibria, while the replication models may admit some oscillation phenomena.

Since evolutionary game theory is a dynamical process (hopefully) ending up in some equilibrium which may also demonstrate robustness against invasion or infection (stability), we consider it as a self-organization process in a very large population of entities that adopt some strategies representing either computer programs trying to prevail in the market, or competing network protocols, or viruses trying to spread all over the Internet. Thus the prime concern of an algorist is to determine the principal rules defining the (global or local) convergence to stable states, to propose computationally efficient algorithms for constructing such stable states, to be able to compare the trajectories of two phenomenically different laws of motion, or even to describe how the underlying infrastructure (eg, the network infrastructure in which a virus might spread) is involved in the evolution of a population. We comment here that our perspective has to do with strategic evolution, ie, the evolution in time of the vector of the strategies' frequencies in the whole population, and not with traditional evolution (under the biological viewpoint) where the sub-populations actually evolve in size.

2 Definitions and Notation

For any $n \in \mathbb{N}$, let $[n] \equiv \{1, 2, \ldots, n\}$, while $\mathcal{P}_n \equiv \left\{\mathbf{z} \in [0,1]^n : \sum_{j\in[n]} z_j = 1\right\}$ is the simplex of probability n-vectors. Assume having a set $N \equiv [n]$ of selfish, non-cooperative players, and n action sets $(S_i \equiv \{s_{i,1}, \ldots, s_{i,m_i}\})_{i\in N}$ for them, where $m_i = |S_i|$. Each player $i \in N$ may adopt either a pure strategy $s_{i,j} \in S_i$ (ie, a fixed allowable action), or a mixed strategy $\mathbf{p_i} \in \mathcal{P}_{m_i} \equiv \left\{(z(s_{i,j}))_{j\in[m_i]} \in [0,1]^{m_i} : \sum_{j=1}^{m_i} z(s_{i,j}) = 1\right\}$, (ie, a probability distribution over its own action set). For simplicity of notation, let $\forall i \in N,\ \forall \mathbf{p_i} \in \mathcal{P}_{m_i},\ \forall j \in [m_i],\ p_i(s_{i,j}) = p_i(j)$.

A set $\mathbf{p} = (\mathbf{p_1}, \ldots, \mathbf{p_n}) \in \times_{i\in N}\mathcal{P}_{m_i} \equiv \mathcal{P}$ of (mixed in general) strategies for all the users of the game is called a *(mixed) strategies profile*. For the special case where $\mathbf{p}$ corresponds to a set $\mathbf{s} \in \times_{i\in N} S_i \equiv S$ of pure strategies for all the players, we have a *pure strategies profile* or a *configuration*, which is also represented by $\mathbf{s}$. $\forall i \in N$, we denote by $\mathbf{p}^{-i} \equiv (\mathbf{p_1}, \ldots, \mathbf{p_{i-1}}, \cdot, \mathbf{p_{i+1}}, \ldots, \mathbf{p_n})$ the mixed strategies profile of all the players but for player i, while we define the operation $\mathbf{p}^{-i} \oplus \mathbf{z} = (\mathbf{p_1}, \ldots, \mathbf{p_{i-1}}, \mathbf{z}, \mathbf{p_{i+1}}, \ldots, \mathbf{p_n})$, $\forall \mathbf{z} \in \mathcal{P}_{m_i}$, $\forall \mathbf{p}^{-i} \in \times_{j\neq i}\mathcal{P}_j \equiv \mathcal{P}^{-i}$.By abusing notation a little bit, we will occasionally allow a pure strategy $s_{i,j} \in S_i$ to be combined by $\oplus$ with $\mathbf{p}^{-i}$, although we should use instead the corresponding (mixed) strategy $\mathbf{z} = \mathbf{e_j} \in \mathcal{P}_{m_i}$, where $\mathbf{e_j}$ is the vector with 1 in its j^{th} position and 0 in every other position.

Consider now that each player $i \in N$ has its own payoff function $U_i : S \mapsto \mathbb{R}$ (this function can also be represented as a 2-dimensional matrix whose rows are labeled by the actions of S_i and columns are labeled by combined actions of the remaining users, ie, elements of $\times_{r\neq i} S_r \equiv S^{-r}$). We extend the utility functions to the domain $\mathcal{P}$ as follows:

$$\forall i \in N,\ \forall \mathbf{p} \in \mathcal{P},\ U_i(\mathbf{p}) = \sum_{\mathbf{s}\in S} P(\mathbf{s}, \mathbf{p}) \cdot U_i(\mathbf{s})$$

where $P(\mathbf{s}, \mathbf{p}) \equiv \prod_{i\in N} p_i(s_i)$ is the occurrence probability of configuration $\mathbf{s}$ under the mixed strategies profile $\mathbf{p}$.

Nash Equilibrium (NE) [5]. A mixed strategies profile $\mathbf{p} \in \mathcal{P}$ is a NE iff each player's strategy is a best-response to the other players' strategies. That is, $\forall i \in N$,

$$\mathbf{p_i} \in \arg\max_{\mathbf{z}\in\mathcal{P}_i} \left\{U_i(\mathbf{p}^{-i} \oplus \mathbf{z})\right\} \equiv BR_i(\mathbf{p}^{-i}) . \quad (1)$$

$\mathbf{p}$ is a *strict* NE iff $\forall i \in N,\ \{\mathbf{p_i}\} = BR_i(\mathbf{p}^{-i})$.

The Replicator Dynamics. Consider having a (very large) population of individuals, each having exactly the same set $N = [n]$ of possible *types*. Let $\mathbf{x} \in \mathcal{P}_n$ be the *population state*, ie, the vector of proportions (or frequencies) of the n different types in the whole population[1]. Consider the dynamical system where in each step two individuals of the population are chosen randomly and then engage in a symmetric bimatrix game whose (common) payoff function is described by the $n \times n$ matrix U. Then, $(U\mathbf{x})_i$ is the expected payoff of a type-i user that is involved in a new game, while $\mathbf{x}^T U\mathbf{x}$ is the average (expected) payoff of a random user (ie, the average payoff in the system) involved in a new game, wrt the current population state $\mathbf{x}$.

Suppose that the type-frequencies vector $\mathbf{x}$ is a vector of differentiable functions of time t (this requires that the population is infinitely large, or that x_i's express the expected frequencies in an ensemble of populations) and postulate a law of motion $\mathbf{x}(t)$. Suppose that for any type of strategies, its proportion in the next generation is related to the proportion of the same type in the current generation, according to the following equation, called the replicator equation [8]:

$$\forall i \in N,\ \dot{x}_i = x_i \left[(U\mathbf{x})_i - \mathbf{x}^T U\mathbf{x}\right] . \quad (2)$$

Observe that this frequency-dependent fitness rule introduces a strategic aspect to evolution: more successful strategies (ie, those having an expected payoff strictly larger than the average expected payoff) increase their proportions in the population, while those with less than average expected payoff loose some of their proportions. Recall that the total change in frequencies is exactly 0, ie, although the frequencies of the types change, the population size remains the same.

Assume now that the strategies possibly adopted by the individuals of the population represent m different mixed strategies on the set of n different types,$(\mathbf{p}(i) \in \mathcal{P}_n)_{i\in[m]}$. The expected payoff of a type $\mathbf{p}(i)$ against another type $\mathbf{p}(j)$ is $(\mathbf{p}(i))^T U\mathbf{p}(j)$, while

[1] Recall that $\mathbf{x}$ corresponds to (more than one) **pure** strategies profiles for the individuals.

for a given frequencies vector $\mathbf{x} \in \mathcal{P}_m$, the average (expected) payoff within the population is $(\mathbf{p}(\mathbf{x}))^T U \mathbf{p}(\mathbf{x})$, where $\mathbf{p}(\mathbf{x}) \equiv \sum_{i\in[m]} x_i \mathbf{p}(i)$.The analogue of (2) for the law of motion of mixed strategies, can be written as

$$\forall i \in [m],\ \dot{x}_i = x_i \left[(\mathbf{p}(i) - \mathbf{p}(\mathbf{x}))^T U \mathbf{p}(\mathbf{x}) \right] . \quad (3)$$

Evolutionary Stable Strategy (ESS) [7]. A (mixed in general) strategy $\mathbf{p} \in \mathcal{P}_n$ is said to be *evolutionary stable* (ESS) if, for any other strategy $\mathbf{q} \in \mathcal{P}_n \setminus \{\mathbf{p}\}$, the induced replicator equation describing the dynamics of the population consisting of **these two types only** (the proportion of *residents* using $\mathbf{p}$ is $1-r$ and the proportion of *invaders* using $\mathbf{q}$ is r, for some $0 < r \le \varepsilon(\mathbf{q}) \ll 1$) leads to the elimination of the invaders, as long as their initial frequency is sufficiently small. That is, $\mathbf{p}$ is an ESS iff $\forall \mathbf{q} \in \mathcal{P} \setminus \{\mathbf{p}\}$ the following two conditions hold:

(I) *Equilibrium:* $\mathbf{q}^T U \mathbf{p} \le \mathbf{p}^T U \mathbf{p}$
(II) *Stability:* $\mathbf{q}^T U \mathbf{p} = \mathbf{p}^T U \mathbf{p} \Rightarrow \mathbf{q}^T U \mathbf{q} < \mathbf{p}^T U \mathbf{q}$

Condition (I) demands that $\mathbf{p}$ is a Nash equilibrium for the bimatrix game determining the law of motion (no invader can do better than the resident against the resident), while condition (II) states that, in case that an invader does equally well with the resident against the resident, then the resident must be strictly better than the invader against the invader. The following theorem provides a simple (yet **inefficient**, due to the possibly too many (mixed) strategies for the individuals) test whether a specific strategy p is an ESS for a replicator equation:

Theorem 1 (Hofbauer, Schuster and Sigmund [2]) *The strategy* $\mathbf{p}$ *is an ESS iff* $\prod_{i\in[m]} x_i^{p(i)}$ *is a strict local Lyapunov function*[2] *for the replicator equation, or equivalently, iff*

$$\forall \mathbf{q} \ne \mathbf{p},\ \mathbf{p}^T U \mathbf{q} > \mathbf{q}^T U \mathbf{q} \quad (4)$$

in some neighbourhood of $\mathbf{p}$.

3 Computational aspects of game-theoretic evolution

Equilibrium Selection. The most popular notion of stability in game theory is the Nash equilibrium (NE) [5]. Although the notion of NE is quite natural, if we restrict ourselves to pure strategies, we may not be able to reach a pure Nash equilibrium (PNE) at all. For example, in the Matching Pennies game shown in figure 1 there is no PNE at all. More importantly, even if we have PNE in a game, we may face a situation where multiple PNE exist, of different payoffs for the players and quite different aggregate performances. The problem then is for a rational player, how to decide which of the several NE is the "right" one to settle upon. To this direction, numerous refinements of the space of (P)NE have been proposed in the literature. In fact, there are so many refinements, that practically every NE may be shown to be the "right choice" of a proper refinement! There is a strong hope that evolutionary game theory will assist this kind of choices. The reason is the systematic way by which the evolution is modeled, as a kind of a reasonable game between an individual and some other individuals (or even against all other individuals). The point is to be able to construct computationally efficient algorithms for quantifying the values of all the NE, or even solving the corresponding optimization problem in the space of NE wrt a given objective function. Recall that for an evolutionary game the reachable NE are a subset of the rest points for the dynamics of this game, which in turn are the endpoints of (continuous in the limit) trajectories starting from some strategy initially adopted by (ie, prevailing among) the users. This should not be confused with the (computationally hard) task of determining the worst/best NE in a traditional game.

	Heads	Tails
Heads	0,1	1, 0
Tails	1,0	0, 1

Fig. 1. The Matching Pennies game: The row player bets on different outcomes showing up, while the column player bets on the same outcomes showing up. For this game there is no PNE.

Additionally it would be extremely interesting to devise an algorithm for detecting the existence of an ESS in an evolutionary game, that only takes as input the $n \times n$ matrix U determining the payoffs of the game and responds with YES/NO, or even better, constructs an ESS in case of existence. Recall that an ESS is a NE of the corresponding bimatrix game with an additional stability property, which does not necessarily hold in a NE. Indeed, one can construct simple examples where a small payoff matrix admits no ESS at all (although the game has at least one NE). A similar computational problem is the description of a computationally efficient evolutionary process (ie, the proper rule for evolution) that will lead as fast as possible to a rest point which is also a NE for the corresponding traditional game, or even an ESS for the evolutionary dynamics.

Bounded vs. Unbounded Rationality. A central assumption of the traditional game theory is that each player adopts a kind of *selfish behaviour*, exploiting the complete knowledge of other players' decisions. For example, in order to be able to assign a cardinal utility function to individual players, one typically as-

[2] A function $F(\mathbf{x})$ is a *Lyapunov function* if $\dot{F}(\mathbf{x}) \ge 0$, for all $\mathbf{x}$. It is a *strict Lyapunov function*, if additionally equality holds only when $\mathbf{x}$ is a rest point.

sumes that each player has a well-defined, consistent set of preferences over the set of "lotteries" over the outcomes which may result from individual choice. Since the number of different lotteries over outcomes is uncountably many, this requires each player to have a well-defined, consistent set of uncountably many preferences, which is typically considered to be infeasible.

In many cases though this is not a feasible situation, either because it is not computationally efficient to handle such an amount of information, or because each player only knows the actual strategies of those players in its own neighborhood. Despite the fact that traditional game theory does not deal with such situations, there is a strong hope that evolutionary game theory will manage eventually to successfully describe and predict the behaviour of such players since it is better equipped to handle these *weaker rationality assumptions* and yet causing the same effects on players' behaviours in the long run as complete-knowledge traditional games. A typical example of such a scenario is when we wish to cut off the iteratively dominated strategies of the players. Depending on the model of evolution, we are able in some cases to assure that all these strategies, that a rational individual should never play in a traditional game, will eventually vanish in the evolutionary version of the game.

The computational point of view demands again to design algorithms dealing with bounded rationality and achieving in polynomial time (possibly good approximations of) the same outcome as the one expected by the corresponding traditional game.

Trajectory Prediction. Due to its nature, evolutionary game theory explicitly models the dynamics present in interactions among individuals in a population. One might try to capture the dynamics of the decision-making process in traditional game theory by modeling the game in its extensive (rather than its strategic) form. However, for most games of reasonable complexity, the extensive form of the game quickly becomes unmanageable. Moreover, in the extensive form of a game, traditional game theory represents an individual's strategy as a specification of what choice that individual would make at each information set in the game. A selection of strategy then corresponds to a selection, prior to game play, of what that individual will do at any possible stage of the game. This representation of strategy selection clearly presupposes hyperrational players and fails to represent the process by which a player observes its opponents' behaviours, learns from these observations and makes a best/better response, replication, imitation, etc choice to what it has learned so far.

We would like to be able to decide in polynomial time whether two phenomenically different evolutionary dynamics actually produce the same trajectories, up to some kind of homeomorphism. This would enable a classification scheme of the evolution schemes into broader categories according to their orbital characteristics. Such kind of classification for replicator dynamics exist only when there are 2 or 3 distinct types of individuals in the population.

Imposing structural properties of a game into the evolutionary dynamics. Our last question deals with some new ways of interaction in the evolutionary dynamics of a game, that will also depict the special structure of the corresponding traditional game. For example, when a virus spreads in a network, the architecture of the network itself and the starting points of the virus in it should affect somehow the success of the virus. The proposed game theoretic models of evolution proposed so far in the literature, mainly focus on the case where the individuals in a population collide with each other in a random fashion. Ie, the underlying "interaction" infrastructure is represented by a clique. What if this is not the case, and we have instead some special graph representing the interactions? We need new evolutionary models to capture such cases, that will somehow encode the structure of this graph in the dynamics, via elementary properties (eg, the connectivity or the expansion of the graph).

References

1. Cressman R. *Evolutionary dynamics and extensive form games*. MIT Press, 2003.
2. Hofbauer J., Schuster P., Sigmund K. A note on evolutionary stable strategies and game dynamics. *Journal of Theoretical Biology*, 81:609–612, 1979.
3. Hofbauer J., Sigmund K. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, 2003.
4. Matsui A., Gilboa I. Social stability and equilibrium. *Econometrica*, 59:859–867, 1991.
5. Nash J. F. Noncooperative games. *Annals of Mathematics*, 54:289–295, 1951.
6. Rosenthal R.W. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
7. Smyth M. J., Price G. The logic of animal conflict. *Nature*, 246:15–18, 1973.
8. Taylor P. D., Jonker L. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156, 1978.