International Game Theory Review
© World Scientific Publishing Company

EVOLUTIONARY GAMES AND LOCAL DYNAMICS

PHILIPPE UYTTENDAELE
FRANK THUIJSMAN
Department of Knowledge Engineering
Maastricht University[*]

Received (Day Month Year)
Revised (Day Month Year)

In this article we examine several options for modeling local interactions within the framework of evolutionary game theory. Several examples show that there is a major difference between population dynamics using local dynamics versus global dynamics. Moreover, different modeling choices may lead to very diverse results.

Keywords: evolutionary games; replicator dynamics; local interactions.

Subject Classification: 91A22.

1. Introduction

Evolutionary games have been introduced by Maynard Smith and Price [1973] in the early 70's with the aim to study the stability of populations in time. In the general model, a population consists of finitely many different types that interact randomly with each other, where each interaction leads to fitness payoffs for each of the types involved. Consequently, the population distribution over the different types is changing in time. When assuming that the fraction of a type changes in a Darwinian way proportionally to its current fraction and proportionally to the difference in fitness with average population members, then the population development can be represented by the so-called replicator dynamics, introduced by Taylor and Jonker [1978].

Although many other population dynamics have been studied for evolutionary games (see Hofbauer and Sigmund [1998] or Sandholm [2011] for a review of all kinds of dynamics), we would like to take the replicator dynamics as our starting point in view of its relation to the Darwinian theory. However, we want to switch from a global to a local perspective, because one of the underlying assumptions in the model of Maynard Smith and Price [1973] and in that of Taylor and Jonker

---

[*] For both authors: Department of Knowledge Engineering, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. Correspondence: f.thuijsman@maastrichtuniversity.nl

[1978] is that population members interact with random individuals from the entire population. This implies that it is assumed that every population member is facing the same population distribution. In this study we want to examine the effect of spatially distributed population members that interact at random in their spatial neighborhood. Population developments that are locally driven in a Darwinian way, seem to be more realistic to us and closer to biological observations. Although analytically these processes become much more complex, current computing power allows for running simulations that can help us understand where these processes may lead to. We think that our approach will also allow modeling of biological phenomena that can not be captured by the global model.

As expected, and illustrated by our examples, local dynamics and global dynamics may lead to different results. A type that would go extinct under global dynamics may survive because of its local conditions: it may do well in symbiosis with another type to compensate its low individual fitness. Similar differences (local versus global interaction) have been observed in biology (cf. Buttel *et al.* [2002]; Kerr *et al.* [2002]; Killingback and Doebeli [1996]; Kirkup and Riley [2004]).

Others have also examined local interactions, especially in theoretical biology. These studies are mainly based on cellular automata (CA) and often relate to models of learning for stylized examples like the Prisoner's Dilemma. A CA model consists of a grid of cells, which changes based on rules given by a finite number of states. For each cell, a neighborhood is defined for interactions. The initial state can be chosen arbitrarily or randomly but based on this starting point, following generations are computed. These new generations are created according to some fixed rule that determines the new type of each cell in terms of its current type and the types of the cells in its neighborhood. The way the generations are computed can be of several orders of complexity, and it can be probabilistic or deterministic. Introduced in the fifties by Ulam [1952] and by Von Neumann [1966], CA models have been deeply analyzed. Wolfram [2002] even suggested that these would lead to a new concept of sciences and that these could be used in many fields of expertise.

The outline of this paper is as follows: In section 2 we introduce the formal model and the different types of local interaction that we compare in this paper. In section 3 we report the experimental results. In section 4 we relate the models to a biological experiment. Section 5 concludes with remarks on further research and briefly reports on instances of local models that exhibit predator-prey behavior, symbiosis and periodically stable population structures.

2. The Model

Evolutionary games are determined by a fitness matrix A, based on which a population distribution will change over time. The population distribution at a given time t is represented as a vector $\mathbf{p}(t) = (p_1(t), p_2(t), ..., p_n(t))$, where $p_i(t) > 0$ for all i (all types are present) and $\sum_{i=1}^{n} p_i(t) = 1$. The fitness matrix A is an $n \times n$ matrix in which the number a_{ij} in entry (i, j) gives the fitness of an individual of type i

when interacting with an individual of type j. This means that given a population $\mathbf{p}$ the average fitness of type i will be $\mathbf{e}_i A\mathbf{p}^\intercal$ and the average fitness of an individual in the population will be $\mathbf{p}A\mathbf{p}^\intercal$.

An evolutionarily stable strategy (ESS) is a population distribution x such that for all other population distributions y the following two conditions apply:

a. $xAx^\intercal \geq yAx^\intercal$
b. if $y \neq x$ and $yAx^\intercal = xAx^\intercal$, then $xAy^\intercal > yAy^\intercal$.

Here, the first condition is the Nash equilibrium condition saying that x should be a best reply to x, while the second condition prevents the population from drifting away from x.
The replicator dynamics, introduced by Taylor and Jonker [1978], is based on two simple observations. It describes, in a Darwinian way, how a population will evolve by changing the population fraction of any type proportionally to its performance compared to the average and proportionally to its current size. The replicator dynamics is given by the following system:

$$\dot{p_i} = p_i(\mathbf{e}_i A\mathbf{p}^\intercal - \mathbf{p}A\mathbf{p}^\intercal) \quad \text{for } i = 1, 2, \ldots, n.$$

A limit point x of the dynamics is called asymptotically stable if for any starting point x_0 in a sufficiently small open neighborhood of x the dynamical process converges to x. An ESS is always asymptotically stable, but not the other way around. Unfortunately these calculations of average fitness are only valid in a global interaction framework. In other words, any individual interacts with the rest of population with the same probability distribution over the types. The logical next step is to introduce spatial information and to let competition take place at a local scale.

We represent the spatial population distribution by a field consisting of hexagons. To avoid border conditions and to make each cell have the same number of neighbors, the field is represented as a torus. This means that the top and the bottom of the field are connected, just like the left and the right. Cells that are equidistant from a specifically chosen cell form rings. For example, all cells that are exactly 2 cells away from a given cell, will form a ring around the selected cell (see Figure 1), that we call ring 2. All cells within a distance of k from a given selected cell, form the k-neighborhood of the selected cell.

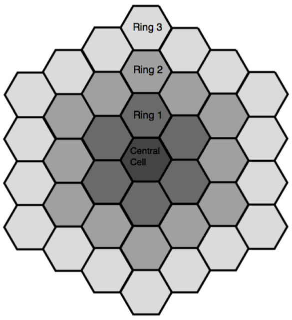

Fig. 1: Rings in a hexagonal field.

With this initial set up, we designed a series of scenarios. These scenarios are ordered from macroscopic to microscopic to analyze the impact of taking the localization more and more into account.

In our scenario studies we simulate how local interactions are changing the population distribution over many generations. In order to compare these scenarios we want to initialize each study by the exact same population distribution over the types. This means that for a given population distribution we will assign fixed numbers of cells to each of the different types, such that in each scenario the initial distribution is represented in the field with exactly the same numbers. For each scenario 10 simulations have been performed for each of 5 different initial distributions.

Our benchmark is the classical process driven by the replicator dynamics. There is no notion of field, and therefore it does not take into account the location of individuals in the population.

Scenario 1 is a process in which the population is placed on a (hexagonal) field. As such there is a finite number of individuals, one per cell in the field. Once placed on the field, the following computations are done simultaneously for all individuals. Each individual interacts with every other individual and all payoffs are recorded per individual. Then the total payoff per type is computed. Next, simultaneously for all cells, the type that has the highest total payoff will take over the cell. As such, a new generation of the field is made.

Here we remark that, in events where there are two or more types with maximum payoff involved in the comparing step, no type change will take place. This is also assumed in the scenarios below.

Scenario 2 is similar to scenario 1 except that, rather than changing all cells at the same time, only one cell is selected to change and next calculations start all over. This makes the process more gradual.

Scenario 3 is in the same line as scenario 1. However instead of interacting with the entire field, each cell is limited to an interaction with each neighbor in ring 1. Then, within the 1-neighborhood the total payoff per type is computed and the central cell will adopt the type with the highest total payoff. This is done simultaneously for all cells in the field. Thus the entire field is updated simultaneously but with local interactions.

Scenario 4 is similar to scenario 3 except that only one cell at the time is selected for changing based on the total payoff per type in its 1-neighborhood. This makes the population change more gradual in time.

Scenario 5 is a scenario in which two adjacent cells are taken randomly, and enrolled in a pairwise competition. Each of the two cells interacts with its 1-neighborhood, and the total payoffs of each of these two cells are compared. The highest total payoff of these two will decide the type of both cells.

Scenario 6 is similar to scenario 5 except that the interaction is purely pairwise. This means that the two selected cells play against each other and the one that has the highest immediate payoff of these will decide the type of both cells.

Each simulation ran until a steady state was reached. Here we wish to remark that scenarios 1 and 3 are completely deterministic. For scenario 2 and 4 the updating of the field depends on the selection of one cell only. Whether or not this cell changes type is completely determined by its neighborhood and can be checked easily. For the two remaining scenarios we could conlude that the field had stabilized because either only one type remained or two types remained that could not affect each other. Of course, some scenarios can imply that the field keeps changing forever, an example of which is provided in section 5. However, such did not happen in the comparative study based on these scenarios that is described in the next section.

3. A Case Study for Different Scenarios

For comparing the 6 different scenarios we used the following fitness matrix:

$$M = \begin{pmatrix} 0 & 3 & -1 \\ -3 & 0 & 1 \\ -1 & 1 & 0 \end{pmatrix}$$

For this game there is only one ESS, namely $p = (1\,,0\,,0)$. For matrix M a flow diagram of the replicator dynamics, our benchmark model, can be seen in Figure 2. In this figure, any point in the triangle corresponds to a population distribution over the three types. Processes that start with a population distribution in the top half (any population that has members of all types and more than 50% members of type 3) will cycle around the rest point $p = (\frac{1}{5}\,,\frac{1}{5}\,,\frac{3}{5})$. Any population that has members of all types and less than 50% members of type 3, will converge to $p = (1\,,0\,,0)$ and a population that has exactly 50% members of type 3 will converge to $p = (\frac{1}{2}\,,\frac{1}{2}\,,0)$, which is an unstable rest point. A population that does not contain any member of type 2, will either converge to $p = (1\,,0\,,0)$ if type 1 has more members in the population, or to $p = (0\,,0\,,1)$ if it is type 3 that has more members in the population. A population that does not contain any member of type 1, will converge to $p = (0\,,\frac{1}{2}\,,\frac{1}{2})$.

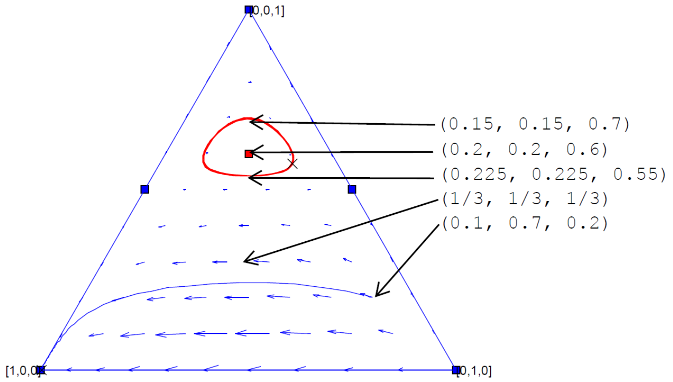

Fig. 2: Flow diagram of the replicator dynamics with the payoff matrix M

The following table summarizes the results of the experiments in which, for each scenario, 10 simulations have been performed for each of 5 different inital distributions.

Table 1: Stabilization for different starting points depending on the scenario.

		Initial Distribution				
		$(0.1\,,0.7\,,0.2)$	$(\frac{1}{3}\,,\frac{1}{3}\,,\frac{1}{3})$	$(0.225\,,0.225\,,0.55)$	$(0.2\,,0.2\,,0.6)$	$(0.15\,,0.15\,,0.7)$
Benchmark		all type 1	all type 1	stable cycle	stable point	stable cycle
Scenario	1	all type 1	all type 1	all type 1	stable initial field	all type 2
	2	all type 1	all type 1	all type 1	stable initial field	$(0\,,1-q\,,q)$
	3	extinction of type 2 and coexistence of types 1 and 3				
	4					
	5	extinction of types 2 and 3; only type 1 survives				
	6	extinction of type 2 and coexistence of types 1 and 3				

When analyzing Table 1 scenario by scenario, we can see that scenario 1 moves to extremes. The reason is that, because all cells of the field are updated simultaneously, based upon the payoff information for the entire field, the same best type is adopted by all cells at the same time, leading to a field containing only one type. As $(0.2\,,0.2\,,0.6)$ is a rest point, all types are equally good and the field does not change at all. Also note that for initial distribution $(0.15\,,0.15\,,0.7)$ type 2 is the surviving type, because at this starting position type 2 is doing best as can be seen in Figure 2.

For scenario 2, the sum of all payoffs per type determines the future type of one cell at the time. For the first three initial distributions in the beginning of the simulation the proportion of type 1 is growing while those for types 2 and 3 are shrinking. Once below the 50% line for type 3, this process goes even faster and leads to extinction of types 2 and 3. For the last initial distribution types 1 and 3 are first loosing from type 2 and then the process leads to extinction of type 1 and stabilizes when this happens.

We can see in Table 1 that scenarios 3, 4, and 6 each lead to a coexistence of types 1 and 3. There is a lot of variance in the stable configurations, but Figure 3 gives an impression of how the process may end.

For scenario 5 only type 1 remains. This is explained by the fitness advantage that it gets by interacting with type 2 leading to a strategic advantage in number to win in pairwise interactions.

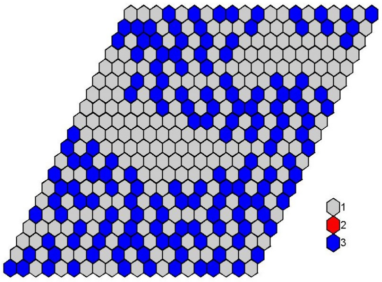

Fig. 3: Coexistence of types 1 and 3 for scenario 4, with initial distribution $(\frac{1}{5}, \frac{1}{5}, \frac{3}{5})$.

4. An Application in Biology

In order to compare our model to other studies, we analyze the observations reported by Kerr *et al.* [2002] in an article titled "Local dispersal promotes biodiversity in a real life game of rock-paper-scissors".

Based on experiments with populations of bacteria, they examine the evolution of three different types of bacteria when put in the same environment. These three types are: a toxic $\mathcal{C}$olicinogenic type ($\mathcal{C}$), a type $\mathcal{R}$ that is $\mathcal{R}$esistant to the toxicity of $\mathcal{C}$, and a type $\mathcal{S}$ that is $\mathcal{S}$ensitive to the poison. The idea is that $\mathcal{C}$ bacteria release a poison that kills any $\mathcal{S}$ bacteria in its immediate neighborhood while $\mathcal{R}$ bacteria are not affected. However $\mathcal{S}$ bacteria have a fitness advantage on $\mathcal{R}$ bacteria and therefore do better in pairwise competition. In a similar way $\mathcal{R}$ has a fitness advantage over $\mathcal{C}$. As such, these types are engaged in a *rock-paper-scissors relationship*. Note however, that these three types do not face a completely symmetric situation as in the classical rock-paper-scissors game.

The fitness matrix

$$F = \begin{array}{c} \\ \mathcal{S} \\ \mathcal{C} \\ \mathcal{R} \end{array} \begin{array}{c} \begin{array}{ccc} \mathcal{S} & \mathcal{C} & \mathcal{R} \end{array} \\ \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ -1 & 1 & 0 \end{pmatrix} \end{array}$$

has the properties that in pairwise competition $\mathcal{C}$ beats $\mathcal{S}$, $\mathcal{S}$ beats $\mathcal{R}$, and $\mathcal{R}$ beats $\mathcal{C}$, so it matches the bacteria relationships in Kerr *et al.* [2002]. Moreover, using the replicator dynamics (global neighborhood), $\mathcal{S}$ will go extinct first, followed by $\mathcal{C}$ and ultimately leading to a population consisting of $\mathcal{R}$ exclusively. This is illustrated by Figure 4 based on our model and by Figure 2.b in Kerr *et al.* [2002].

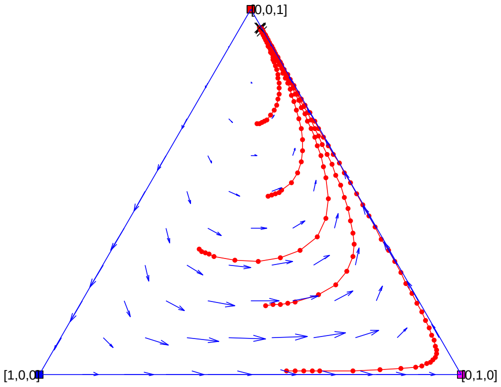

Fig. 4: Flow diagram of the replicator dynamics for fitness matrix F.
Type 1 is $\mathcal{S}$ (bottom left), type 2 is $\mathcal{C}$ (bottom right), and type 3 is $\mathcal{R}$ (top).

The following table shows what the different scenarios from section 2 lead to for fitness matrix F and initial distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

Table 2: Analysis of the different scenarios based on matrix F.

		Initial Distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
Benchmark		only type 3, $\mathcal{R}$, survives
Scenario	1	only type 2, $\mathcal{C}$, survives
	2	only type 3, $\mathcal{R}$, survives
	3	only type 3, $\mathcal{R}$, survives
	4	only type 3, $\mathcal{R}$, survives
	5	only type 1, $\mathcal{S}$, or type 3, $\mathcal{R}$, survives
	6	all three types survive

Just like in the scenario analysis of the previous section, Scenario 1 leads to a field consisting of a unique type surviving, which is $\mathcal{C}$ for fitness matrix F and initial distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

When using the dynamics of scenarios 2, 3 or 4, then the $\mathcal{R}$ type will take over the entire field, as described for the mixed plate experiment of Figure 2.c in Kerr *et al.* [2002].

For scenario 5 we observed that sometimes only $\mathcal{R}$ survives and sometimes only $\mathcal{S}$. This depends on whether $\mathcal{C}$ or $\mathcal{S}$ goes extinct first.

When using scenario 6 we observe that all three types remain coexisting in approximately equal shares over time. This reflects the experiments reported by

Kerr *et al.* [2002] on a static plate as illustrated in their Figure 2.a.

5. Concluding Remarks

In the previous section we have seen that scenario 6 fits best for the static plate experiment reported in Kerr *et al.* [2002], while other scenarios seem to be more appropriate for circumstances which depend less on locality. However, whether or not scenario 6 always provides a best matching with experimental data, remains for further explorations.

We would also like to mention that taking into account local circumstances, in calculating and simulating population development, opens the way to analyze much more complex interaction schemes. This is now possible because of highly improved computational possibilities. We can now use local fitness matrices, which allow to include local environmental features. Moreover, these fitness matrices can even depend on time as well as on local population distributions. As such, we can for example incorporate issues like resource depletion or symbiotic relations into the model. One can also depart from the hexagonal structure and examine more general network structures that represent the interactions among individuals. These issues will be part of further studies.

We would like to end this paper by mentioning two studies on locality and time dependencies reported by our former student Tak [2012] based on scenario 3. For the following symmetric fitness matrix T and the initial field displayed top-left in Figure 5, she found that this leads to a repetitive pattern:

$$T = \begin{pmatrix} 1 & 8 & 1 \\ 1 & 1 & 8 \\ 8 & 1 & 1 \end{pmatrix}.$$

This pattern is similar to migratory predator prey behavior: type 1 chases type 2, which in turn chases type 3, which again chases type 1, as is clear from matrix T.

As an example of a study that uses fitness matrices which depend on locality as well as on time, Tak [2012] worked with

$$A(y,t) = \begin{pmatrix} 0 & 0 & 2+\sin\left(\frac{t+y}{20}\cdot 2\pi\right) \\ 0.7 & 0 & 0.7 \\ 2-\sin\left(\frac{t+y}{20}\cdot 2\pi\right) & 0 & 0 \end{pmatrix}$$

where y is the vertical cell coordinate in a 20×20 field and where t is time. This lead to a periodically stable heart-shaped configuration of all three types. A movie of this dynamic process can be observed at
`http://youtu.be/_pJCcVkdL40`

The latter study is an adaptation of earlier work done on replicator dynamics with periodic fitness functions (cf. Uyttendaele *et al.* [2012]).

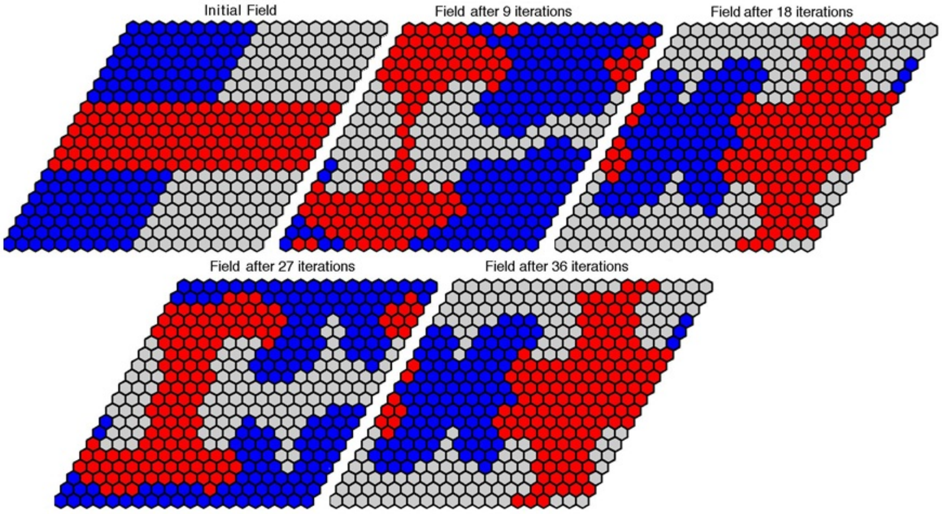

Fig. 5: Migratory predator prey behavior

Acknowledgement

We would like to thank Ronald Westra and Mandy Tak for their contributions in an early stage of this study.

References

Buttel, L.A., Durrett, R. and Levin, S.A. [2002] Competition and species packing in patchy environments. *Theoretical Population Biology* **61** (3), 265–276.

Hofbauer, J. and Sigmund, K. [1998] *Evolutionary Games and Population Dynamics.* Cambridge University Press.

Kerr, B., Riley, M.A., Feldman, M.W. and Bohannan, B.J.M. [2002] Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature* **418** (6894), 171–174.

Killingback, T. and Doebeli, M. [1996] Spatial evolutionary game theory: Hawks and doves revisited. *Proceedings of the Royal Society of London. Series B: Biological Sciences* **263** (1374), 1135–1144.

Kirkup, B.C. and Riley, M.A. [2004] Antibiotic-mediated antagonism leads to a bacterial game of rock–paper–scissors in vivo. *Nature* **428** (6981), 412–414.

Sandholm, W.H. [2011] *Population Games and Evolutionary Dynamics.* MIT Press.

Maynard Smith, J. and Price, G.R. [1973] The logic of animal conflict. *Nature* **246**, 15–18.

Tak, M. [2012] Experiments in local evolutionary dynamics. MSc Thesis, Maastricht University.

Taylor, P.D. and Jonker, L.B. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* **40** (1), 145–156.

Ulam, S. [1952] Random processes and transformations. *Proceedings of the International Congress on Mathematics* **2**, 264–275.

Uyttendaele, P., Thuijsman, F., Collins, P., Peeters, R., Schoenmakers, G. and Westra, R.

[2012] Evolutionary games and periodic fitness. *Dynamic Games and Applications* **2** (3), 335–345.

Von Neumann, J. [1966]. *Theory of Self-Reproducing Automata*. Edited by A.W. Burks, University of Illinois Press.

Wolfram, S. [2002] *A New Kind of Science*. Wolfram Media Inc.