

Fagiolo, Giorgio; Marengo, Luigi; Valente, Marco

**Working Paper**

## Endogenous networks in random population games

LEM Working Paper Series, No. 2003/03

**Provided in Cooperation with:**

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies

*Suggested Citation:* Fagiolo, Giorgio; Marengo, Luigi; Valente, Marco (2003) : Endogenous networks in random population games, LEM Working Paper Series, No. 2003/03, Scuola Superiore Sant'Anna, Laboratory of Economics and Management (LEM), Pisa

This Version is available at:

<https://hdl.handle.net/10419/89379>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



Laboratory of Economics and Management  
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà - I-56127 PISA (Italy)

Tel. +39-050-883-341 Fax +39-050-883-344

Email: lem@sssup.it Web Page: <http://www.sssup.it/~lem>

# LEM

## Working Paper Series

*Endogenous Networks in Random  
Population Games*

*Giorgio Fagiolo*<sup>\*</sup>  
*Luigi Marengo*<sup>°</sup>  
*and*  
*Marco Valente*<sup>+</sup>

<sup>\*</sup> *Sant Anna School of Advanced Studies, Pisa (Italy)*

<sup>°</sup> *University of Teramo, DSGSS, Teramo (Italy)*

<sup>+</sup> *University of L'Aquila, Faculty of Economics, L'Aquila (Italy)*

**2003/03**

**January 2003**

# Endogenous Networks in Random Population Games

Giorgio Fagiolo\*      Luigi Marengo†      Marco Valente‡

January 30, 2003

## Abstract

Population learning in dynamic economies has been traditionally studied in over-simplified settings where payoff landscapes are very smooth. Indeed, in these models, all agents play the same bilateral stage-game against any opponent and stage-game payoffs reflect very simple strategic situations (e.g. coordination). In this paper, we address a preliminary investigation of dynamic population games over ‘rugged’ landscapes, where agents face a strong uncertainty about expected payoffs from bilateral interactions. We propose a simple model where individual payoffs from playing a binary action against everyone else are distributed as a i.i.d.  $U[0,1]$  r.v.. We call this setting a ‘random population game’ and we study population adaptation over time when agents can update both actions and partners using deterministic, myopic, best reply rules. We assume that agents evaluate payoffs associated to networks where an agent is not linked with everyone else by using simple rules (i.e. statistics) computed on the distributions of payoffs associated to all possible action combinations performed by agents outside the interaction set. We investigate the long-run properties of the system by means of computer simulations. We show that: (i) allowing for endogenous networks implies higher average payoff as compared to “frozen” networks; (ii) the statistics employed to evaluate payoffs strongly affect the efficiency of the system, i.e. convergence to a unique (multiple) steady-state(s) or not; (iii) for some class of statistics (e.g. MIN or MAX), the likelihood of efficient population learning strongly depends on whether agents are change-averse or not in discriminating between options delivering the same expected payoff.

**Keywords:** Dynamic Population Games, Bounded Rationality, Endogenous Networks, Fitness Landscapes, Evolutionary Environments, Adaptive Expectations.

**JEL Classification:** C72, C73, D80.

---

\**Sant’Anna School of Advanced Studies*, Laboratory of Economics and Management (LEM), Piazza Martiri della Libertà 33, 56127 Pisa, Italy, e-mail: [fagiolo@sssup.it](mailto:fagiolo@sssup.it)

†*University of Teramo*, DSGSS, Loc. Colleparco, 64100 Teramo, Italy, e-mail: [marengo@unite.it](mailto:marengo@unite.it)

‡*University of L’Aquila*, Facoltà di Economia, Piazza del Santuario 19, 67040 Roio Poggio, L’Aquila, Italy, e-mail: [mv@business.auc.dk](mailto:mv@business.auc.dk)

# 1 Introduction

The last decade has witnessed the emergence of a large class of models aiming at understanding the evolution of behaviors in a population of interacting economic actors (cf. e.g. Kirman (1997) and Fagiolo (1998)). A typical exercise consists in modelling a decentralized economy as composed of many entities (consumers, firms, etc.) who repeatedly interact over time. Interactions are generated by some form of externalities because the decision problem of any single agent takes as parameters the behaviors of other agents in the population (see Blume & Durlauf (2001)). These parameters might be interpreted as individual expectations. If the latter are modelled as myopic (or adaptive), individual decision problems depend on past (observed) behaviors of others actors in the system. Therefore, any aggregate measure of the behavior of the system (e.g. some statistics computed on individual choices) will follow a Markovian process (cf. Brock & Durlauf (2001)).

Within this framework, much effort has been devoted to the study of repeated, dynamic, population games, cf. Blume (1993). A dynamic population game is an interaction-based description of a decentralized economy where agents play non-cooperative, simple games against other agents in the population. Interactions (and externalities) are modelled through payoff matrices describing the outcome of bilateral games. Agents hold myopic or adaptive expectations about the behavior of their opponents in the games and employ simple (boundedly rational) decision rules to choose the strategy to perform. Examples range from deterministic and myopic best-response to stochastic rules (e.g. log-linear rule, best-reply with noise, etc.). A key assumption is that behaviors are reversible. Therefore, agents are allowed, from time to time, to revise their current choice on the basis of the observation of actions performed by their opponents in the past.

These models allows one to address three related questions: (i) Under which conditions the system converges to some equilibria (either as a steady-state or as some statistical regularity) in the long-run? (ii) Which efficiency properties do long-run aggregate outcomes

possess?<sup>1</sup> And: (iii) To what extent direct interactions, as well as individual rationality, affect population dynamics and long-run properties?

In this perspective, two broad classes of models might be singled out. First, many scholars have been attempting to study population games where the *interaction structure* (i.e. the map that defines who plays the game with whom at each point in time) is not allowed to evolve through time. Here, the basic exercise involves assessing how different interaction structures are able to affect the long-run behavior of the system. Interaction structures might range from *global* ones (i.e. each agent has a positive probability of playing the game with any other agent in the population) to *local* ones (i.e. each agent always plays the game with her ‘nearest neighbors’). In the latter case, players are placed in some metric space (e.g. regular lattices) and only interact with agents located in their neighborhood. The spatial dimension of the economy is taken to reflect some underlying socio-economic dissimilarity, defined in a space of unobserved variables that are assumed to change very slowly as compared to the pace at which individual actions are allowed to be revised. In this case, the assumption of static interaction structure is justified because the frequency at which agents are allowed to update their partners in the game is so small that this additional revision process would not affect the properties of the one governing strategy updating<sup>2</sup>.

Second, a complementary and more recent line of research has originated from the observation that, whenever the frequency at which agents are allowed to revise the choice about their interacting partners is at least comparable to that at which they update their strategy in the game, it becomes crucial to study in a co-evolutionary manner the interplay between the two revision processes (cf. Goyal & Vega-Redondo (2001), Jackson & Watts (2000), Droste, Gilles & Johnson (2000) and Fagiolo (2001)).

---

<sup>1</sup>For instance, in population games where all agents play a coordination game, one is often interested in studying whether the system converges to a configuration where all agents play the Pareto-efficient Nash equilibrium vs. a risk-dominant one, cf. Kandori, Mailath & Rob (1993).

<sup>2</sup>Kandori et al. (1993), Young (1996) and Young (1998) study dynamic population games with static interaction structures where players employ stochastic decision rules. On the contrary, spatial population games with local and static interaction structures are investigated in Ellison (1993), Young (1998), Nowak, Bonhoefer & May (1994) and Nowak et al. (1994).

Models which adopt the latter perspective describe dynamic settings where agents have the option of repeatedly updating (whether simultaneously or not) both the strategy to be played in the game against their current partners and the set of partners in the game<sup>3</sup>. Network updating becomes endogenous and often occurs on the basis of expected payoffs in different alternative networks. Once again, a crucial assumption concerns whether agents are able to freely select any other agent in the population (i.e. a sort of global matching process) or not (e.g. they can only locally adapt the set of their interacting partners in some space endowed with a metrics defined on the basis of some underlying, slow-moving, variables).

Notwithstanding all that, the entire body of models studying dynamic, repeated, population games, while concentrating its attention almost exclusively on the roles played by the amount of rationality to be imputed to economic agents and the structure of (either static or endogenous) interactions, has paid virtually no attention in exploring economies where the payoff landscape generated by individual stage-games played by the agents is not ‘smooth’. In fact, the literature has extensively studied population games where bilateral stage-game payoffs reflect very simple, strategic situations (e.g. coordination, ‘prisoner dilemma’, ‘hawk-dove’, etc.) which are directly interpretable and, so to speak, very easy to learn both by individuals and by the population. In all these settings, individual payoffs are common knowledge (there is no uncertainty whatsoever about payoffs) and each agent plays the same game against any other agent in the population. Moreover, the learning process (albeit rudimentary) acts at a population-level on smooth landscapes where the payoff of any single agent depends on some average levels of the behaviors of players belonging to her interacting set. For instance, in population coordination games, individual payoffs are a linear function of the number (or the frequency) of coordinated agents in the network. This implies that individual payoffs are invariant to permutations which preserve the frequency of agents currently playing a given strategy in the network. As a result, the

---

<sup>3</sup>See also Goyal & Janssen (1997), Skyrms & Pemantle (2000) and Mailath, Samuelson & Shaked (2000). Dynamic models of non-cooperative network formation only (i.e. without simultaneous choice of a strategic variable) are studied in Bala & Goyal (2000), Watts (2001) and Jackson & Watts (2002).

payoff landscape is relatively smooth because it does not change very much across configurations that are characterized by the same distribution of local frequencies (e.g. number of agents playing coordinated actions across different networks).

On the contrary, in many real-world settings agents face a strong uncertainty about ‘which game to play with whom’ in any time period and, consequently, about the payoff that they might expect from any bilateral interaction. If the ‘type’ of each prospective opponent is (at least at the beginning of the process) unknown and agents are allowed to change the game they play both over time and across different meetings, the metaphor of a smooth payoff landscape with homogenous stage-game payoff matrices might be misleading (for a quite similar perspective cf. Bednar & Page (2002), Calvert & Johnson (1997) and Taylor (1987)). Indeed, in such situations, one typically observes strong heterogeneity of stage-game payoffs matrices and, in turn, a high variability of payoffs experienced by agents after each bilateral game. More specifically, expected (and realized) payoffs of each economic agent might be extremely sensitive to small changes in the configuration of actions currently performed by actors in her network.

In this paper, we propose a preliminary model that attempts to capture population game settings where agents face high uncertainty about expected payoffs from bilateral interactions. We argue that, at least as a first approximation, individual payoff from playing a certain strategy (given the current configuration of population choices) may be modelled as being i.i.d. random variables. In particular, we assume that if an agent interacts with everyone else in the population, the payoff she receives (conditional to any possible combination of actions performed by the others) is distributed as a uniform (i.i.d.) random variable with support  $[0, 1]$ . We call this setting a ‘random population game’.

We choose not to model the process through which agents learn e.g. how to ‘play the right game’, as they discover the ‘types’ of their opponents (see Bednar & Page (2002) for an alternative approach based on ‘individual learning’). On the contrary, we study the process through which the population adapts over time, when agents are allowed to *both* adjust their actions *and* their network (i.e. their opponents in the game). More

specifically, we assume that from time to time agents might either delete single links that they currently maintain active with other agents or add single, new links with currently disconnected agents<sup>4</sup>. We assume that agents hold myopic expectations (i.e. based on last-period observation) and employ deterministic best reply rules (i.e. maximization of expected payoffs) to choose the strategy to play and whether adding or deleting links. Furthermore, we assume that maintaining a link is costless and that both link addition and link deletion require mutual consent. Finally, we allow payoff tie-breaking rules to account either for *change-adverse* players (who always stick to current choice when a tie occurs) or *change-lover* players (who always accept to change even if their payoff is the same).

In order to compute expected payoffs associated to local networks implying a number of links smaller than the maximum one (i.e. a player *not* interacting with all the others), we suppose that players use simple statistics computed on the distribution of payoffs associated to all possible combinations of actions performed by agents outside her local network. In particular, we will study four simple statistics: average of such payoffs (MEAN henceforth), their maximum value (MAX), their minimum value (MIN) and a random draw thereof (RND). These rules (or criteria) might be interpreted as heuristics employed by individuals to form (myopic) expectations (or more generally to learn or forecast) on the payoff associated to radically new environments (i.e. when the network is different from the complete one). Along the same lines, a rule may be understood as the way agents cope with the uncertainty of the system: using a given rule might then represent the way agents synthesize the information provided by the distribution of payoffs associated to the complete network. For example, if an agent uses a MIN rule to compute the expected payoff of the network associated to a link deletion, she might be labelled as being ‘conservative’ or ‘pessimistic’ (and vice versa if she employs the MAX rule).

In turn, the above criteria have a direct interpretation in terms of the trade-off be-

---

<sup>4</sup>Network updating rules is quite similar to that in employed in Jackson (2001), Watts (2001) and Jackson & Watts (2000). Unlike these models, however, we impose mutual consent in both link addition and deletion.



tween mean and spread of experienced payoffs. Since the cardinality of the set of payoffs associated to all possible choices of agents outside the network *increases* as the number of her links *decreases*, the distribution of the expected payoffs of an agent using e.g. a MIN criterion becomes *more concentrated* around a *decreasing mean* as the agent shrinks her network. Conversely, expected payoffs of an agent using e.g. a MAX criterion become *more concentrated* around an *increasing mean* as she keeps deleting links. Therefore, different criteria will have non trivial consequences on individual (and population-wide) fitness landscapes.

We study the long-run behavior of the model in different setups. *First*, as a benchmark exercise, we explore the economy where *networks are exogenously given and time-invariant*<sup>5</sup>. Our original contribution here is to study what happens to aggregate statistics (average payoffs, etc.) when one compares systems characterized by an increasing average number of links (i.e. from very disconnected and sparse networks toward complete, fully-connectivity ones) and different payoff criteria (MIN, MAX, etc.). Computer simulations show that if the average number of links is sufficiently large, then no steady-state is ever reached. All populations continue to explore the landscape and serial correlation between average payoffs is not statistically different from zero. The long-run relationship between the distribution of individual payoffs and number of links held by agents is driven (for any payoff rule) by the trade-off between average and spread induced by any chosen criterion. Since the set of payoffs associated to all possible choices of agents outside the network becomes larger as the number of links an agent maintains decreases, then agents holding a lower (higher) number of links get a payoff with mean 0.5 and highly concentrated (highly dispersed).

Second, we explore random population games with endogenous networks. We run Montecarlo exercises to investigate the effect of initial conditions (i.e. initial network and strategy configurations), payoff criteria and tie-breaking rules on long-run average

---

<sup>5</sup>In this case, the model has a structure quite similar (but also with some substantial differences, see below) to that of Kauffman's NK class of formalizations (see Kauffman (1993)).

outcomes. We show that both the long-run behavior of the system (e.g. convergence to steady-states) and its short-run dynamic properties are strongly affected by: (i) the payoff rule employed; (ii) whether players are change-adverse or not. We find that if agents use the MEAN rule, then, irrespective of the change-aversion regime, the system displays multiplicity of steady-states. Populations always climb local optima by first using action- and network-updating together and then network-updating only. Climbing occurs through successful adaptation and generates long-run positive correlation between number of links and average payoffs. With MIN or MAX rules, the long-run behavior of the system is instead affected by whether players are change-adverse. If they are, and employ the MIN rule, then the network converges to a steady-state where all agents are (almost) fully connected but strategies are not, so that average payoffs oscillate. If agents employ the MAX rule then the system displays many steady-states (in both networks and actions) characterized by few links and different levels of average payoff. Finally, if agents are change-lovers, then the population can explore a larger portion of the landscape. Therefore, with agents using the MIN rule, the network will converge to the complete one, but from then on exploration on strategies will go on forever. If they employ the MAX rule, then the system will display a unique optimum. All populations converge to the same payoff distribution but neutral NU will continue forever (without affecting realized payoffs).

The rest of the paper is organized as follows. In Section 2 we briefly discuss some relevant pieces of literature and we introduce informally the model presented in Section 3. An overview of simulation results is contained in Section 4. Finally, Section 5 concludes discussing extensions of the model and future research.

## 2 Population Games, Payoffs and Endogenous Networks

In the last years, a large body of literature has been extensively studying the outcome of decentralized decisions undertaken by large populations of boundedly rational agents who repeatedly and directly interact over time. Within this large class of formal models, many authors have been focusing on *dynamic population games*, i.e. games played over time by large populations boundedly rational players. In what follows, we will first briefly survey this vast literature. Next, we will discuss some motivations underlying the model which we will informally describe at the end of this Section.

### 2.1 Dynamic Population Games: A Brief Survey

The standard framework common to dynamic population games consists of a set of  $N$  individuals who play games in discrete time. In any time period  $t \geq 1$ , each individual  $i$  plays a strategy (or strategy)  $s_i^t \in S$ . Assume for the sake of simplicity that  $S = \{-1, +1\}$  and that whenever two agents (say  $i$  and  $j$ ) meet, they play a bilateral game whose (symmetric) payoff matrix (to agent  $i$ ) is given by:

$$G = \begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

with  $a, b, c \in \mathbb{R}$ .

In this setup, all agents play the same game, they know that also the others will do the same and  $G$  is common knowledge. Finally, assume that at each point in time, any agent  $i$  only meets agents  $j$  belonging to the set  $V_i^t \subseteq I$  (i.e. plays bilateral games defined by  $G$  against all and only agents in  $V_i^t$ ). The collection  $(V_i^t)_{i \in I}$  is called (time  $t$ ) *interaction structure*. The dynamics is typically defined as follows. At any  $t$ , an agent (say  $i$ ) is chosen at random from  $I$ <sup>6</sup> to revise her current state (i.e.  $s_i^t$ , or  $V_i^t$ , or both). Agent  $i$

---

<sup>6</sup>This scheme is known as *asynchronous* updating. The consequences of assuming *synchronous* updating

then forms myopic expectations about her next-period payoff under different alternatives (e.g. actions and/or interaction structures) and employs a decision rule to choose her next-period state (i.e. either  $s_i^{t+1}$ , or  $V_i^{t+1}$ , or both). Decision rules typically differ as to whether they introduce some idiosyncratic noise (interpreted either as the possibility of experimentation or mistakes) or not. In the latter case, the rule is deterministic and is usually based on local best-reply dynamics. For instance, if  $V_i^t$  are exogenously given and do not change over time, agents will pick their next-period strategy by employing the rule:

$$s_i^{t+1} = \arg \max_{s \in \{+1, -1\}} w(s; s_j^t, j \in V_i^t),$$

where total individual payoffs to  $i$  from interacting with agents in the network - i.e.  $w(s; s_j^t, j \in V_i^t)$  - are simply defined as the sum (or the average) of all payoffs from bilateral games. Introducing a stochastic term might then reverse, with some small probability, the decision that maximizes local payoffs. The size of a ‘mistake’ may in turn be either constant as the relative frequencies of players choosing  $+1$  or  $-1$  in each  $V_i^t$  change (cf. *noisy best-reply* rules used in Ellison (1993)) or state-dependent (e.g. the *log-linear* rule, where the probability of choosing against the majority becomes very small - although not null - as the “size of the majority” gets larger, cf. Brock & Durlauf (2001) and Blume (1993)).

A dynamic population game is therefore completely defined once one specifies: (i) the payoffs in the matrix  $G$ ; (ii) the interaction structure in place at each given point in time (or the rule that governs how the interaction structure changes over time).

With some minor exceptions<sup>7</sup>, the vast majority of models have been focusing on either *coordination* or *prisoner dilemma* games with *static* (either global or local) interaction structures (i.e. games where  $V_i^t = V_i$ , all  $t \geq 1$ ). More recently, however, a few contribu-

---

schemes (i.e. where agents are allowed perform a parallel updating) or *incentive-based* updating schemes (where which agent is allowed to revise her choice depends on the state of the system) are studied in Page (1997).

<sup>7</sup>Cf. for instance Nowak & May (1993), Nowak et al. (1994) and Herz (1994). See also Jackson & Watts (2000) in the context of endogenous networks.

tions have pointed out that endogenous network dynamics may play a non-trivial role in shaping long-run equilibrium patterns and their efficiency properties<sup>8</sup>.

For example, in the context of coordination games, ‘static interactions’ literature has been able to provide, in an evolutionary game perspective, a quite robust microfoundation for equilibrium selection. For instance, whenever the underlying  $2 \times 2$  game  $G$  has two Nash equilibria, one of which is Pareto efficient and the other is risk-dominant, Kandori et al. (1993), Young (1996), Blume (1993) and Ellison (1993) have shown that the unique long-run equilibrium is the risk-dominant one and that local interactions can speed up the rate of convergence (as compared to global ones). Convergence to an efficient outcome is however the case either when non-exclusive conventions are assumed (so that agents can pay to remain flexible, i.e. choosing not to choose, cf. Goyal & Janssen (1997)); or when players are mobile (cf. Bhaskar & Vega-Redondo (1996), Ely (1996), Oechssler (1997) and Dieckmann (1999))<sup>9</sup>; or, more importantly, when interaction structures are fully endogenized and agents are free to select as a partner any other agent in the population (cf. Goyal & Vega-Redondo (2001)). On the contrary, when geographical barriers prevent individuals to build links with anyone else in the system, population learning converges with a higher likelihood to risk-dominant equilibria (see Fagiolo (2001)).

Results in a similar vein have been obtained also in the context of dynamic population games where agents play ‘prisoner dilemmas’<sup>10</sup>. No matter whether interactions are

---

<sup>8</sup>This class of contributions posits dynamic population games where agents, from time to time, may have access to a network updating decision, often modeled as a best-reply rule (either deterministic or noisy). A crucial ingredient is whether agents (when choosing their next-period network  $V_i^{t+1}$ ) are able to compare expected payoffs from every possible network (cf. Goyal & Vega-Redondo (2001)) or they just have the option of adding/deleting a small number of links in any choice-stage (cf. Jackson & Watts (2000)). In this latter case, agents might be allowed to choose with a positive (although not necessarily homogeneous) probability any other agent in the population as a new partner (cf. Droste et al. (2000)) or being constrained in choosing their network by some underlying geographical structure that restricts networks to become neighborhoods (see Fagiolo (2001)).

<sup>9</sup>In such models, partner selection is to some extent endogenized by assuming the existence of a fixed number of spatial locations. Players are mobile and can *indirectly* select their future partners by picking the place they want to move to, on the basis of the expected net payoff of each location.

<sup>10</sup>Dynamic prisoner dilemma (DPD) population games with static interaction structures have been studied in Axelrod (1984), Herz (1994), Nowak & May (1993), Nowak et al. (1994), Oliphant (1994), Oltra & Schenk (1998) and Tieman, Houba & Van Der Laan (1998). Population games where agents play DPD and can choose not to interact with an opponent (i.e. PD with ostracism or refusal) are instead investigated by Hirschleifer & Rasmusen (1989), Kitcher (1993), Smucker, Stanley & Ashlock (1994), Stanley, Ashlock

static (globally or locally) or endogenously evolving, iterated prisoner dilemma played by myopic agents allows cooperation to be sustained (unlike in standard game-theoretic models). Furthermore, if players have the option to refuse interactions with other individuals (i.e. if interaction structures become to some extent endogenous), the population tends to cooperate more than in the case of a compulsory prisoner dilemma (cf. also Zimmermann, Eguíluz & San Miguel (2001) and Hanaki & Peterhansl (2002)).

Despite these promising results, the literature on dynamic population games (with both exogenous and endogenous networks) has been largely neglecting the analysis of decentralized economies where agents face more complicated payoff structures. Indeed, the baseline model of a dynamic population game shares (at least) the following strong assumptions: (a) all agents know that everyone is going to play the same game; (ii) stage-game payoffs are common knowledge. This implies that the payoff landscape over which population learning takes place is quite ‘smooth’ (i.e. individual payoffs display a small variability as one slightly changes the frequency of agents in the network who currently play the same strategy). Therefore, the long-run behavior of the system is directly interpretable in terms of bilateral games played by individuals. For example, in coordination population games, the purported across-agent homogeneity of stage-games, as well as the nature of the game to be played, together imply that the payoff to any agent  $i$  playing, say, +1 is a linear function of the number of agents in  $i$ ’s network currently playing +1. Thus, individual payoffs are invariant to any permutation of individual choices that keeps local frequency constant. Hence, the aggregate state of the system (e.g. all playing +1) can be interpreted as a state of maximum coordination.

## 2.2 The Model: An Informal Description

In many real-world settings that we might conceive, however, the situation is quite different. First, an agent is not always aware in advance about which game her opponent will play in the next meeting. Second, and partly as a consequence, stage-game payoffs might be

---

& Tesfatsion (1994), Ashlock, Smucker, Stanley & Tesfatsion (1996) and Hauk (1996).

highly uncertain and, possibly, endogenously changing.

In order to model such complicated environments, one could choose to describe learning processes from an *individual* or a *population* perspective. In the first case, the goal would be the investigation of the properties of adaptive individual learning about: (i) the type of game an opponent will play tomorrow; (ii) the type of game to play against a given opponent; and: (iii) selecting opponents that are expected to play the ‘right’ game.

In this paper, on the contrary, we focus on *population* learning dynamics. We model environments where network payoffs are modeled so as to capture (albeit in a rudimentary way) strong across-agent interdependencies and high sensitivity of payoff landscapes (with respect to small changes in the current configuration), generated by an extreme uncertainty about the game played by other agents in the network (and, consequently, by a large across-time volatility of stage-game payoffs). More precisely, we choose to model payoff landscapes in the following way. Suppose to have a population of  $N = 3$  agents (i.e.  $I = \{1, 2, 3\}$ ) and that the strategy space is binary:  $s_i^t \in \{-1, +1\}$ . Agents live (or *think* to live) in a world where everyone is in principle connected with anyone else. They expect their payoff to change in unpredictable ways whenever a small change in the current configuration of the system occurs (i.e. an agent changes either her strategy, or her partners, or both). Consider first the payoff landscape defined over the complete-network population game. Since the payoff of any agent  $i$  depends on the choices of the remaining  $k_i = 2$  agents, each one faces  $2^{N-k_i+1} = 4$  possible configurations of the world for each  $s_i^t \in \{-1, +1\}$  and thus  $2 \cdot 2^{N-k_i+1} = 8$  individual payoffs  $\pi_i(s_1^t, s_2^t, s_3^t)$ . In order to model the extreme underlying uncertainty characterizing our system, we suppose that  $\pi_i(s_1^t, s_2^t, s_3^t)$  are i.i.d., uniformly distributed (over the unit interval), random variables<sup>11</sup>. An example of a complete payoff

---

<sup>11</sup>If we let the system adapt using an asynchronous updating mechanism and individual best-reply decision rules, population dynamics is similar to standard adaptation over rugged fitness landscapes explored by Kauffman’s NK model where  $K = N - 1$ , cf. Kauffman (1993). However the latter uses global payoff (fitness) signals to drive adaptation: a new configuration is chosen if its global fitness is higher. On the contrary, as we will see below, we use local payoff criteria: a link is established or deleted - and a strategy is switched - if the payoff of the agent(s) directly involved gets higher, regardless what happens to the rest of the population. This difference has important consequences on the dynamics and in particular on the likelihood of lock-in into local optima (see below). Finally, note that in Kauffman’s model the assumption of complete connectivity is taken to reflect the highest possible level of ruggedness of the fitness landscape

landscape is in Table 1.

$(s_1^t, s_2^t, s_3^t)$	$\pi_1$	$\pi_2$	$\pi_3$
$(-1, -1, -1)$	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

Table 1. An Example of the Payoff Landscape with  $N = 3$ .

Experienced payoffs depend on how many links an agent actually maintains with the others. We suppose that agents can hold bilateral, costless, links with any other agents in the population. A player may hold any number of links between 0 (isolated agents) and  $N - 1$  (complete connectivity). Any two agents may then hold a different number of links (i.e. the resulting graph is not homogeneous). Therefore, given the payoff landscape  $\pi_i(s_1^t, s_2^t, s_3^t)$ , each agent is fully characterized in each time period by her current choice  $s_i^t$  and the set  $V_i^t \subseteq I$  containing all her current partners (i.e. agents with whom she currently maintain a link).

If a link between  $i$  and  $j$  exists, experienced payoffs of both  $i$  and  $j$  are affected by whether the other chooses  $-1$  or  $+1$ . Therefore, if agent  $i = 1$  is connected with both 2 and 3 and currently plays, say,  $s_1^t = -1$ , she will face a payoff which changes as soon as either

---

over which population adaptation takes place.



one of her partners switches her strategy (according to  $\pi_i(s_1^t, s_2^t, s_3^t)$  entries in Table 1). If an agent, conversely, is not connected with all  $N - 1$  other agents in the population, she will use a simple statistics in order to compute expected payoffs from playing a given strategy. This statistics will be computed on payoffs associated to all combinations of strategies that could be played by agents outside the current network. For example, suppose that  $s_1^t = -1$  and that the statistics employed by agent 1 is the arithmetic mean (MEAN)<sup>12</sup>. If 1 is currently connected with 2 only, then she faces only two possible payoffs (according to what her partner plays). Each payoff will be then computed as the MEAN of payoffs associated to any possible choice for  $s_3^t$ . Thus:

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

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

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

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

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

Suppose that at time  $t = 0$  a random initial configuration  $(s_i^0, V_i^0)_{i=1}^N$  is given. We study a population dynamics governed by an asynchronous updating process both on strategies and (bilaterally) on local networks. More precisely, in each  $t \geq 1$  a pair of agents (say  $h$  and  $k$ ) is firstly randomly drawn to attempt graph updating given current strategy configuration. If a link connecting  $h$  and  $k$  already exists (respectively, does not exist), the link is tentatively removed (respectively, tentatively added). Once again, decisions are made

---

<sup>12</sup>Alternative choices might be: MAX, MIN or RND (i.e. picking at random one of the  $2^{N-k_i+1}$  payoff entries). See below.

by deterministically best-responding to current local configurations under the proposed alternatives. We consider two different tie-break rules. In the first one (henceforth, network updating *without neutrality*), a link is removed (added) if and only if both agents are strictly better-off under the change. This implies that agents are change-averse, because they prefer to stick to their current choices unless the alternative option is associated to strictly larger payoffs for both. In the second one (henceforth, network updating *with neutrality*), a change is accepted even if it implies the same payoff that both agents were experiencing before the change. This implies that agents are change-lovers, because they prefer to change their current choices even if the new one is associated to the same payoff for both<sup>13</sup>.

Secondly, an agent (say  $i$ ) is firstly drawn at random to update her current strategy  $s_i^t$  given the network configuration obtained in the first stage. We assume that agent  $i$  will best reply deterministically to the current local configuration (i.e. she chooses the strategy  $s$  that strictly maximizes  $\pi_i(s; s_j^t, j \in V_i^t)$ , while keeping  $s_i^t$  if a tie occurs).

For instance, suppose that the current configuration is  $(+1, +1, +1)$ . 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 her current payoff  $\frac{1}{2}(0.04 + 0.80)$  with the payoff after link addition, namely 0.80. Agent 3 will compare her 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 better off (strictly) under the change and the link will be added. After network updating, player 3 is connected with 1 only. If she were called to strategy updating, she would switch to  $-1$  since  $\frac{1}{2}(0.70 + 0.78) > \frac{1}{2}(0.25 + 0.44)$ .

These simple updating rules define a population dynamics whose long-run outcome (absorbing state, cycle, etc.) and its efficiency properties (e.g. aggregate payoff, across-agents distribution, etc.) might be strongly affected by the choice of the criterion employed

---

<sup>13</sup>The term *neutrality* refers to the fact that under the associated tie-breaking rule agents accept *neutral* changes (i.e. they add or delete a link even if they both continue to get the same payoff). Notice also that decisions are made on the basis of current payoffs (i.e. those observed by agents at the time of the choice). This implies that we are assuming agents holding myopic expectations.

to form myopic expectations over networks different from the complete one (as well as by whether updating rules allow for *neutrality* or not).

After having more formally described the model (see next Section), we will turn to explore if (and under which conditions) this preliminary conjecture holds true by means of extensive computer simulation exercises.

### 3 The Model

Consider a fixed population of agents  $I = \{1, 2, \dots, N\}$ ,  $N \geq 3$ . Time is discrete, i.e.  $t = 0, 1, 2, \dots$ . At any time  $t$ , an agent  $i \in I$  is completely characterized by the pair  $(s_i^t, V_i^t)$  where  $s_i^t \in \{-1, +1\}$  is her current strategy in the population game and  $V_i^t \subseteq I - \{i\}$  is the set of agents of her opponents in the game (or her partners). We study a population where the underlying links connecting agents are bilateral (i.e.  $i \in V_j^t \Leftrightarrow j \in V_i^t$ ) and maintaining a link is costless for both agents. Define the size of  $V_i^t$  as  $k_i^t = |V_i^t| \in K = \{0, 1, 2, \dots, N-1\}$  and let  $\bar{V}_i^t = V_i^t \cup \{i\}$  (of course,  $\bar{k}_i^t = k_i^t + 1$ ). The collection of sets  $\{V_i^t, i \in I\}$  induces at any  $t$  a non-directed graph  $G_t \in P(N)$ , where  $P(N)$  is the set of all non-directed graphs over  $I$ . We denote with  $ij \in G^t$  the fact that in the network  $G^t$  agents  $i$  and  $j$  are linked. At any given time, the system is therefore characterized by the pair  $\{\Omega^t, G^t\}$ , where  $\Omega^t = (s_i^t)_{i \in I} \in \{-1, +1\}^N$ .

A “random population game” is defined through the following payoff structure. Let  $u_i^t(\Omega^t, G^t)$  the payoff to agent  $i$  at time  $t$ , given the current state of the system. For any subset  $J \subseteq I$ , define a  $J$ -restricted action configuration by  $\Omega^t(J) = (a_j^t)_{j \in J}$ . We assume that:

$$u_i^t(\Omega^t, G^t) = \pi_i^t(\Omega^t(\bar{V}_i^t)) = \pi_{i,t}(s_i^t; s_j^t, j \in V_i^t),$$

i.e. that the payoff to  $i$  at  $t$  does not change if the strategies currently performed by all agents in  $I$  who are not partners of  $i$  (under the current graph  $G_t$ ) change. To model very uncertain environments, we suppose that if all agents were connected with anyone else (i.e.

$\bar{V}_i^t = I$ , all  $i \in I$ ; or, equivalently, if  $G^t$  were completely connected:  $k_i^t = N - 1, \forall i$ ), then:

$$\pi_i^t(\Omega^t(\bar{V}_i^t)) \sim X,$$

where  $X$  are i.i.d. random variables with p.d.f.  $F$ . In what follows, we will assume that  $X \sim U[0, 1]$ .

To form expectations about payoffs  $\pi_i^t(\Omega^t(\bar{V}_i^t))$  associated to local networks  $\bar{V}_i^t$  s.t.  $\bar{k}_i^t < N$ , we suppose instead that agents use some ‘‘criterion’’ (or statistics)  $\mathfrak{R}$ . More precisely, we posit that  $\mathfrak{R}$  is computed over individual payoffs agent  $i$  would have earned if individual strategies within her network  $V_i^t$  were fixed, while individual strategies of agents  $j \in I - \{V_i^t\}$  were allowed to freely change in  $\{-1, +1\}^{N-k_i^t-1}$ . More formally, let:

$$I = \bar{V}_i^t \cup W_i^t, \bar{V}_i^t \cap W_i^t = \emptyset,$$

where  $\bar{V}_i^t = \{i, j_1, \dots, j_{k_i^t}\}$  and  $W_i^t = \{h_1, \dots, h_{N-k_i^t-1}\}$ . Next, define  $P_i^t \subseteq \{-1, +1\}^N$  as the set of all possible configurations  $\Omega^t = (s_i^t)_{i \in I} \in \{-1, +1\}^N$  for which  $\Omega^t(\bar{V}_i^t)$  are kept constant while we allow  $\Omega^t(W_i^t)$  to freely vary. We assume that if  $k_i^t = 0, \dots, N - 1$  then:

$$\pi_i^t(\Omega^t(\bar{V}_i^t)) = \mathfrak{R}(\pi_i^t(\Omega^t), \Omega^t \in P_i^t),$$

i.e.  $\pi_i^t(\Omega^t(\bar{V}_i^t))$  are computed by employing the criterion (calculating the statistics)  $\mathfrak{R}$  over all possible configurations where only the strategies of agents outside  $\bar{V}_i^t$  are allowed to assume all possible values.

We will suppose throughout that agents employ one out of the following four criteria to evaluate such payoffs: (i)  $\mathfrak{R} = MAX$ ; (ii)  $\mathfrak{R} = MIN$ ; (iii)  $\mathfrak{R} = MEAN$ ; (iv)  $\mathfrak{R} = RND$ . In this last case, agent  $i$  computes her payoff by picking at random one out of all available payoffs.

The model is completely defined once one has described the processes governing individual strategy and network updating. Suppose at time  $t = 0$ , a strategy configuration

$\Omega^0 \in \{-1, +1\}^N$  and a graph  $G^0$  over  $I$  are randomly drawn. We assume that in any time period  $t \geq 1$ , first, agents update networks (given  $\Omega^t$ ); second, agents update strategies given the new (just updated) graph. More specifically, given  $(\Omega_t, G_t)$ , any two agents (say  $i$  and  $j$ , either connected in  $G_t$  or not) are picked at random from  $I$ . They evaluate current payoffs as:

$$\begin{aligned} w_i^t &= \pi_i^t(\Omega_t(\bar{V}_i^t)), \\ w_j^t &= \pi_j^t(\Omega_t(\bar{V}_j^t)). \end{aligned}$$

If  $i$  and  $j$  are not connected ( $ij \notin G_t$ ) then define  $\tilde{V}_i^t = \bar{V}_i^t \cup \{j\}$  and  $\tilde{V}_j^t = \bar{V}_j^t \cup \{i\}$ . On the contrary, if  $i$  and  $j$  are already connected ( $ij \in G_t$ ) then define  $\tilde{V}_i^t = \bar{V}_i^t - \{j\}$  and  $\tilde{V}_j^t = \bar{V}_j^t - \{i\}$ . Accordingly, define payoffs under the proposed change (i.e. addition or deletion):

$$\begin{aligned} \tilde{w}_i^t &= \pi_i^t(\Omega_t(\tilde{V}_i^t)), \\ \tilde{w}_j^t &= \pi_j^t(\Omega_t(\tilde{V}_j^t)). \end{aligned}$$

We suppose that agents decide (bilaterally) whether to add (or delete) the link by simply comparing current payoffs before and after the proposed change and picking the network associated to the largest one (i.e. deterministic, myopic, best-reply). We consider two alternative tie-braking rules:

1. Accept the change if and only if both agents are strictly better off under the change, i.e. add (or delete)  $ij$  if and only if  $\tilde{w}_i^t > w_i^t$  and  $\tilde{w}_j^t > w_j^t$  (we call this rule tie-break *without neutrality*);
2. Accept the change if and only if no agent is strictly worse off under the change, i.e. add (or delete)  $ij$  if and only if  $\tilde{w}_i^t \geq w_i^t$  and  $\tilde{w}_j^t \geq w_j^t$  (we call this rule tie-break *with neutrality*).

Network updating is therefore defined as follows:

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

After network updating, strategy updating takes place given  $G_{t+1}$ . We assume that an agent (say  $i$ ) is drawn at random from  $I$ . Given  $(s_i^t, V_{i,t+1})$  and her current payoff  $w_i^t = \pi_i^t(s_i^t, \Omega^t(V_i^{t+1}))$ , she will switch to  $-s_i^t$  at the beginning of period  $t + 1$  if and only if:

$$\pi_i^t(-s_i^t, \Omega^t(V_i^{t+1})) > w_i^t,$$

(i.e. neutral updates are not accepted).

Let us turn now to provide some preliminary results about the long-run behavior of the process governing the evolution of the pair  $\{\Omega^t, G^t\}$  - and of some statistics thereof.

## 4 Simulating the System: Some Preliminary Results

In this section we present some results which can be obtained by simulating our model. Unless otherwise specified, all results refer to Montecarlo averages across 50 independent simulations characterized by the same parameters (i.e. the random values in the payoff matrix) but different initial states for the agents in the population (actions and networks). In each repetition we generate a population containing 15 agents, each using an independently drawn payoff matrix containing  $2^{15}$  entries, uniformly distributed in the  $[0,1]$  interval, for a total of  $2^{15} \times 15$  random values. As mentioned, such payoff matrices are kept unchanged across all repetitions of the same simulation. Nevertheless, given distributional assumptions on (and the huge number of) payoff values, our results do not qualitatively change if one employs different payoff landscapes across independent sets of simulations.

We first discuss population dynamics in static networks (i.e. where agents can only

update their action). Second, we will allow for local, endogenous, changes in the network architecture and we will present results when agents are assumed to use different statistics in order to evaluate their payoff.

## 4.1 Static networks

In order to appreciate how the connectivity of the network influences agents' payoffs we first run a set of simulations in which agents can only modify their strategies but not their links. We generated 100 random network structures in which networks are initialized (and then kept fixed throughout the rest of the simulation). We let the probability that, taken two agents, a link is established between them to range from 0.01 (almost totally unconnected network) to 1.00 (fully connected one). For each population we perform 10 repetitions with the same network structure and payoffs matrices and randomly varying agents starting actions.

Simulations show that, if the average number of links is sufficiently large, then no steady-state is ever reached and all populations continue to explore the payoff landscape. Figure 1 summarizes the main results when MEAN evaluation statistics is employed. The figure plots long-run payoffs against long-run number of links for all the agents from the 100 populations. The expected value of payoff is always 0.5, but the variance increases with the number of links. Agents with no or few links have always payoffs very close to 0.5, because this value is computed as the average over a large number of random values. As the number of links increases, payoffs become more and more scattered on the entire support  $[0, 1]$ , because the MEAN statistics is computed on smaller and smaller sets of uniformly distributed random values.

The exercises using MIN and MAX payoff rules produce similar results, with the concentration of payoff for unconnected agents concentrated around 0 and 1 for minimum and maximum respectively. In turn, simulations with RND payoff rule exhibit a payoffs spread over all the range  $[0,1]$ , independently on the number of links.

Besides the absolute values of payoffs, the number of links an agent holds also affect how many updates improving the (population) average payoff are performed. In fact, when agents with fewer links switch strategy, they affect only the few agents they are connected to. Strategy changes have therefore little impact on the population average payoff. Conversely, populations of highly connected agents have highly volatile average payoffs, because any strategy mutation modifies the payoff of many agents. In Figure 2 we report the phase diagram for the population average payoff when agents are fully connected<sup>14</sup>. The graph plots the population average payoff at time  $t + 1$  as a function of the same variable at time  $t$ . Since we are assuming agents who are change-averse with respect to action-updating, points on the diagonal indicate all cases in which no strategy switch has taken place. When instead a switch has happened (off-diagonal), the distribution of points shows no statistically significant, different from zero, correlation.

The MIN and MAX evaluation statistics show similar patterns. Of course here the expected payoff distribution for agents holding few links will be concentrated around 0 (MIN) or around 1 (MAX), as Figures 3 and 4 show. Therefore, when static networks are assumed, the long-run relationship between the distribution of individual payoffs and number of links held by the agents is driven (for any payoff rule) by the trade-off between average and spread induced by the chosen criterion.

## 4.2 Dynamic networks

### 4.2.1 MEAN and RND evaluation rules

We now move to the analysis of the full-fledged model and allow agents to adaptively change not only their strategies but also the set of neighbors with whom they interact by adding and deleting connections. We start again by considering the MEAN evaluation statistics and we initialize our simulations with totally unconnected networks and random initial strategies.

---

<sup>14</sup>Again, the graph refers to a simulation with MEAN payoff rule, although practically the same result is obtained also for any other type of payoff rule



In this case we find that, *irrespective of whether neutrality is assumed or not*, the system displays multiplicity of steady-states. As we can see in Figure 5, all populations display a similar dynamic pattern characterized by distinct phases (though the duration of these phases may vary). At the outset, all agents in the population produce successful updating both on strategies and on graphs, by exploring different strategies and networking structures. Second, they stop updating their strategies (i.e. strategy switches stop to provide any improvement), but they continue to make some network modification, adding or removing links. Finally, also network updating stops, with the agents making no further modifications. It is important to notice that payoffs and number of links for the 15 agents in the steady-state configuration widely differ within and across simulations. A typical results shows that the number of links can range from 2 to 14 and payoffs from 0.1 to 0.98. Population averages range between 7.7 to 12.1 for the number links, and between 0.48 to 0.69 for payoffs (cf. Figure 5). Since these heterogeneous stable points reached by our populations differ from one another, the system will displays many ‘local optima’ which create lock-in points for updating rules.

In Figure 6 we report the scatter plot of agents’ payoffs as a function of the number of their links. As compared to the ‘frozen networks’ case (see Figure 1), endogenous network updating allows agents with wider neighborhood sets to reach persistently higher payoffs: here agents adapt so to concentrate on the higher portion of the payoff space.

Notice also that these results are not affected by the initial settings (e.g. fully connected networks). The dynamic behavior of a system based on a RND evaluation rule is instead quite similar to that displayed by static networks. Indeed, a RND rule fails to exhibit a positive relationship between number of links and average payoffs. Moreover, strategy and network updating continue indefinitely and a stable structure is never achieved. This result is consistent with the observation we made when analyzing static networks: agents who employ this payoff evaluation rule have no reason to prefer few to many links, since the payoff is always picked randomly.

### 4.2.2 MIN evaluation rule

A richer set of results can be obtained when one introduces the hypothesis that agents use the MIN payoff evaluation rule, i.e. they compute their payoff as the minimum of the values in their payoff matrix corresponding to the strategies used by the agents with whom there is no direct interaction. We can consider this payoff rule as the one played by “pessimistic” agents, in that they try to maximize the payoff in the worst possible case, considering these as the unobservable strategies (from the non-linked agents).

Contrary to the previous cases, here results heavily depend upon acceptance of neutral network updating (i.e. whether agents add or remove a link although their payoff remains unchanged). In fact, increasing or decreasing the number of links held by an agent respectively narrows or widens the set of payoffs from which the minimum is taken. Quite often, it may indeed happen that adding or deleting one more links of an agent does not modify the minimum value of such a set. Therefore these network changes are “neutral” with respect to the payoff evaluation. Whether such neutral changes are acceptable or not is thus of a paramount importance to the dynamic of the network structure.

Let us begin by examining the case in which a payoff-neutral network updating is always accepted (i.e. agents are change-lovers). Figures 7 and 8 present respectively the time series of the average number of links for the 50 simulated populations and the plot of the average payoff against the number of links. In this case, irrespective of the connectedness of the initial graph, network structure and strategy choices never converge to a steady-state and the long-run population average number of links fluctuates around very high values (between about 11 and 14). Furthermore, some positive correlation between average payoffs and average number of links (although weak) still emerges (see Figure 8). If on the contrary agents do not accept neutral network updating (i.e. they are change-adverse), we observe convergence to multiple network steady-states, but strategies and payoffs never settle.

To see why this happens, notice that while a link between any two agents can never be deleted (as such an operation cannot increase the minimum payoff)<sup>15</sup>, new links can always

---

<sup>15</sup>This is obviously the case because, if an agent removes a link, the pool of payoffs across which the

be formed. Thus, the network becomes quickly highly connected and almost all agents maintain a number of links close to 14 (cf. Figures 7 and 9 for an illustration). When agents accept neutral updates, further changes in the network structure are very likely, but this in turn heavily changes individual payoffs. Therefore, the system never climbs a local optimum. Conversely, if agents are change-adverse as far as network updating is concerned, they typically get stuck in one of the many local maxima. Strategy updating instead never stops, because changes of an agent’s strategy induces considerable changes in the other agents’ payoffs, and in general offers opportunities for some of them to change, in turn, their strategy.

In some sense, our result could be easily interpreted by saying that “pessimistic” agents, who base actions on worst-case considerations, tend to increase more and more their “span of control” in order to reduce the risk of unexpected worst cases. However, this behavior collectively generates a very unstable environment, especially when payoff-neutral network updating is always accepted (i.e. when players can explore a larger portion of the payoff landscape).

### 4.2.3 MAX evaluation rule

When using the MAX payoff evaluation criterion, agents adopt a sort of “optimistic” criterion, as they base their decisions upon the best payoffs which could derive from the behavior of *unobservable* agents (i.e. those outside their own neighborhood). Alike the MIN criterion, also in this case results depend upon whether neutral network updating is accepted or not.

Let us first consider the case in which it is accepted. Here agents tend to develop networks with very few connections, because adding a link can never (strictly) increase the maximum payoff<sup>16</sup>. On the contrary, link deletion might increase the maximum payoff and

---

minimum is computed becomes strictly larger than (and includes) the payoff set associated to the network with that link still in place.

<sup>16</sup>The new payoff is computed as the maximum of a set which is strictly contained in the one associated to the network containing that link. Thus, only if the maximum is still contained in the subset can the link addition be accepted under *neutrality*.

therefore be accepted. As the population tends to produce scarcely connected networks, agents can tune their actions in order to reach the globally optimum payoff. Figure 10 reports the average payoff for the 50 populations and shows that each population always converges to the same highest payoff (very close to 1). Furthermore, if one disaggregates average payoff in each population, it is interesting to notice that all populations converge (for the same payoff matrix) to the same strategy profile and the same (maximum) individual payoffs, from any initial condition. Once the optimal strategy profile has been reached, some payoff neutral network changes are still possible and therefore the network structure never stabilizes. Notice, however, that all network structure over which the population keeps cycling indefinitely are payoff equivalent.

If instead we do not allow agents to make strategy or network changes unless they imply a strictly higher payoff (*non neutrality*), we observe rather different results. First, and trivially, if we begin with an unconnected network, no link is ever established as none of them can be strictly payoff-increasing. If instead we begin with a fully or highly connected network, not accepting payoff neutral changes implies that agents may not climb the entire payoff landscape and may end up locked into local optima. The deletion of a link in fact requires that both agents are strictly better off after deletion: if one of them is indifferent, the link cannot be removed and the other agent may be locked into a suboptimal set of neighbors. Figure 11 plots the average payoffs for the 50 populations given the same payoff matrix. It can be noticed that populations lock into different local optima depending on the initial conditions<sup>17</sup>. Therefore, allowing for payoff neutral network changes implies the emergence of many steady-states (in both networks and actions) characterized by few links and different levels of average payoff.

---

<sup>17</sup>Average payoffs are anyway very high because the MAX rule is used.

## 5 Conclusions

Population learning in dynamic economies has been traditionally studied in over-simplified settings where all agents play the same bilateral stage-game against any opponent and stage-game payoffs reflect very simple strategic situations (e.g. coordination). In this paper, we have addressed a preliminary investigation of dynamic population games over ‘rugged’ landscapes, where agents face a strong uncertainty about expected payoffs from bilateral interactions. We have proposed a very simple model where payoff landscapes are modeled through ‘random population games’ and the population adapts (over both strategies and network structures) using deterministic, myopic, best reply rules. The key assumption of the model concerns how agents evaluate payoffs associated to networks which imply “local” interactions. We have explored settings where players use very simple statistics (such as MIN, MAX and MEAN) which are computed on the distributions of payoffs associated to all possible action combinations performed by agents outside the interaction set. Preliminary computer simulations have shown that: (i) allowing for endogenous networks implies higher average payoff as compared to “frozen” networks; (ii) the statistics employed to evaluate payoffs strongly affect the efficiency of the system, i.e. convergence to a unique (multiple) steady-state(s) or not; (iii) for some class of statistics (e.g. MIN or MAX), the likelihood of efficient population learning strongly depends on whether agents are change-averse or not in discriminating between options delivering the same expected payoff.

Given the preliminary nature of this study, many issues remain to be explored. First, one needs a more careful investigation of the robustness of our results with respect to larger population sizes ( $N$ ). In fact, increasing  $N$  might have non-trivial effects on long-run system behavior because - for any given evaluation rule - the trade-off between average and spread (or variance) of payoff distribution may be heavily affected.

Second, alternative strategy/network updating processes may be introduced. For example, the introduction of idiosyncratic (low-probability) flips in strategy updating (or more generally of stochastic best-reply rules such as log-linear rules) may allow agents to better

explore the environment and avoid lock-ins. Moreover, network updating schemes which do not require mutual consent (i.e. unilateral link addition-deletion) might be employed. Along, the same lines, one may think to introduce some concepts of ‘power’ or ‘hierarchy’ in order to study systems where some agents can place a *veto* on payoff-decreasing changes within their network.

Third, one might study the effects of introducing (across-agent) heterogeneous, time-varying and/or endogenously changing evaluation rules. For instance: What 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?

Finally, the model seems well suited to address the issue of social efficiency of different network structures in complex (random) environments populated by adaptive agents. In a perspective quite similar to Jackson (2003), alternative efficiency criteria may be studied against different strategy/network updating schemes. Likewise, a social planner approach might be conceived in order to address optimal network structure analyses.

## References

- Ashlock, D., Smucker, M., Stanley, E. & Tesfatsion, L. (1996), ‘Preferential partner selection in an evolutionary study of prisoner’s dilemma’, *BioSystems* **37**, 99–125.
- Axelrod, R. (1984), *The Evolution of Cooperation*, New York, Basic Books.
- Bala, V. & Goyal, S. (2000), ‘A non-cooperative model of network formation’, *Econometrica* **68**(5), 1181–1229.
- Bednar, J. & Page, S. (2002), Game theory and culture. mimeo.
- Bhaskar, V. & Vega-Redondo, F. (1996), Migration and the evolution of conventions, Discussion Paper WP 354.96, Universitat Autònoma de Barcelona, Departament d’Economia i d’Historia.
- Blume, L. (1993), ‘The statistical mechanics of strategic interaction’, *Games and Economic Behaviour* **5**, 387–424.
- Blume, L. & Durlauf, S. (2001), The interaction-based approach to socioeconomic behavior, *in* S. Durlauf & H. Young, eds, ‘Social Dynamics’, The MIT Press, Cambridge, Mass.

- Brock, W. & Durlauf, S. (2001), ‘Discrete choice with social interactions’, *Review of Economic Studies* **68**, 235–260.
- Calvert, R. & Johnson, J. (1997), Interpretation and coordination in constitutional politics. mimeo.
- Dieckemann, T. (1999), ‘The evolution of conventions with mobile players’, *Journal of Economic Behavior and Organization* **38**, 93–111.
- Droste, E., Gilles, R. & 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**(5), 1047–1071.
- Ely, J. (1996), Local conventions. Unpublished Manuscript, Department of Economics, Northwestern University.
- Fagiolo, G. (1998), Spatial interactions in dynamic decentralized economies: a review, *in* P. Cohendet, P. Llerena, H. Stahn & G. Umbhauer, eds, ‘The Economics of Networks. Interaction and Behaviours’, Berlin - Heidelberg, Springer Verlag.
- Fagiolo, G. (2001), Coordination, local interactions and endogenous neighborhood formation, LEM Working Paper 2001/15, Sant’Anna School of Advanced Studies, Pisa, Italy.
- Goyal, S. & Janssen, M. (1997), ‘Non-exclusive conventions and social coordination’, *Journal of Economic Theory* **77**, 34–57.
- Goyal, S. & Vega-Redondo, F. (2001), Learning, network formation and coordination. Unpublished Manuscript, Erasmus University, Rotterdam.
- Hanaki, N. & Peterhansl, A. (2002), Viability of cooperation in evolving social interaction structures. mimeo.
- Hauk, E. (1996), ‘Leaving the prison: A discussion of the iterated prisoner’s dilemma with ostracism’, Ph.D. Thesis, European University Institute, Florence, Italy.
- Herz, A. (1994), ‘Collective phenomena in spatially extended evolutionary games’, *Journal Theoretical Biology* **169**, 65–87.
- Hirschleifer, D. & Rasmusen, E. (1989), ‘Cooperation in a repeated prisoners’ dilemma with ostracism’, *Journal of Economic Behavior and Organization* **12**, 87–106.
- Jackson, M. (2001), The stability and efficiency of directed communication networks. Unpublished Manuscript, Division of Humanities and Social Sciences, Caltech, Pasadena, CA.
- Jackson, M. (2003), The stability and efficiency of social and economic networks, *in* B. Dutta & M. Jackson, eds, ‘Networks and Groups: Models of Strategic Formation’, Springer-Verlag, Forthcoming.

- Jackson, M. & Watts, A. (2000), On the formation of interaction networks in social coordination games. Unpublished Manuscript, Division of Humanities and Social Sciences, Caltech, Pasadena, CA and Vanderbilt University, Nashville, TN.
- Jackson, M. & Watts, A. (2002), ‘The evolution of social and economic networks’, *Journal of Economic Theory* **106**, 265–295.
- Kandori, M., Mailath, G. & Rob, R. (1993), ‘Learning, mutation and long run equilibria in games’, *Econometrica* **61**(1), 29–56.
- 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.
- Kitcher, P. (1993), ‘The evolution of human altruism’, *Journal of Philosophy* **90**, 497–516.
- Mailath, G., Samuelson, L. & Shaked, A. (2000), Evolution and endogenous interactions, in U. Pagano & A. Nicita, eds, ‘The Evolution of Economic Diversity’, Siena, Routledge.
- Nowak, M., Bonhoefer, S. & May, R. (1994), ‘More spatial games’, *International Journal of Bifurcation and Chaos* **4**(1), 33–56.
- Nowak, M. & May, R. (1993), ‘The spatial dilemmas of evolution’, *International Journal of Bifurcation and Chaos* **3**(1), 35–78.
- Oechssler, J. (1997), ‘Decentralization and the coordination problem’, *Journal of Economic Behavior and Organization* **32**, 119–135.
- Oliphant, M. (1994), Evolving cooperation in the non-iterated prisoner’s dilemma: The importance of spatial organization, in R. Brooks & P. Maes, eds, ‘Artificial Life IV’, Cambridge, MIT Press.
- Oltra, V. & Schenk, E. (1998), Evolution of cooperation with local interactions and imitation, in P. Cohendet, P. Llerena, H. Stahn & G. Umbhauer, eds, ‘The Economics of Networks. Interaction and Behaviours’, Berlin - Heidelberg, Springer Verlag.
- Page, S. (1997), ‘On incentives and updating in agent based models’, *Computational Economics* **10**, 67–87.
- Skyrms, B. & Pemantle, R. (2000), ‘A dynamic model of social network formation’, *Proceedings of the National Academy of Sciences* **97**, 9340–9346.
- Smucker, M., Stanley, E. & Ashlock, D. (1994), Analyzing social network structures in the iterated prisoner’s dilemma with choice and refusal, Technical Report CS-TR-95-1259, University of Wisconsin, Madison, Department of Computer Sciences.
- Stanley, E., Ashlock, D. & Tesfatsion, L. (1994), Iterated prisoner’s dilemma with choice and refusal of partners, in C. Langton, ed., ‘Artificial Life III’, Santa Fe Institute, Santa Fe and Reading, MA, Addison-Wesley.



- Taylor, M. (1987), *The Possibility of Cooperation*, Cambridge, MA, Cambridge University Press.
- Tieman, A., Houba, H. & Van Der Laan, G. (1998), Cooperation in a multi-dimensional local interaction model. Unpublished Manuscript.
- Watts, D. (2001), ‘A dynamic model of network formation’, *Games and Economic Behavior* **34**, 331–341.
- Young, H. (1996), ‘The economics of convention’, *Journal of Economic Perspectives* **10**(2), 105–122.
- Young, H. (1998), *Individual Strategy and Social Structure*, Princeton, Princeton University Press.
- Zimmermann, M., Eguíluz, V. & San Miguel, M. (2001), Cooperation, adaptation and the emergence of leadership, *in* A. Kirman & J.-B. Zimmermann, eds, ‘Economics with Interacting Agents’, Lecture Notes in Economics and Mathematical Series, 503, Berlin, Springer.

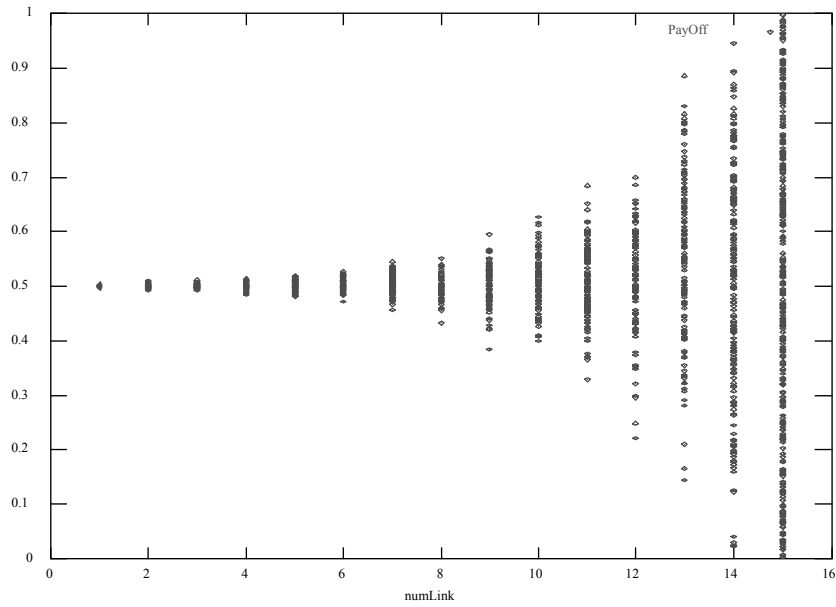


Figure 1: Fixed Networks with players employing the MEAN evaluation rule. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the long-run (across 50 populations of 15 agents).

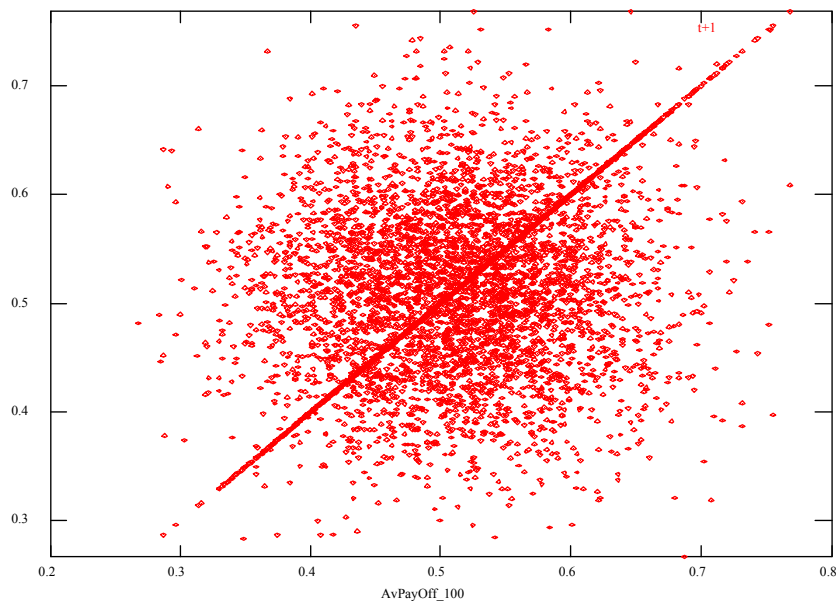


Figure 2: Fixed Networks with players employing MEAN evaluation rules. Scatter-plot of population average payoffs at time  $t+1$  (y-axis) vs. population average payoffs at time  $t$  (x-axis), after 10000 time steps and across 50 populations of 15 agents.

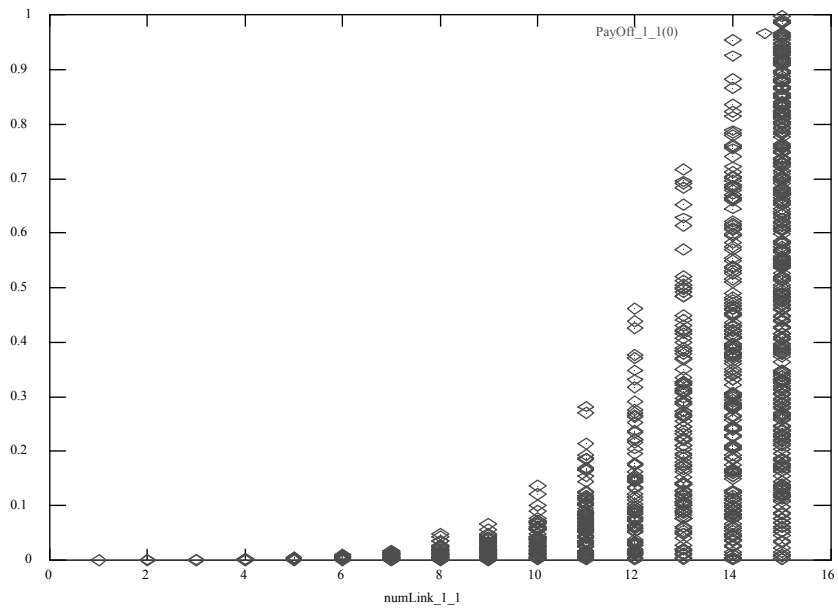


Figure 3: Fixed Networks with players employing MIN evaluation rules. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the long-run (across 50 populations of 15 agents).

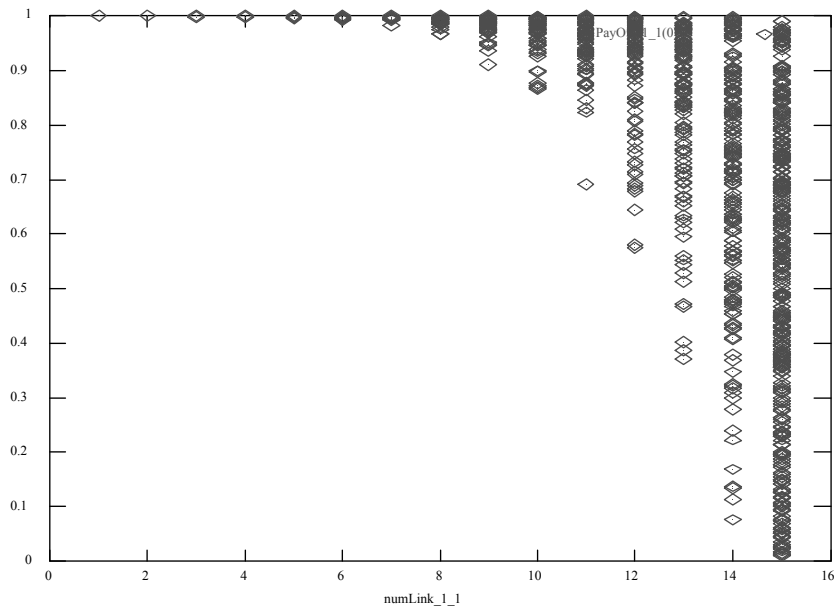


Figure 4: Fixed Networks with players employing MAX evaluation rules. Scatter-plot of individual payoffs (y-axis) vs. individual number of links (x-axis) in the long-run (across 50 populations of 15 agents).

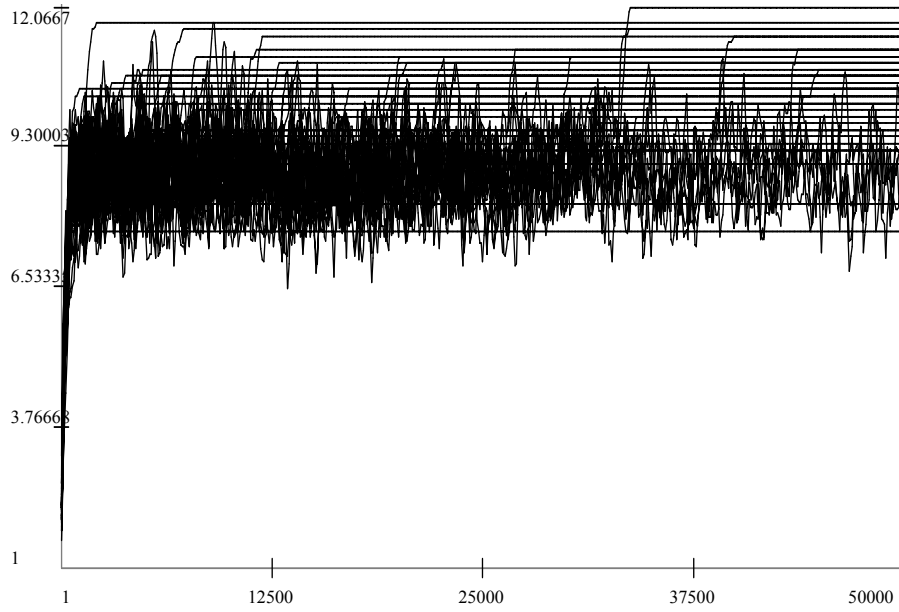


Figure 5: Endogenous Networks with players employing MEAN evaluation rules. Evolution over time of population average number of links across 50 populations with 15 agents each.

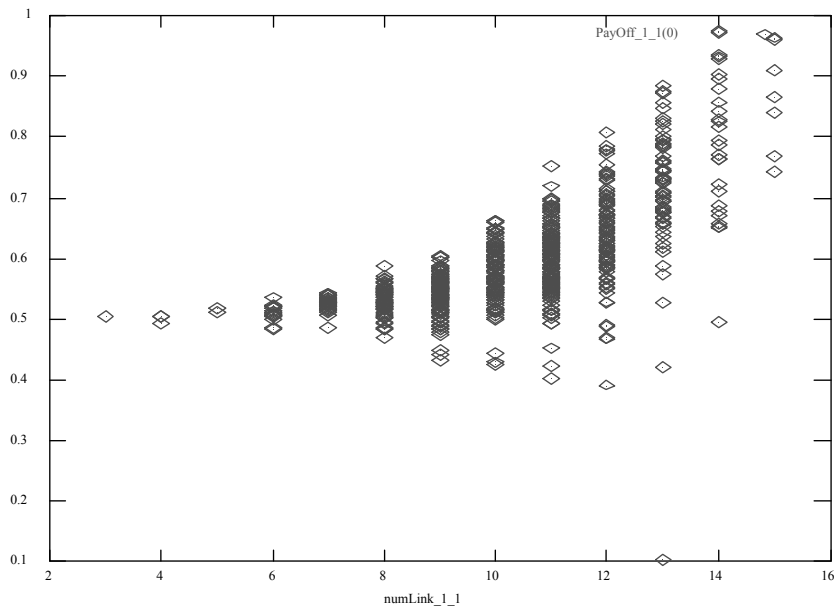


Figure 6: Endogenous Networks with players employing MEAN evaluation rules. Scatter-plot of individual payoffs ( $y$ -axis) vs. individual number of links ( $x$ -axis) in the long-run (across 50 populations of 15 agents).

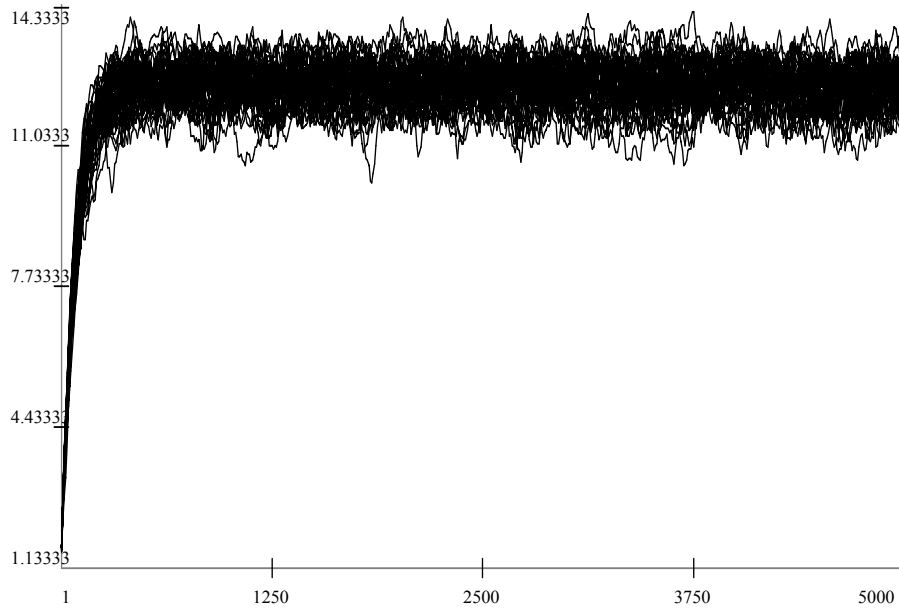


Figure 7: Endogenous Networks with players employing MIN evaluation rules and accepting payoff-neutral network changes. Evolution over time of population average number of links across 50 populations with 15 agents each.

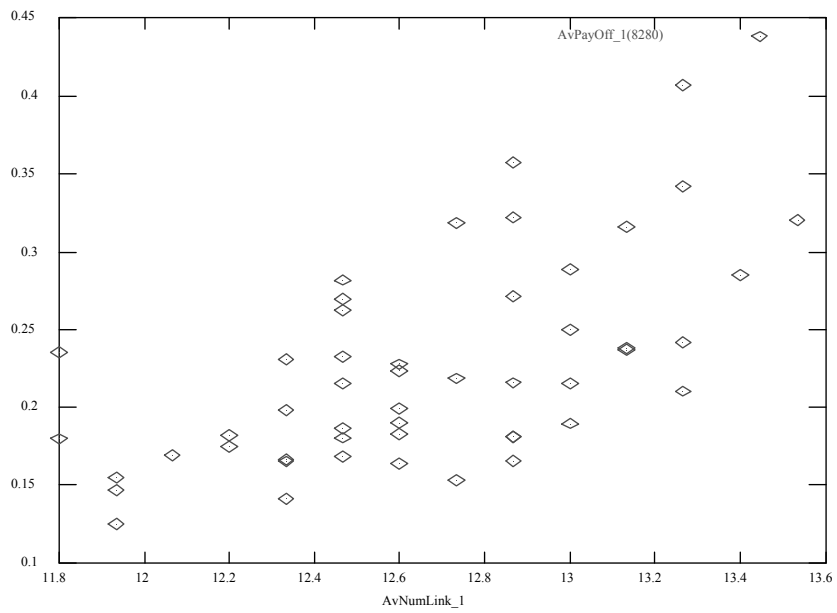


Figure 8: Endogenous Networks with players employing MIN evaluation rules and accepting payoff-neutral network changes. Scatter plot of average payoff (y-axis) as a function of the average number of links (x-axis) in the long-run (across 50 populations of 15 agents).

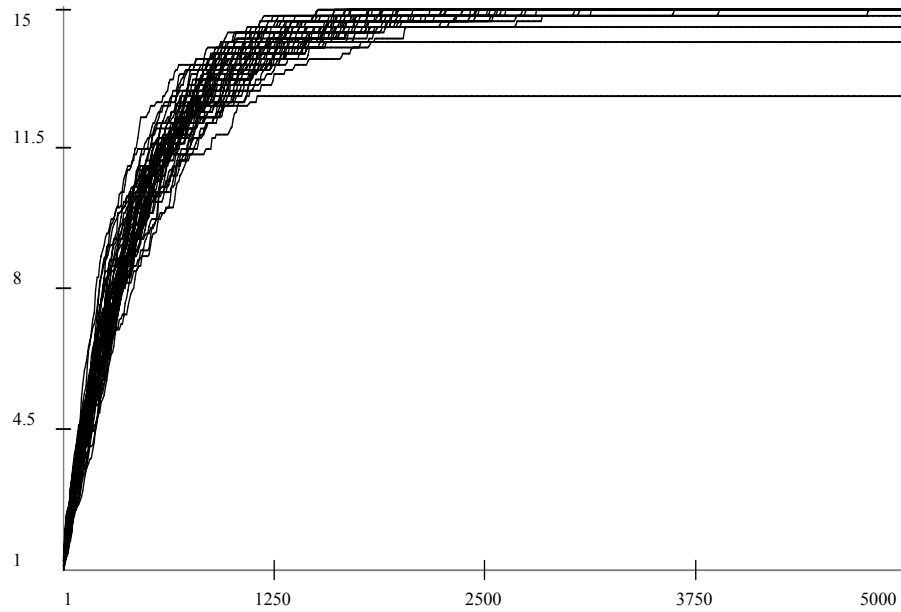


Figure 9: Endogenous Networks with players employing MIN evaluation rules but NOT accepting payoff-neutral network changes. Evolution over time of population average number of links across 50 populations with 15 agents each.

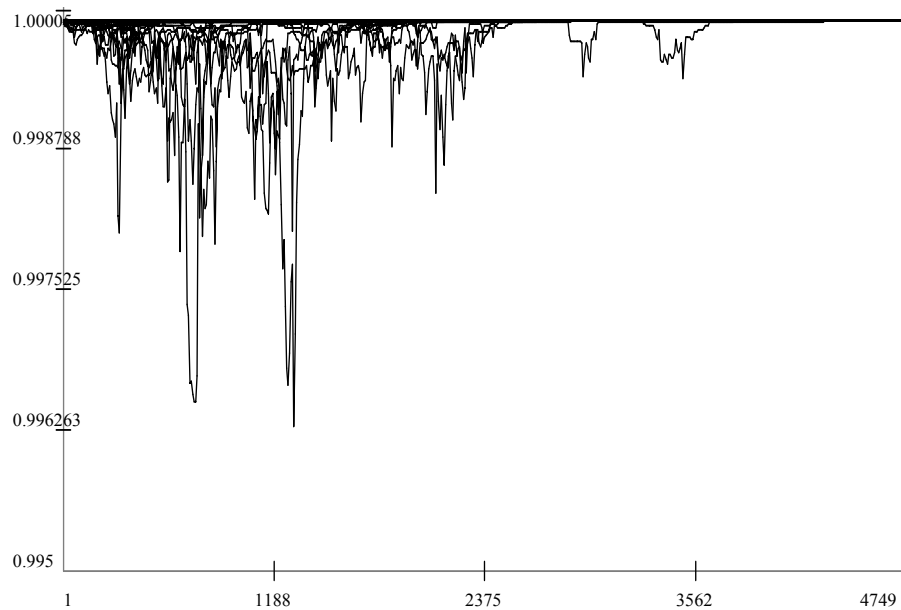


Figure 10: Endogenous Networks with players employing MAX evaluation rules and accepting payoff-neutral network changes. Evolution over time of population average payoffs across 50 populations with 15 agents each.

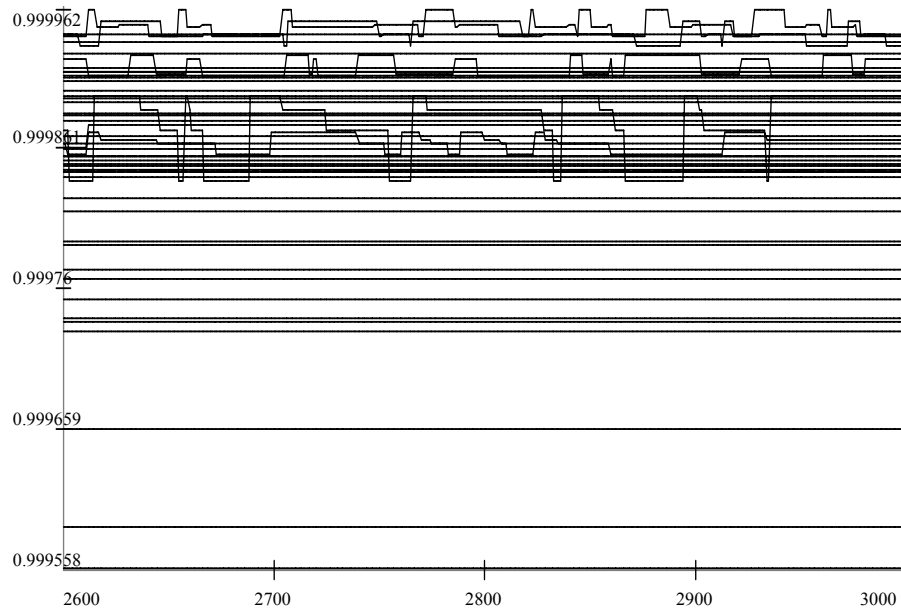


Figure 11: Endogenous Networks with players employing MAX evaluation rules but NOT accepting payoff-neutral network changes. Evolution over time of population average payoffs across 50 populations with 15 agents each (zoom on few time steps) .