Computational Economics **24:** 383–408, 2004.

Population Learning in a Model with Random Payoff Landscapes and Endogenous Networks

GIORGIO FAGIOLO[1,*], LUIGI MARENGO[2] and MARCO VALENTE[3]
[1]*Sant'Anna School of Advanced Studies, LEM, Pisa, Italy;* [2]*University of Teramo, DSGSS, Teramo, Italy;* [3]*University of L'Aquila, Faculty of Economics, L'Aquila, Italy.* **author for correspondence; e-mail: fagiolo@sssup.it*

Accepted 11 November 2004

Abstract. Population learning in dynamic economies with endogenous network formation has been traditionally studied in basic settings where agents face quite simple and predictable strategic situations (e.g. coordination). In this paper, we start instead to explore economies where the payoff landscape is very complicated (rugged). We propose a model where the payoff to any agent changes in an unpredictable way as soon as any small variation in the strategy configuration within its network occurs. We study population learning where agents: (i) are allowed to periodically adjust both the strategy they play in the game and their interaction network; (ii) employ some simple criteria (e.g. statistics such as MIN, MAX, MEAN, etc.) to myopically form expectations about their payoff under alternative strategy and network configurations. Computer simulations show that: (i) allowing for endogenous networks implies higher average payoff as compared to static networks; (ii) populations learn by employing network updating as a "global learning" device, while strategy updating is used to perform "fine tuning"; (iii) the statistics employed to evaluate payoffs strongly affect the efficiency of the system, i.e. convergence to a unique (multiple) steady-state(s); (iv) for some class of statistics (e.g. MIN or MAX), the likelihood of efficient population learning strongly depends on whether agents are change-averse in discriminating between options associated to the same expected payoff.

Key words: adaptive expectations, dynamic population games, endogenous networks, fitness landscapes, population learning

1. Introduction

The study of population learning in dynamic economies with many interacting agents has received an increasing attention in the last two decades (Kirman, 1997; Fagiolo, 1998). In particular, a major effort has been recently devoted to the investigation of dynamic population games with endogenous network formation.[1] In these models, agents are placed on the nodes of a graph and receive a payoff from repeatedly playing some bilateral, non-cooperative, game (e.g. coordination) against the agents they are currently connected with. Links may be costly and can be established (or removed) either unilaterally or bilaterally. Any agent is periodically allowed to revise the (pure) strategy it currently plays in the game. Moreover, from time to time, agents are called to adjust the set of agents they interact with by adding or removing single links. Strategy and link adjustments are typically carried out through simple best-response rules based on expected payoffs (Blume, 1993). Expectations are formed myopically: agents compute tomorrow's expected payoff by observing today's strategy and network configurations.

These studies have shed light on the importance of interactions and individual behaviors in shaping the outcomes of population learning. For instance, the type of behavioral rules employed by the agents to adjust their states and to form their expectations has been shown to affect dramatically the long-run behavior of the system, both in terms of its qualitative features (e.g. convergence to steady states, cycles, etc.) and in terms of its quantitative properties (e.g. efficiency).[2] Moreover, the geography of interactions has been proved to strongly influence the outcomes of population learning, especially when agents can only adjust their strategy over time for an exogenously given network.[3]

However, existing works have been exclusively focusing on economies where population learning takes place in very simple, predictable environments. Indeed, in all settings studied in the literature, individual players face very simple strategic situations (e.g. coordination). Furthermore, there is no uncertainty whatsoever in the economy: individual payoffs are common knowledge and each agent always plays the same game against any other agent in the population. Finally, the payoff of any single agent typically depends on some average levels of the behaviors of players belonging to its interacting set. This implies that individual payoffs are relatively invariant to permutations which preserve the frequency of agents currently playing a given strategy in the group. As a result, the overall payoff landscape over which population learning takes place is typically very *smooth* and is characterized by few peaks (i.e. local optima).

In this paper, on the contrary, we start to explore dynamic population games with endogenous network formation in the presence of very complicated (rugged) payoff landscapes. We focus on economies where individual payoffs are extremely sensitive to any small variation in the configuration of strategies and change almost unpredictably. This analysis is motivated by the observation that, in real-world situations, agents are typically uncertain about 'which game to play with whom' in any time period and, consequently, about the payoff that they might expect from any bilateral interaction (cf. Taylor (1987) for a quite similar perspective).

We propose a model where N agents are repeatedly called to revise *both* their binary strategy *and* their interaction group, that is the links with their opponents in the game. In order to describe population games where agents face complicated, multi-peaked, payoff landscapes, we employ the following three key assumptions. First, we suppose that there exists an underlying payoff structure defined as a random map. This map associates to each of the 2^N possible combinations of strategies a vector of N payoff realizations (one for each agent) of an i.i.d. random variable uniformly-distributed over $[0, 1]$. This underlying payoff landscape is given once and for all and represents a metaphor of an extremely complicated strategic situation.

Second, we suppose that agents can periodically revise both the strategy they play in the game and the set of agents they interact with by employing simple best-reply rules. To do so, they form myopically their expectations about next-period payoffs by observing current strategies and networks. In line with existing literature, we assume that the realized (and expected) payoff to any agent *i* only depends on the strategy played by *i* and the (expected) strategies played by those who interact with *i*.

Third, we assume that agents form their expectations about next period payoff by summarizing the information about payoffs. Indeed, any agent *i* who interacts with *h* other agents faces a *distribution* of payoffs containing 2^{N-h-1} values (one for each combination of strategies played by the remaining $N - h - 1$ agents). We suppose that agents try to summarize this large amount of information by computing some simple statistics on such distribution (e.g. MEAN, MAX, MIN). These rules (or criteria) might be interpreted as heuristics employed by the agents to cope with the uncertainty of the system. For instance, agents using the MIN (resp. MAX) criterion might be labeled as "pessimistic" (resp. "optimistic").

We study the long-run behavior of the model in two different settings. First, as a benchmark exercise, we explore an economy where agents can only update their strategy over time for an exogenously, initially given, interaction structure (static networks). This exercise allows us to study how the structure of the interaction network (e.g. its connectivity) affects population learning. Second, we study random population games with *endogenous networks*, i.e. a setup where agents can both adjust their strategies and decide whether to build or remove links with other agents in response of expected payoffs.

Due to the *ruggedness* of the underlying payoff landscape, the coevolution between strategies and networks induces a non-trivial adaptation process for the population of agents. In particular, we are interested in answering the following questions: Does endogenous network formation favor population learning, as compared to static networks? Under which behavioral conditions (e.g. expectation criterion employed) is the system able to climb (and settle upon) local or global optima? Which types of networks are likely to emerge in the long-run (e.g. highly vs. weakly connected)?

Computer simulations show that, for any given expectation rule, populations that can endogenously adjust networks as well as strategies systematically reach higher average payoff levels, as compared to populations that adapt over static interaction structures. We also find that agents employ network updating as a "global learning" device that allows them to attain major payoff improvements through *jumps* in the payoff landscape. On the contrary, strategy updating is typically used as a *fine tuning* device. Agents employ it to climb local (or global) optima once network adjustment has found a promising region of the payoff landscape. We then study how expectation criteria affect learning dynamics. Simulations indicate that "pessimistic" agents (i.e. using MIN) tend to develop highly connected networks but they never reach local (or global) optima. However, they attain a higher performance if they are change-averse (i.e. if they refuse network changes delivering the same payoff). On the contrary, populations with "optimistic" agents (i.e. using MAX) converge to weakly connected networks and are able to reach a global optimum only if individuals are change-lovers (i.e. they accept network changes even if they deliver the same payoff).

The rest of the paper is organized as follows. In Section 2 we formally describe the model and we present a simple example to illustrate its main features. In Section 3 we illustrate simulation results and we discuss their robustness with respect to departures from the basic model. Finally, Section 4 draws some conclusions and sketches future research.

2. The Model

Consider a finite population of agents $I = \{1, 2, \ldots, N\}$, $N \geq 3$, placed on the nodes (or vertices) of a non-directed graph. Time is discrete. In any period $t = 0, 1, 2, \ldots$ each agent $i \in I$ plays a game with binary pure-strategy set $S = \{-1, +1\}$ against all (and only) agents it is currently connected with through a bilateral link (i.e. an edge). We call *interaction group* the set $V_i^t \subseteq I - \{i\}$ of players with whom i plays the game at time t. Since we study a system where interactions are symmetric (i.e. $i \in V_j^t \Leftrightarrow j \in V_i^t$), at any t there exists a one-to-one relation between the *interaction structure* $(V_i^t)_{i=1}^N$ and the non-directed graph G_t containing all bilateral links currently in place in I. We also assume that interactions are costless for both agents (i.e. links can be built and removed for free).[4] Let us now turn to describe the payoff structure and the dynamics of the system.

2.1. PAYOFFS

In our economy, individual payoffs are extremely sensitive to any small variation in the configuration of strategies and change almost unpredictably. More formally, define by:

$$\pi_i : \{-1, +1\}^N \to [0, 1] \tag{1}$$

the payoff map which associates to each strategy configuration $\Omega \in \{-1, +1\}^N$ a payoff x (to agent i) defined as a realization of a random variable $X \sim U[0, 1]$, i.i.d. both across i's and Ω.[5] In what follows, we suppose that $(\pi_i)_{i=1}^N$ are given once and for all from the beginning of the process.

The underlying payoff structure $(\pi_i)_{i=1}^N$ sets the stage where expected and realized payoffs are defined. Indeed, in line with existing literature, we assume that realized payoff $u_i^t(\Omega^t, G^t)$ of any agent i who plays $s \in \{-1, +1\}$ and interacts with agents in $V \subseteq I - \{i\}$ only depends on the configuration of strategies currently played by the agents in $\bar{V} = V \cup \{i\}$:

$$u_i^t(\Omega^t, G^t) = u_i^t(\Omega^t(\bar{V})) := u_i^t\big(s; s_j^t, j \in V\big). \tag{2}$$

Agents are allowed to revise, at the beginning of $(t-1, t]$, their current strategy s_i^{t-1} and interaction group V_i^{t-1} on the basis of expected payoffs $E_i^{t-1}[u_i^t(s; s_j^t, j \in V)]$, where (s, V) are some alternative strategy and interaction group[6] and E_i^{t-1} denotes the expectation of agent i at the beginning of the time interval $(t-1, t]$. Since agents do not know the strategy their current and prospective partners will play at t, they must form an expectation about s_j^t, $j \in V$. We suppose that agents are myopic, i.e. that $E_i^{t-1}(s_j^t) = s_j^{t-1}$, $j \in I$, $j \neq i$. Therefore:

$$E_i^{t-1}\left[u_i^t(\Omega^t(\bar{V}))\right] = u_i^{t-1}(\Omega^{t-1}(\bar{V})). \tag{3}$$

How do agents evaluate payoffs $u_i^{t-1}(\Omega^{t-1}(\bar{V}))$? We assume that individuals try to summarize the information coming from the underlying payoff landscape (i.e. the maps π_i) by employing some "criterion" (or statistics) $\Re$. In fact, given the strategy configuration prevailing in $\bar{V} = \{i, j_1, \ldots, j_m\}$ at time $t-1$, any agent i faces 2^{N-m-1} payoff values, each one associated to any possible combinations of strategies played by the $N-m-1$ agents not currently linked with it.[7] To summarize this distribution, agents compute expected payoffs as:

$$u_i^{t-1}(\Omega^{t-1}(\bar{V})) = \Re\left(\Phi_i^{t-1}(\bar{V})\right), \tag{4}$$

where $\Phi_i^{t-1}(\bar{V}) = \{\pi_i(\Omega), \Omega \in \Theta_i^{t-1}(\bar{V})\}$ and $\Theta_i^{t-1}(\bar{V}) \subseteq \{-1, +1\}^N$ is the set of all configurations $\Omega = (\omega_i)_{i=1}^N \in \{-1, +1\}^N$ for which $w_j = s_j^{t-1}$, $j \in \bar{V}$, while strategies ω_j, $j \notin \bar{V}$ are allowed to vary in $\{-1, +1\}^{N-|\bar{V}|-1}$.

We suppose that all agents employ the same criterion $\Re$ during the entire dynamic process. We will experiment with the following three statistics: (i) $\Re$ = MEAN; (ii) $\Re$ = MIN; (iii) $\Re$ = MAX. Notice that if agent i interacts with all the others, the statistics will be trivially applied to a single value (i.e. the underlying payoff value associated to the current global configuration Ω^{t-1}). Thus, in this case, all statistics coincide. Conversely, if an agent is currently isolated (i.e. $\bar{V} = \{i\}$), then $\Re$ will be computed over a distribution whose size is 2^{N-1}. In general, the higher the number of links held by any agent, the smaller the cardinality of Θ_i^{t-1}.

2.2. DYNAMICS

Let us now describe the dynamic process governing individual strategy and network updating. Suppose that at time $t = 0$ a strategy configuration $\Omega^0 \in \{-1, +1\}^N$ and a graph G^0 over I are randomly drawn. In line with existing literature (Goyal and Vega-Redondo, 2001), we assume that at the beginning of any subsequent time interval $(t-1, t]$, $t \geq 1$ agents undergo a two-stage asynchronous updating process. In the first stage, two agents are drawn at random to update their networks given the current strategy-configuration Ω^{t-1}. In the second stage, one agent is drawn at random to update its strategy given the graph resulting from first-stage decisions.

In the *first* stage, two agents meet at random (say *i* and *j*). They evaluate their current payoffs (i.e. their expected payoff under the current state):

$$w_h^{t-1} = u_h^{t-1}\big(\Omega^{t-1}\big(\bar{V}_h^{t-1}\big)\big), \quad h = i, j. \tag{5}$$

If *i* and *j* are not connected, a link between them is tentatively added. To do so, they define their alternative interacting sets: $\tilde{V}_h = \bar{V}_h^{t-1} \cup \{k\}, h, k = i, j, k \neq h$. On the contrary, if *i* and *j* are already connected, they consider whether to remove the existing link and define $\tilde{V}_h = \bar{V}_h^{t-1} - \{k\}, h, k = i, j, k \neq h$. Then, both agents compute expected payoffs under the proposed change as:

$$\tilde{w}_h^{t-1} = u_h(\Omega^{t-1}(\tilde{V}_h)), \quad h = i, j. \tag{6}$$

We suppose that agents make their decisions by employing a deterministic, myopic, best-reply rule. In other words, agents simply compare payoffs before and after the proposed change and pick the choice delivering the highest expectation.

We consider two alternative tie-breaking rules (TBRs), namely: (i) TBR *without neutrality*: accept the change if and only if both agents are strictly better off under the change, i.e. add/delete *ij* if and only if $\tilde{w}_h^{t-1} > w_h^{t-1}, h = i, j$; (ii) TBR *with neutrality*: Accept the change if and only if no agent is strictly worse off under the change, i.e. add/delete *ij* if and only if $\tilde{w}_h^{t-1} \geq w_h^{t-1}, h = i, j$. Notice that according to the first TBR (*without neutrality*) agents might be labeled as *change-averse*, because they prefer to stick to their current choices unless the alternative option is associated to strictly larger payoffs for both. Conversely, agents using the second TBR (*with neutrality*) might be labeled as *change-lovers*, because they prefer to change even if the new choice is associated to the same payoff, i.e. they accept *payoff-neutral* changes.[8]

Network updating is therefore defined as follows:

$$G^t = \begin{cases} G^{t-1} \cup \{ij\} & \text{if the link has been added} \\ G^{t-1} - \{ij\} & \text{if the link has been deleted,} \\ G^{t-1} & \text{otherwise} \end{cases}$$

where we denote by $ij \in G^t$ the fact that in the graph G^t agents *i* and *j* are linked.

In the *second* stage, strategy updating takes place given G^t. We assume that an agent (say *i*) is drawn at random from *I*. Given (s_i^{t-1}, V_i^t) and current payoff $u_i^{t-1}(\Omega^{t-1}(\bar{V}_i^t)) := u_i^{t-1}(s_i^{t-1}, \Omega^{t-1}(V_i^t))$, it will switch to $-s_i^{t-1}$ at the beginning of period *t* if and only if:

$$u_i\big(-s_i^{t-1}, \Omega^{t-1}\big(V_i^t\big)\big) > u_i^{t-1}\big(s_i^{t-1}, \Omega^{t-1}\big(V_i^t\big)\big). \tag{7}$$

After the two updating stages, next time iteration starts given the new state of the system (Ω^t, G^t).[9]

2.3. A SIMPLE EXAMPLE

In order to further clarify how the model works, let us present a simple example. Let us consider a population of $N = 3$ agents: $I = \{1, 2, 3\}$. The underlying payoff map $(\pi_i)_{i=1}^3$ faced by any agent i is defined once and for all by assigning to each combination of strategies played by the 3 agents a realization of a i.i.d. random variable uniformly distributed over [0, 1]. Since each agent faces $2^{N-1} = 4$ possible configurations for each $s_i \in \{-1, +1\}$, the underlying payoff associated to the complete graph will be a matrix with $2 \cdot 2^{N-1} = 8$ rows (i.e. $\pi_i(s_1, s_2, s_3)$) and three columns, see Table I for an example.

Individual payoffs to any agent i depend on both the strategy s_i^{t-1} it currently plays and the strategies currently played by the agents which i is connected with. Therefore, if a link between i and j exists, payoffs of both i and j are affected by whether the other chooses -1 or $+1$. In order to compute expected payoffs and revise its current state, it computes statistics on the distribution of payoffs associated to all combinations of strategies played outside its interaction group. For instance, if agent $i = 1$ is connected with both 2 and 3 and currently plays $s_1^{t-1} = -1$, it will face (irrespective of the expectation rule employed) a payoff:

$$w_1\left(-1; s_2^{t-1}, s_3^{t-1}\right) = \begin{cases} 0.56 & \text{if } \left(s_2^{t-1}, s_3^{t-1}\right) = (-1, -1) \\ 0.77 & \text{if } \left(s_2^{t-1}, s_3^{t-1}\right) = (-1, +1) \\ 0.58 & \text{if } \left(s_2^{t-1}, s_3^{t-1}\right) = (+1, -1) \\ 0.55 & \text{if } \left(s_2^{t-1}, s_3^{t-1}\right) = (+1, +1) \end{cases}.$$

Let us assume that $s_1^{t-1} = -1$ and that the statistics employed by agent 1 is the arithmetic mean (MEAN). If 1 is currently connected with 2 only, then for each

Table I. An Example of the payoff landscape with $N = 3$.

(s_1, s_2, s_3)	π	π_2	π_3
$(-1, -1, -1)$	0.56	0.11	0.24
$(-1, -1, +1)$	0.77	0.54	0.17
$(-1, +1, -1)$	0.58	0.31	0.45
$(-1, +1, +1)$	0.55	0.59	0.19
$(+1, -1, -1)$	0.78	0.42	0.70
$(+1, -1, +1)$	0.32	0.54	0.25
$(+1, +1, -1)$	0.04	0.44	0.78
$(+1, +1, +1)$	0.80	0.67	0.44

possible choice of agent 2, it faces a distribution of only two payoffs (each one associated to the strategy that agent 3, the not connected one, might play). Thus:

$$w_1\left(-1; s_2^{t-1}\right) = \begin{cases} \frac{1}{2}(0.56 + 0.77) & \text{if } s_2^{t-1} = -1 \\ \frac{1}{2}(0.58 + 0.55) & \text{if } s_2^{t-1} = +1 \end{cases}.$$

Along the same line, if 1 is currently connected with no other player, its payoff will be:

$$w_1(-1; \cdot) = \frac{1}{4}(0.56 + 0.77 + 0.58 + 0.55).$$

Finally, if agent 1 decides to consider payoffs associated to choosing +1 and connecting with 3 only, it will get:

$$w_1\left(+1; s_3^{t-1}\right) = \begin{cases} \frac{1}{2}(0.78 + 0.04) & \text{if } s_3^{t-1} = -1 \\ \frac{1}{2}(0.32 + 0.80) & \text{if } s_3^{t-1} = +1 \end{cases}.$$

Suppose now that at some time $t - 1 > 0$ the current strategy configuration is (+1, +1, +1) and that agent 1 is connected with 2 but not with 3, while 3 is isolated. If 1 and 3 are drawn for a network update, they will consider to add a link between them. Agent 1 will compare its current payoff $\frac{1}{2}(0.04 + 0.80)$ with the payoff after link addition, namely 0.80. Agent 3 will compare its old payoff $\frac{1}{4}(0.17 + 0.19 + 0.25 + 0.44)$ with the new one $\frac{1}{2}(0.25 + 0.44)$. In this case both agents are (strictly) better off under the change and the link will be added. After network updating, player 3 is connected with 1 only. If it were called to strategy updating, it would switch to -1 since $\frac{1}{2}(0.70 + 0.78) > \frac{1}{2}(0.25 + 0.44)$.

It is worth pointing out that the coevolution between strategies and networks will induce a non-trivial adaptation process for the population of agents. Due to the ruggedness of the underlying payoff landscape, interesting questions concern the ability of the population to climb local (or global) optima and whether the effectiveness of population learning is affected by the class of rules employed by the agents to form their expectations about payoffs.

3. Simulation Results

Unless otherwise specified, we employ the following simulation design. For a given population size N, we first generate an underlying payoff landscape $\Pi = (\pi_i)_{i=1}^N$. To allow for easier comparisons, we study adaptation of populations over the same Π.[10] We run Montecarlo experiments to explore how different behavioral assumptions affect population learning. Therefore, our *independent* variables will be the expectation rule $\Re$ and the kind of TBR (neutral vs. non-neutral) employed by the agents. For each choice of our independent variables, we let M populations evolve given different initial conditions (i.e. strategy and network configurations). Our *dependent* variables, i.e. the statistics we shall study, are the number of links held by each agent, its current payoff, as well as population averages of these variables. These statistics are typically computed when the system has reached either a steady-state or a sufficiently stable behavior. In what follows, all results refer to $N = 15$ and $10 \leq M \leq 100$.[11]

3.1. STATIC NETWORKS

In order to appreciate how networks affect the dynamics of the system, we first run a set of simulations where agents can only modify their strategies for a given, initially drawn, interaction structure (i.e. the first stage of updating process is suppressed).[12]

In this case, it is easy to see that a population of totally disconnected agents will converge to a steady state very quickly. Since individual payoffs do not directly depend on the strategies played by the others, each agent will compute its expected payoff over the largest set of values (2^{N-1}). As these values are drawn from a uniform random variable, all payoffs will be very close to the expected value of the statistics (0.5 if $\Re$ = MEAN, 0 if $\Re$ = MIN and 1 if $\Re$ = MAX). On the other hand, if one initializes the system with a complete network (i.e. all agents hold $N - 1$ links), the population never settles to a stable state because any single strategy switch influences all others' payoffs in a random way. As a result, a fully connected population will continuously move across the payoff landscape. Population learning will not be very effective, as average payoffs will keep oscillating in [0, 1]. Again, the criterion employed will not affect these results because all statistics always return the same value if a complete network is in place.

But what happens in the intermediate cases? To answer this question, we have studied how the long-run distribution of individual payoffs changes as one increases the average number of links in the economy.[13] Figures 1–3 show the distribution of long-run individual payoffs stemming from all $M = 100$ populations in each of the three behavioral setups (MEAN, MIN, MAX) as a function of the connectivity of the interaction structure. As one might have expected, the results show that for all three statistics the higher the number of links, the smaller the size of the payoff pool, and the more dispersed the payoffs.

Network connectivity also affects the dynamic behavior of the system in the long-run. Indeed, populations with a relative high number of links do not tend to converge to a steady-state. Agents always prefer to switch their state, thus influencing each other's payoff. This prevents population learning from reaching any stable strategy profile. On the contrary, populations characterized by few links tend to learn easily how to climb on a local optimum and stay there, because each agent can act in a relatively autonomous way.

Nevertheless, even in weakly connected populations, there can be cases where a stable state is never reached. Figure 4 provides an example where the average payoff

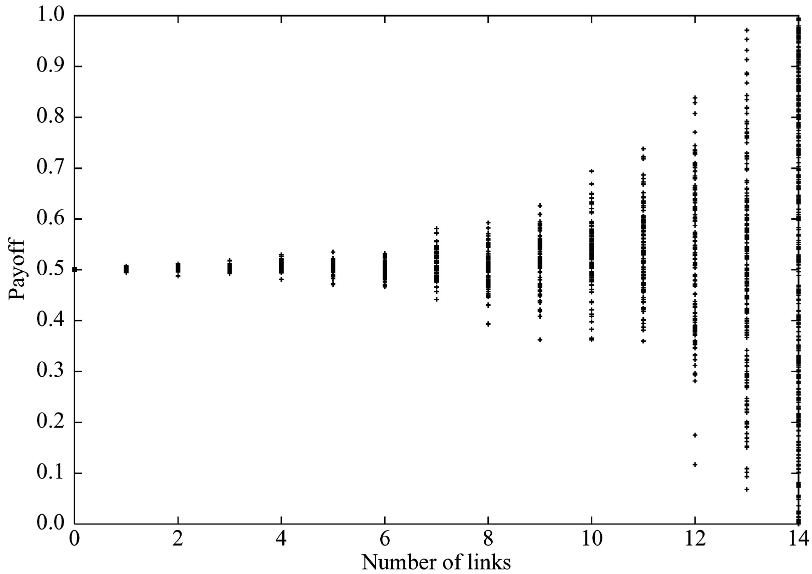

Figure 1. Static Networks with agents employing the MEAN evaluation rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the steady-state (across 100 populations of 15 agents).

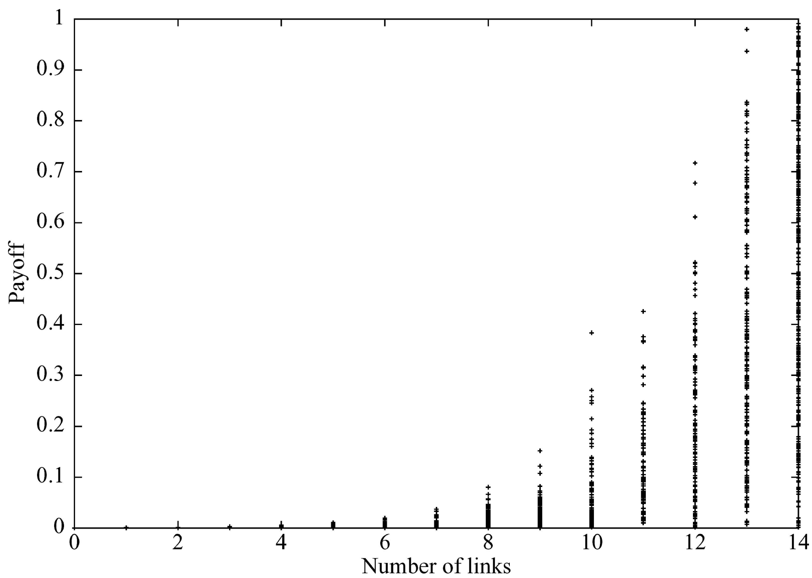

Figure 2. Static Networks with agents employing the MIN evaluation rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the steady-state (across 100 populations of 15 agents).

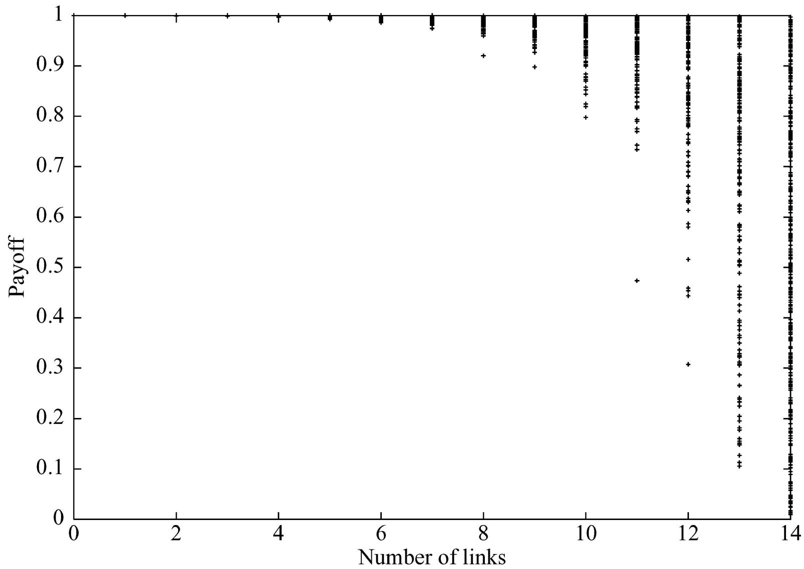

Figure 3. Static Networks with agents employing the MAX evaluation rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the steady-state (across 100 populations of 15 agents each).

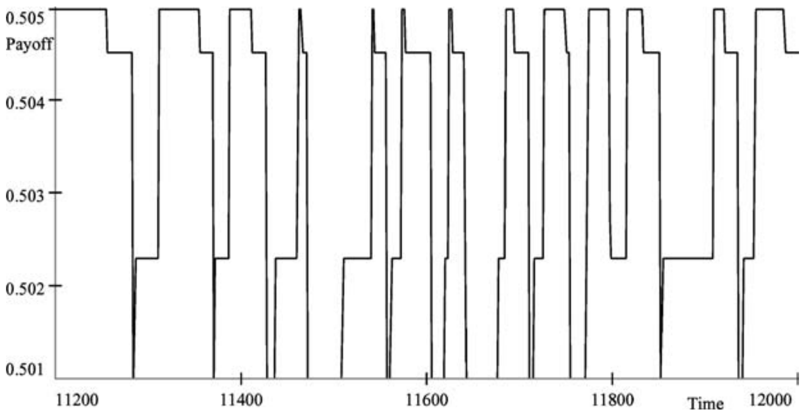

Figure 4. Static Networks with agents employing the MEAN rule. Time-series of population average payoffs for a weakly connected population where two linked agents have incompatible payoffs (see Table II at page 12).

of a population characterized by very few links (on average 0.5) cycles among four payoff levels. To see why this can happen, suppose at one extreme that only two agents in the population are linked and that their underlying payoff matrices are as in Table II.

Inspection of Table II shows that, irrespective of the expectation rule employed, the population will cycle forever among the four strategy combinations associated to the two linked agents.[14]

Table II. A payoff matrix for two linked agents leading to a never ending cycle. Payoffs in bold text indicate the agent who prefers to switch state.

Strategies		Payoffs	
s_l	s_2	Agent 1	Agent 2
0	0	0.7	**0.2**
0	1	**0.2**	0.7
1	1	0.4	**0.3**
1	0	**0.3**	0.4

3.2. ENDOGENOUS NETWORKS

Let us now turn to the more general case where agents are allowed to perform both strategy and network updating. We are interested in two main issues. First, we ask whether, for a given statistics $\Re$, the performance of a system with endogenous network updating is higher than in the case where networks were static (e.g. in terms of average payoffs). Second, we study the behavior of the system in different settings. Notice that in the presence of network updating the latter exercise is not only confined to the comparison of populations endowed with different expectation criteria. In fact, depending on the statistics employed, whether agents use neutral vs. non-neutral TBRs in network updating may have important consequences. To see why, simply recall that the addition (resp. deletion) of a link causes the set of values over which the statistics $\Re$ is computed to shrink (resp. enlarge). Agents using the MEAN statistics will therefore obtain different payoffs when their networking structure is changed: whether they are change-averse or not in undertaking network decisions is completely irrelevant. Conversely, by using MIN or MAX it is well possible that payoffs remain constant when a link is added or removed. In these cases, agents employing neutral vs. non-neutral TBRs might behave differently.

3.2.1. *MEAN Payoff Statistics*

Consider first a population composed of agents employing the MEAN statistics to compute expected payoffs. The first question we ask is: Do endogenous networks systems favor population learning, as compared to static networks? In Figure 5 we report average payoffs for two groups of 10 populations each. Both groups are initialized with the same average number of links ($\frac{1}{2}(N-1) = 7$ in this exercise). In the first group, all populations contain agents which are only allowed to update their strategies, while in the second group agents can modify their links as well as strategies. The group of populations with endogenous networks clearly outperforms their peers living in a static network. Thus, the possibility to act upon the link structure proves to be an advantage as compared to agents adapting over fixed networks. This is because endogenous networking

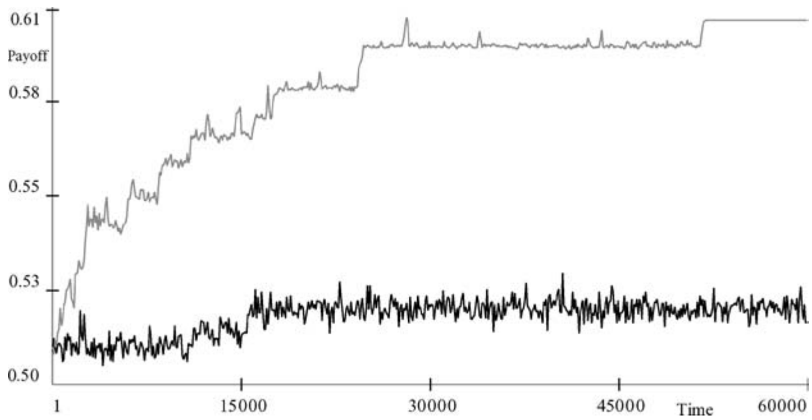

Figure 5. Time-series of population average payoffs for two groups of populations where agents use the MEAN rule. The first group (black series) contains populations adapting over static networks. The second group (grey series) contains agents where endogenous network updating is allowed. For each group, data refers to averages across 10 populations initialized with 7 links (on average).

agents have more ways to improve their payoff (for a similar finding, cf. Fagiolo (2004)).

Although this result may seem obvious, it is nevertheless interesting to spell out the sources of such an advantage. We run for this purpose a simulation using 50 populations living in an endogenous network system. Each population is initialized with a different number of links per agent. All populations with endogenous networks reach a steady state, albeit at different times. However, each population settles in a different global configuration, showing that there are a multiplicity of local optima. In these local optima agents tend to build networks with intermediate connectivity. Indeed, in our simulations all populations end up with a number of average links per agent ranging between 6 and 11, irrespective of the initial number of links.

Furthermore, the average steady-state number of links tends to be positively correlated with average steady-state payoffs, as Figure 6 shows. Therefore, populations which succeed in reaching a steady-state with highly connected networks enjoy higher payoffs on average. We have seen however that in these cases agents find it more difficult to coordinate. This is why we typically observe many populations that get stuck in weakly connected networks. By reducing interdependencies agents are able to find a Pareto-efficient configuration. More in detail, Figure 7 provides the scatter plot of steady-state individual payoffs as a function of steady-state individual number of links. If we compare this plot with its equivalent in the case of static networks (cf. Figure 1), it is easy to see that endogenous networks allow agents to occupy the upper half of the payoff distribution.

A less obvious result concerns the characteristics of the dynamic process through which population learning approaches a steady-state. All populations exploit for a

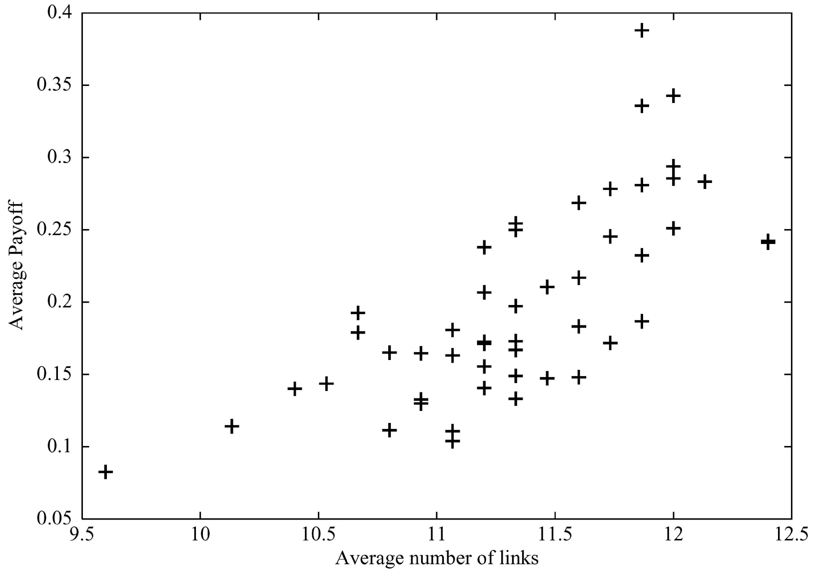

Figure 6. Endogenous networks with agents using the MEAN rule. Scatter-plot of population average payoffs (y-axis) vs. population average number of links (x-axis). Data refers to steady-state values of 50 populations of 15 agents each.

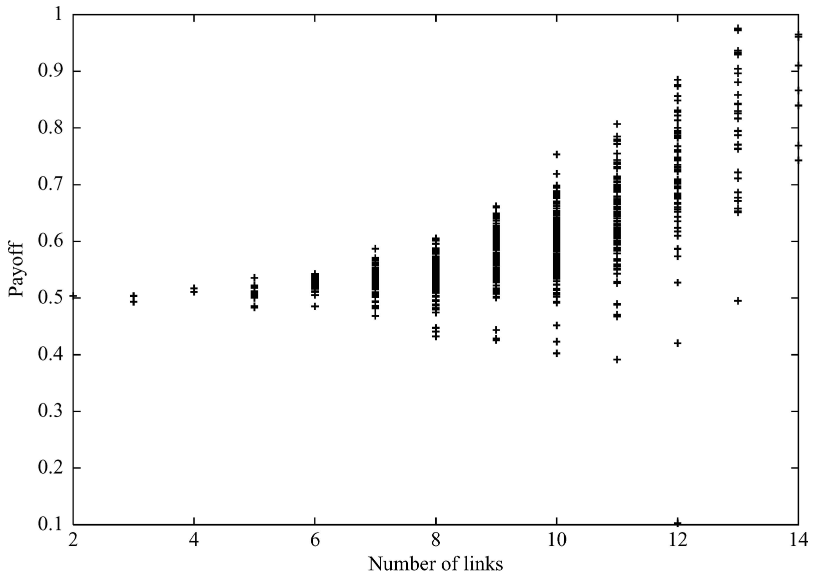

Figure 7. Endogenous networks with agents using the MEAN rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis). Data refers to steady-state values of 50 populations of 15 agents each.

while both network and strategy adjustments. In all cases, the number of successful network innovations is larger than successful strategy switches (about 24% more). This is quite surprising if we recall that network updating need mutual consent, while strategy switches are unilateral decisions. Moreover, while approaching a stable state all populations reach firstly a stable network structure. From then on they perform only strategy switches until a local optimum is found.

The robustness of such result deserves a closer investigation. In fact, irrespective of the criterion employed, our populations seem to use network updating as their main tool for getting major learning improvements (i.e. big jumps in the payoff landscape). Only when big improvements are not possible anymore, they start to "fine tune" their global configuration by exploring the strategy space for a given network structure. Of course, it may frequently happen that in the last step of local search, the best strategy configuration is not compatible with the current network structure for some agents. In these cases, network exploration begins once again. This result supports the hypothesis that partner selection is more effective than strategy revision as far as exploration of completely new regions of the payoff landscape is concerned. However, strategy updating becomes the best device when the population needs to exploit realized improvements and perform cumulative learning.

3.2.2. *MIN Payoff Statistics*

A richer set of results can be obtained when one introduces the hypothesis that agents employ the MIN payoff evaluation rule. We can consider this payoff rule as the one played by "pessimistic" individuals. Indeed, take all combinations of strategies played by the agents currently not linked to i as the ones outside i's control. If i employs the MIN rule, it will try to maximize its payoff in the worst of the uncontrolled cases.

Contrary to what happens with the MEAN rule, here the system (almost) never reaches a steady-state.[15] Moreover, results heavily depend, as expected, upon acceptance of neutral network updating. We begin by comparing the behavior of three groups consisting of 10 populations each. In the first group networks are static, while in the other two groups endogenous networks are assumed. The latter two groups differ because the second uses a TBR *with* neutrality (agents are change-lovers) while the third uses a TBR *without* neutrality (agents are change-averse).

Figure 8 shows the average payoffs for the three groups. The highest average payoff is reached by populations adapting over endogenous networks which refuse neutral network changes; the intermediate level refers to the populations with endogenous networks and accepting neutral changes; the lowest average payoff is produced by populations with static networks.

Some considerations are in order. First, notice that populations with endogenous networks tend to increase the number of links up to (almost) its maximum. In fact, the statistics used to compute expected payoffs cannot decrease by removing a link, while it might increase by adding a link. Thus, the group of populations with

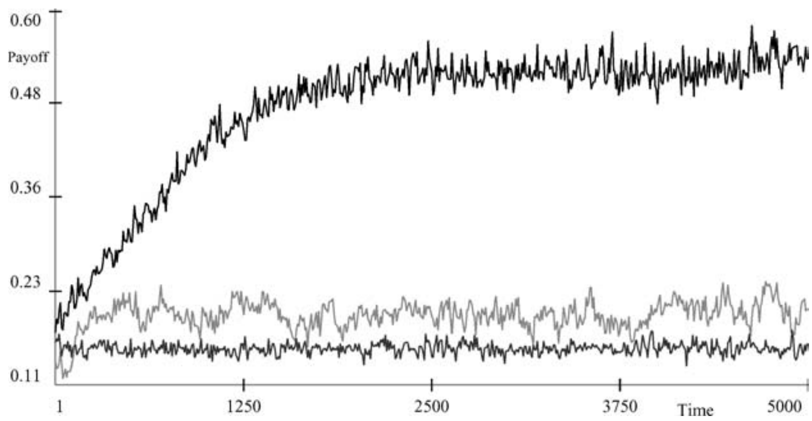

Figure 8. Endogenous networks. Time-series of population average payoffs for three groups of populations where agents use the MIN rule. Black series: endogenous networks and refusal of neutral changes. Light grey series: endogenous networks and acceptance of neutral changes. Dark grey series: static networks. For each group, data refers to averages over 10 populations.

constant networks suffers from the existence of populations characterized by an initially weakly connected network.

Second, our results show that refusing payoff neutral network changes provides a strong advantage. The reason is that TBRs without neutrality prevent the agents to reach stable, fully connected networks. In fact, it can happen that agents holding too many links reach the same payoff by removing a single link. However, this change enlarges the set over which the minimum is computed and makes it easier for subsequent strategy switches to generate a sensible drop in the payoff. Simulations indeed show that populations accepting neutral changes experience payoff drops which are typically larger than those where agents are change-averse. In this latter case, agents quickly form many links. Although continuously changing, their payoffs will be always the MIN over singleton sets, and therefore, on average, higher than in the case the MIN is computed over larger sets.

Finally, note that the "pessimistic" attitude of agents using a MIN criterion is well in line with the result that agents tend to develop many links. Indeed, agents use a decision criterion based upon the comparison between the worst possible conditions out of their control. Therefore, the resulting behavior should consist in extending as much as possible the ability to observe others' strategies. In other words, a "pessimistic" attitude generates a tendency towards over-control. At the population level, this in turn implies persistent payoff oscillations and poor learning.

3.2.3. *MAX Payoff Statistics*

When using the MAX payoff evaluation criterion, agents adopt a sort of "optimistic" criterion, as they base their decisions upon the best value which they could derive

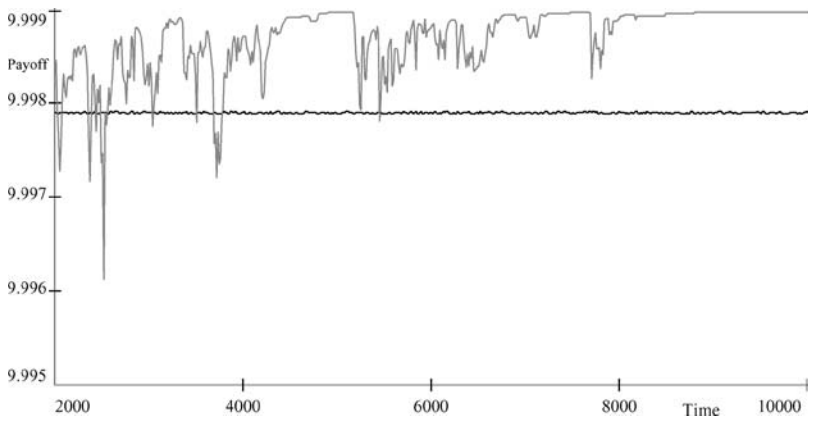

Figure 9. Endogenous networks. Time-series of population average payoffs (×10) for two groups of populations where agents use the MAX rule. Black series: endogenous networks and refusal of neutral changes. Grey series: endogenous networks and acceptance of neutral changes. For each group, data refer to averages over 10 populations.

from the behavior of *unobservable* agents. For similar albeit antithetical reasons, also in this case endogenous network updating favors population learning as compared to static networks. Agents adopting the MAX statistics are never worse-off by removing a link, since this enlarges the set over which the criterion is computed. As a result, populations with endogenous networks quickly reduce their links and improve their payoff in the early stages of the process.

Contrary to the MIN case, change-lovers agents earn *higher* payoff values (see Figure 9) as compared to change-averse ones. Indeed, by accepting neutral network modifications, agents continue to change their network structure over time. Nevertheless, all populations, for all initial conditions, reach the *same* (global) optimum. In these steady-states, all agents get a payoff very close to 1 because of the MAX statistics. To see why this happens, consider agent h and call s_h^* the strategy associated to the highest payoff value π_h^* in its payoff map π_h. If any two agents i and j do not currently play (s_i^*, s_j^*), it is sufficient for them to avoid forming a link. If they instead play the two compatible strategies (s_i^*, s_j^*), they can always reach (π_i^*, π_j^*) regardless the presence of a link, because they are "optimistic" and change-lovers. This result is depicted in Figure 10, where we plot individual steady-state payoff values attained by change-lover agents in 10 populations. On the horizontal axis, we distribute our 10 populations side by side. The cyclical pattern proves that each agent gains the same, highest as possible, identical payoff in each population.[16]

Conversely, populations where agents are change-averse do not reach the unique global maximum, but only a local one. Which local optimum is reached depends on initial conditions. Agents in such populations keep removing their initially assigned links until no further improvement is possible. Typically, any two linked agents get stuck in a state where they enjoy a relatively high payoff but the only

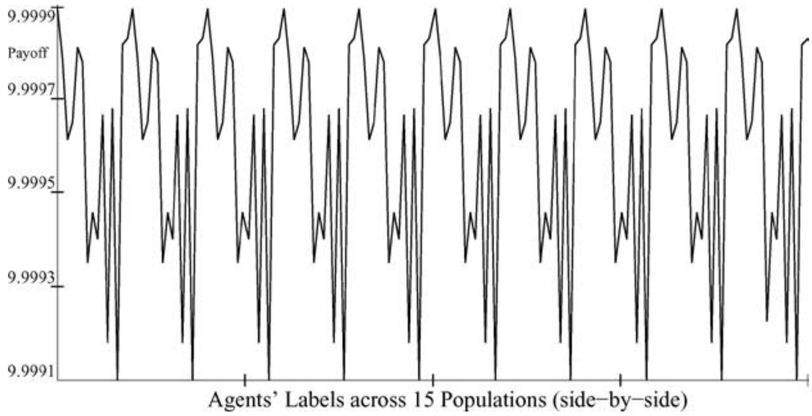

Figure 10. Endogenous networks when agents use the MAX rule and accept neutral changes. Individual Steady-State Payoffs (y-axis) for 10 populations ($\times 10$). On the x-axis we put alongside blocks of agents' labels $1, 2, \ldots, 15$ within each population $k = 1, \ldots, 10$. The cyclical pattern shows that each single agent i gets in the steady-state the same payoff in all 10 populations $k = 1, \ldots, 10$.

possible improvement is associated to a reciprocal strategy switch. Since link removal requires a strict improvement on both sides and strategy updating cannot be simultaneously performed, the "final" jump to the global optimum can never be attained.

Notice that "optimistic" agents act as if the strategies played by those who are not their partners were the most profitable ones. Hence, they typically tend to ignore the actual behavior of the others and to avoid the exploration of large parts of the payoff landscape. However, in order to obtain the best from this behavior it is necessary to act as "change-lovers" and accept network changes even when they do not appear to engender immediate, strictly positive, improvements. These apparently useless changes may allow to escape local optima and to climb the global one. On the contrary, change-averse agents get trapped in local optima because, by refusing apparently neutral changes, they prevent themselves from exploring other, possibly more profitable, regions of the payoff landscape.

3.3. BEYOND THE BASIC MODEL

In this Section we discuss some alternative specifications of our basic model. We are interested in checking the robustness of the foregoing results under alternative assumptions concerning the mechanisms governing network updating and expectation formation.

3.3.1. *Unilateral Link Removal*

We first modify the network updating rule to allow for unilateral link deletion. In the model presented above (see Section 2), both link formation and deletion

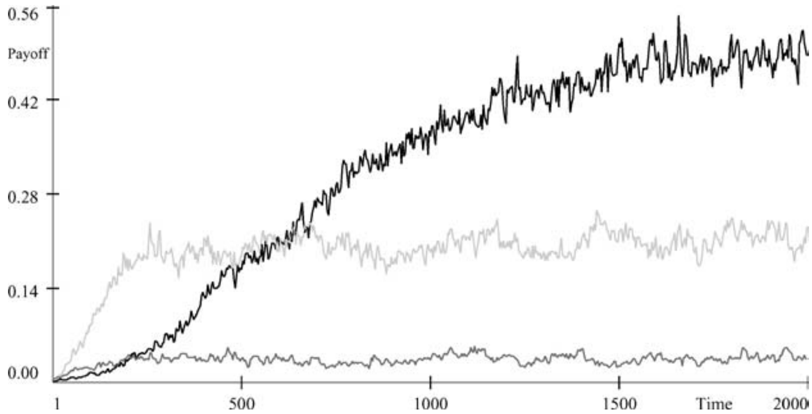

Figure 11. Unilateral Link Removal. Time Series of population average payoffs for 3 groups of populations where agents use the MIN rule. Black series: Bilateral link removal and agents refusing neutral changes. Light grey series: Bilateral link removal and agents accepting neutral changes. Dark grey series: Unilateral link removal. For each group, data refers to averages over 10 populations.

require mutual consent. This assumption is consistent with idea that whether any non-directed links *ij* is in place or not depends on common agreement by the two nodes. Suppose instead that adding a link still requires both agents to agree, while for a link removal it suffices that at least one agent is better-off under the change.[17] Network updating with *neutral TBR* would then read:

$$G^t = \begin{cases} G^{t-1} \cup \{ij\} & \text{if } \tilde{w}_j^{t-1} \geq w_h^{t-1}, \quad h = i, j \\ G^{t-1} - \{ij\} & \text{if } \tilde{w}_i^{t-1} \geq w_i^{t-1} \quad \text{or} \quad \tilde{w}_j^{t-1} \geq w_j^{t-1}, \\ G^{t-1} & \text{otherwise} \end{cases}$$

where $\tilde{w}_h^{t-1}$ (resp. w_h^{t-1}) are expected payoffs with (resp. without) the proposed change.[18] In this new setup, sustaining highly connected networks becomes harder, as any agent tends to avoid interactions as soon as they become unprofitable. Thus, an interesting question concerns whether "pessimistic" populations still attempt to over-control their environment. In fact, simulations show that agents using the MIN statistics tend to build less links and attain a lower payoff as compared to their peers performing bilateral link removal (see Figure 11). A similar result can be obtained also for agents using the MEAN criterion, cf. Figure 12. On the contrary, unilateral link formation does not affect "optimistic" agents. In particular, those accepting neutral payoff changes are still able to climb the global optimum.

3.3.2. *Costly Links*

Suppose now that links are costly. In line with Goyal and Vega-Redondo (2001) and Jackson and Watts (2002), we assume that each agent must pay a cost $c \in (0, 1)$ to

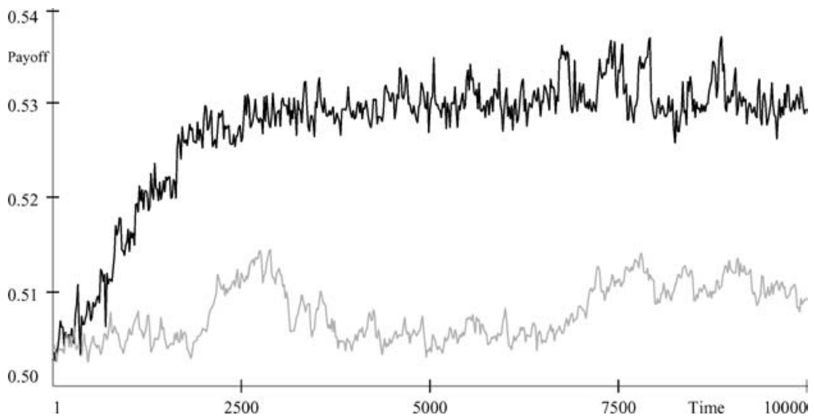

Figure 12. Unilateral Link Removal. Time Series of population average payoffs for 2 groups of populations where agents use the MEAN rule. Black series: Bilateral link removal. Grey series: Unilateral link removal. For each group, data refers to averages over 10 populations.

form a link. Of course, this additional assumption does not affect strategy updating, as the latter is performed *given* networks. On the contrary, our network updating rule changes as follows:

$$G^t = \begin{cases} G^{t-1} \cup \{ij\} & \text{if } \tilde{w}_h^{t-1} - w_h^{t-1} \geq c, \quad h = i, j \\ G^{t-1} - \{ij\} & \text{if } \tilde{w}_h^{t-1} - w_i^{t-1} \geq -c, \quad h = i, j, \\ G^{t-1} & \text{otherwise} \end{cases}$$

where $\tilde{w}_h^{t-1}$ and w_h^{t-1} are expected payoffs with and without the proposed change. Notice that now whether agents accept *neutral* network updating or not is irrelevant even if agents use MIN or MAX since *c* is picked on the real unit interval. Moreover, as *c* increases, link formation becomes less frequent because it requires larger and larger payoff improvements. Conversely, agents should perform more link deletions since they are willing to delete a link even if the net payoff change turns out to be negative but greater than $-c$.[19] For these reasons, a costly-link system behaves similarly to an economy where a unilateral link removal rule is assumed. More specifically, as one increases c, the behaviors of populations employing MEAN, MIN and MAX rules converge to a long-run steady-states characterized by "isolated" agents, with average payoffs close to their expected values.

3.3.3. *Sampling and Expectation Formation*

In our model, agents form expectations about the payoff they could get under any alternative state (s, V) by computing statistics on the distribution of payoffs associated to *all* possible strategy combinations that agents outside *V* can notionally play. This means that any agent *i* computes the criterion $\Re$ on $2^{N-|V|-1}$ payoff values, where $|V|$ is the number of links held by *i*. Assume instead that agents can only

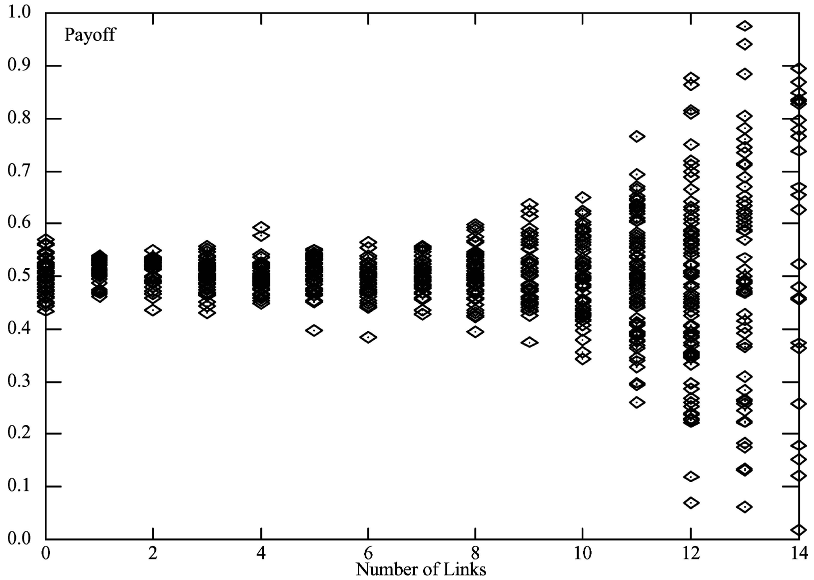

Figure 13. Payoff sampling where agents have a capacity bound $m = 100$. Static Networks with agents employing the MEAN evaluation rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the steady-state (across 50 populations of 15 agents).

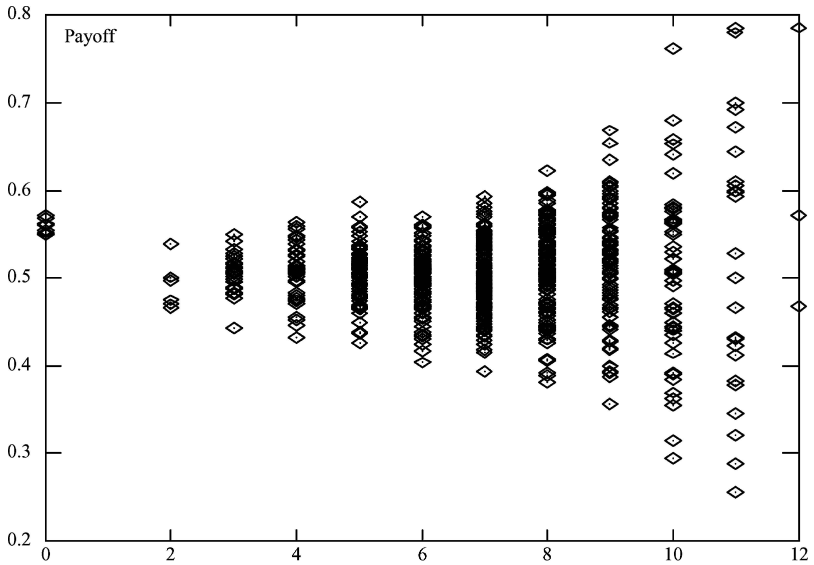

Figure 14. Payoff sampling where agents have a capacity bound $m = 100$. Endogenous Networks with agents employing the MEAN evaluation rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the steady-state (across 50 populations of 15 agents).

take into account *no more* than m such values, where $1 \leq m < 2^{N-1}$. In fact, agents can have a bounded computing and/or memory capability (or some budget constraint) that prevents them from recalling all possible values if their number exceeds m.

We will then assume that agents now form their payoff expectations as follows. If V is such that $2^{N-|V|-1} \leq m$, they simply consider $2^{N-|V|-1}$ possible payoff values and behave as before. If, conversely, the computational burden required exceeds their capacity m, agents randomly sample (with replacement) m payoffs values and compute $\Re$ on such a restricted pool.

To investigate this setup, we study how long-run individual payoffs depend on the individual number of links in 50 populations using the MEAN statistics and adapting over static networks for $m = 100$ (cf. Figure 13). It is easy to note that agents holding a low or an intermediate number of links now suffer from a higher payoff variability (compare to Figure 1). To see why, recall that for any (sufficiently small) value of m, agents now face a "double" trade-off between holding few and many links. We already know that if $m = 2^{N-|V|-1}$ for any V, agents with a smaller number of links face a payoff distribution which is closer to the "true" one. Conversely, agents with no capacity constraints and many links form their expectations using a distribution involving a lot of payoff variability. If however such capacity is bounded, there exists another source of payoff variability: if an agent holds few links but m is small, it may well be the case that the sampled distribution departs remarkably from a $U[0, 1]$ even if the one it would have faced if $m = 2^{N-|V|-1}$ is very close to the "true" one. For the same reason, strongly connected agents behave as if they had no constraints whatsoever.

Consider now the case of endogenous networks. Our simulations show that the introduction of a capacity constraint destroys the positive correlation between individual payoffs and individual number of links that we found in Figure 7. Indeed, as depicted in Figure 14, agents employing the MEAN criterion still tend to develop an intermediate number of links, but only those who end up isolated are able to converge to steady-states where the average payoff is larger than 0.5. For such agents, the capacity bound seems to be beneficial.

4. Conclusions

In this paper, we have proposed a preliminary investigation of population learning over 'rugged' landscapes, where agents face a strong uncertainty about expected payoffs from bilateral interactions. Our results seem to indicate that, in presence of complicated payoff landscapes and rather uncertain rewards, the details of the process governing the co-evolution between networks and strategies become crucial to understand the outcomes of population learning. In particular, the properties of the expectation formation rule strongly affect long-run network connectivity. The latter is indeed higher when agents are "pessimistic" and lower when they are "optimistic". Notice however that, unlike standard literature (e.g. Goyal and Vega-Redondo (2001)), here highly connected networks typically imply *more* payoff volatility and coordination problems. Therefore, populations that are able to build weakly connected networks can also reach higher local or even global optima. If this is the case (e.g. agents using MAX), population learning becomes efficient only if agents manage to avoid lock-ins by accepting changes which are payoff neutral. In general this can be done by employing neutral network updating to attain major payoff improvements through *jumps* in the payoff landscape.

Although our findings seem to be quite robust to alternative specifications of the basic model, many issues remain to be explored.

First, one might consider alternative formulations for the underlying payoff landscape $(\pi_i)_{i=1}^N$. So far, we have assumed that population learning takes place over a payoff landscape characterized by the highest possible ruggedness. Indeed, any occurrence of π_i is supposed to stem from a uniform random variable and to be totally uncorrelated across agents and strategy configurations. This generates an environment where the underlying games have no structure whatsoever and population learning (with endogenous networks) can be studied as if individual rewards from any bilateral interaction were completely unpredictable.

Such extreme assumptions can be weakened in two ways. On the one hand, one might assume that occurrences of π_i stemming from the same strategy configuration are correlated across agents (but still completely uncorrelated across configurations for each agent). For example, one might assume that across-agents correlation is 1, i.e. that each agent has the same payoff map $\pi_i \equiv \pi$. The resulting payoff landscape will then admit (for N sufficiently large) a unique maximum, which is associated to the strategy configuration $(s_i)_{i=1}^N$ (i.e. to the row of the matrix) delivering the highest level of π. An interesting question here concerns whether MAX populations endowed with neutral TBRs are able to climb this global maximum (as they do in correlated landscapes).

On the other hand, one might explore a setup where for each agent there exists some correlation across configurations (but for any given configuration payoffs are uncorrelated across agents). For instance, we can suppose in line with Kauffman (1993) that: (i) each agent is ex-ante connected through bilateral links with (on average) K other agents, where $0 \leq K < N - 1$; (ii) for each agent i, all 2^{N-K-1} payoffs values π_i associated to all strategy combinations played by the non-linked agents are constant; (iii) the 2^{K+1} payoff values associated to all strategy combinations played by i and by its partners are i.i.d. $U[0, 1]$. Here K becomes a parameter governing the smoothness of the underlying payoff landscape. The higher K, the more rugged the payoff landscape over which population learning takes place (in our basic model $K = N - 1$). Put it differently, by considering lower values for K we could model economies where the underlying payoff structure is less interdependent and individual payoffs are actually affected by subgroups of other players (e.g. firms in different industries). An interesting question here concerns whether population learning is able to recover the underlying network structure, both from an aggregate point of view (i.e.: Does the long-run average number of links converge to K?) and from an individual point of view (i.e.: To what extent all agents hold exactly K links in the long-run?).

Second, one needs a more careful investigation of the robustness of our results with respect to alternative strategy/network updating processes. For example, the introduction of idiosyncratic (low-probability) flips in strategy updating (or more generally of stochastic best-reply rules such as log-linear ones) may allow agents to better explore the environment and avoid lock-ins.

Finally, one might study the effects of introducing (across-agent) heterogeneous, time-varying and/or endogenously changing evaluation rules. For instance: What happens if we split the population in two subsets, one using MIN and the other MAX throughout the entire process? What happens when one introduces exogenous mutation in criteria? And, similarly: Which is the effect of allowing for endogenously changing (e.g. imitation-driven) evaluation rules?

Acknowledgements

Thanks to Christophe Deissenberg, Hans Hamman, Koen Frenken, Dan Levinthal, Nick Vriend, and two anonymous referees, for their precious comments and suggestions. The paper has also benefitted from discussions with participants of the following conferences: "Complex Behavior in Economics: Modelling, Computing, and Mastering Complexity", Aix en Provence, May 2003; "8th Annual Workshop on the Economics with Heterogeneous Interacting Agents (WEHIA)", Kiel, May 2003; and "9th International Conference of the Society for computational economics, Computing in Economics and Finance", Seattle, June 2003. L. M. and M. V. gratefully acknowledge financial contribution from the project NORMEC (SERD-2000-00316), funded by the European Commission, Research Directorate, 5th framework programme.

Notes

[1] See for instance Goyal and Vega-Redondo (2001), Jackson and Watts (2002), Droste, Gilles and Johnson (2000) and Fagiolo (2004).

[2] An interesting example concerns the consequences of assuming stochastic vs. deterministic individual best-reply rules, see Brock and Durlauf (2001).

[3] Cf. local-interaction models where agents are placed on regular lattices and play a game with their nearest-neighbors. See, *inter alia*, Young (1998), Ellison (1993), Nowak and May (1993) and Nowak, Bonnhoefer and May (1994).

[4] This assumption will be relaxed in Section 3.3.2.

[5] We briefly discuss the consequences of employing correlated payoff maps in Section 4.

[6] Typically,$s = -s_i^{t-1}$ and V differs from V_i^{t-1} by a single link addition or removal (see below).

[7] In Section 3.3.3 we explore an alternative specification of the model where agents can only observe (sample) a limited number of payoff values.

[8] Both link formation and deletion require mutual consent. In Section 3.3.1 we introduce an alternative specification where link deletion can be performed unilaterally as in Jackson and Dutta (2000).

[9]Whether agents accept payoff *neutral* updates or not does not affect our results.

[10]Notice that since Π contains $N \cdot 2^N$ i.i.d. realizations of a $U[0, 1]$, the associated payoff distribution is very close to the expected one even for small N (e.g. $N \geq 10$). Therefore, there is no gain whatsoever in allowing for heterogeneous initial payoff landscapes across Montecarlo replications.

[11]These values have been chosen in order to trade-off a sufficiently high variability of the payoff landscape with a reasonable computational burden. All results presented below are not substantially altered by increasing N and M.

[12]Such static networks systems look quite similar to those explored by Kauffman in his NK model (with $K = N - 1$, cf. Kauffman, 1993). Notice however that in the NK model *global* (instead of *individual*) payoff signals drive adaptation: a new configuration is chosen if its *global* (e.g. average) fitness is higher.

[13]More precisely, each set of M replications is characterized by populations whose initial network is drawn at random and each link has the same probability $p \in (0, 1)$ to be in place.

[14]Cycles typically arise in our model whenever networks are static and some competition emerges among linked agents who hold incompatible payoffs. This kind of behavior is quite similar to *frustration* emerging in "spin glasses" models, see Fischer and Hertz (1991). Note, however, that if network updating is allowed agents might want to remove such links. This can indeed destroy cycles.

[15]The size of the subset of initial conditions for which there exists at least a population which finds a unique steady state is very close to zero. Thus, we do not investigate its properties further.

[16]Recall that underlying payoffs are the same for all populations. Therefore, all ith agents are identical as far as their payoff vector π_i is concerned.

[17]Notice that in this case network updating implies some form of selfish behavior on the side of the agents, because anyone can unilaterally refrain from interacting with their partners. Indeed, for our agents choosing an interaction network means choosing the game to play. Requiring common agreement in both link formation and link deletion (as we do in the basic model) implies that agents commit themselves to play a certain game until there is no common interest in playing another game.

[18]To get the corresponding rule under *non-neutral* TBR simply replace weak with strict inequalities.

[19]Notice that we are implicitly assuming that each agent earns a net payoff where total networking costs are linear in the number of links it holds. Since gross payoffs (before networking costs) are not monotonically increasing with the number of links even if c is very small, we are not positing any positive network externality effect *à la* Goyal and Vega-Redondo (2001). More specifically, allowing for a positive link cost in our model introduces strong *negative* network externality effects. Indeed, the higher c, the more agents prefer to keep a smaller number of links. See Fagiolo (2004) for a complementary perspective in coordination games.

References

Blume, L. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behaviour*, **5**, 387–424.

Brock, W. and Durlauf S. (2001). Discrete choice with social interactions. *Review of Economic Studies*, **68**, 235–260.

Droste, E., Gilles, R. and Johnson, C. (2000). Evolution of conventions in endogenous social networks. Unpublished Manuscript, CentER, Tilburg University, The Netherlands.

Ellison, G. (1993). Learning, local interaction and coordination. *Econometrica*, **61**, 1047–1071.

Fagiolo, G. (1998). Spatial interactions in dynamic decentralized economies: A review. In: P. Cohendet, P. Llerena, H. Stahn, and G. Umbhauer (eds.), *The Economics of Networks. Interaction and Behaviours*. Berlin-Heidelberg, Springer Verlag.

Fagiolo, G. (2004). Endogenous neighborhood formation in a local coordination model with negative network externalities. *Journal of Economic Dynamics and Control*. Forthcoming.

Fischer, K. and Hertz, J. (1991). *Spin Glasses*. Cambridge University Press, Cambridge.

Goyal, S. and Vega-Redondo, F. (2001). *Learning, Network Formation and Coordination*. Erasmus University, Rotterdam. Unpublished.

Jackson, M. and Dutta, B. (2000). The stability and efficiency of directed communication networks. *Review of Economic Design*, **5**, 251–272.

Jackson, M. and Watts, A. (2002). On the formation of interaction networks in social coordination games. *Games and Economic Behavioral*, **41** 265–291.

Kauffman, S. (1993). *The Origins of Order*. Oxford, Oxford University Press.

Kirman, A. (1997). The economy as an evolving network. *Journal of Evolutionary Economics*, **7**, 339–353.

Nowak, M., Bonnhoefer, S. and May, R. (1994). More spatial games. *International Journal of Bifurcation and Chaos*, **4**, 33–56.

Nowak, M. and May, R. (1993). The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos*, **3**, 35–78.

Taylor, M. (1987). *The Possibility of Cooperation*. Cambridge University Press, Cambridge.

Young, H. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.