Networks emerging in a volatile world

George Ehrhardt
The Abdus Salam International Centre for Theoretical Physics

Matteo Marsili
The Abdus Salam International Centre for Theoretical Physics

Fernando Vega-Redondo
Universidad de Alicante,
Instituto Valenciano de Investigaciones Económicas,
and University of Essex

August 31, 2006

Abstract

The paper proposes a model to study the conditions under which complex networks emerge (or not) when agents are involved in a dynamic coordination setup. In contrast with existing literature, however, our main focus is not on the entailed issue of equilibrium selection. Instead, our aim is to shed light on how agents' efforts to coordinate affect the process of network formation in a large and complex environment. The model posits that, over time, new links are created if they are profitable, and existing links disappear due to exogenous decay. Alongside this struggle between link creation and link destruction, agents' choices in the coordination game adapt to their current local conditions and thus coevolve with the social network. We characterize analytically the long-run behavior of the system and show that, as a function of the underlying parameters, the process displays sharp (discontinuous) transitions, resilient network transformations, and equilibrium multiplicity. As it turns out, these are features observed in a wide number of network phenomena in social environments.

JEL Classification nos.: C73, D83, D85.

1 Introduction

Social networks constitute the backbone underlying much of the interaction conducted in socio-economic environments.[1] Therefore, when this interaction

[1] Some important examples include labor markets (Granovetter (1974), Montgomery (1991), or Calvó-Armengol and Jackson (2004)), informal insurance (Murgai *et al.* (2002) or Fafchamps and Lund (2003)), trade arrangements (Kirman *et al.* (2000) or Kranton and

attains a *global* reach it must have, as its counterpart, the emergence of a social network with a wide range of overall (possibly indirect) connectivity. Naturally, for such a social network to emerge, agents must be in a position to link profitably. But this in turn demands that they display a similar – at least compatible – behavior. Thus, for example, they must use coherent communication procedures, share key social conventions, or have similar technical ability. Here, we may quote the influential work of Castells (1996):

> Networks are open structures, able to expand without limits, integrating new nodes as long as they are able to communicate within the network, namely as long as they share the same communication codes (for example, values or performance goals).

Reciprocally, of course, such convergence of behavior is facilitated by the range of social interaction being global rather than local or fragmented. This suggests the idea that the buildup of a global social network might be understood as the outcome of twin cross-reinforcing processes: one that facilitates the convergence of norms and behavior, and another that extends the range of social connectivity. This mutual reinforcement also suggests that if such a global transition indeed takes place, it should be relatively *fast* and *resilient*.

To fix ideas, consider the case of what might be labelled "knowledge networks." Prominent instances are the networks of scientific coauthorship among academics,[2] R&D collaboration agreements among firms,[3] or the more informal (but not less important) case of industrial districts and cities.[4] In all these cases, available empirical evidence suggests that the buildup of dense networks often occurs very fast, introducing a qualitative change in an earlier, much slower trend. By way of illustration, Goyal *et al.* (2003) reports a steep (three-fold) increase in the per capita number of collaborations among academic economists in the last three decades, or Hagerdoorn (2002) reports an even sharper (ten-fold) increase for R&D partnerships among firms during the decade 1980-1990. A similar sense of a rapidly unfolding process of network formation is gained from, say, the excellent account found in Castilla *et al.* (2000) of the rise to prominence of one of the most paradigmatic industrial districts of modern times: the semi-conductor firms located in Silicon Valley. They explain that, in roughly

Minehart (2001)), the performance of firms and other organizations (Krackhardt and Hanson (1993), van Alstyne (1997), and Burt (1992)). On the other hand, for the important class of so-called knowledge networks, see the discussion below.

[2]For example, Newman (2001) has studied the fields of physics, biomedical research, and computer science; Grossman (2002) has focused on mathematics; and Goyal *et al* (2003) has centered on academic economists

[3]Interestingly, R&D collaboration appears to be most prevalent among the most technologically dynamic sectors of modern economies such as those of electronics, pharmaceuticals, information technology, and aerospace and defense (cf. Delapierre and Mytelka (1998), Hagedoorn *et al* (2000), Orsenigo *et al* (2001)).

[4]See e.g. Saxenian (1994) for a comparative study of different industrial districts arising in the computer industry during the last decades of the twentieth century. On the other hand, concerning the unique role of cities as dense networks of economic activity, a lively account can be found in Jacobs (1984) while Glaeser *et al.* (1992) provide one of the earliest systematic empirical analyses of the phenomenon.

five years after the creation of Intel in 1968, much of the density of interpersonal connections that underlay the lively flow of ideas across the Valley was largely in place.[5]

Another important feature that transpires from these examples is that of *resilience*. In most of the knowledge networks mentioned above, the knowledge base underlying the social network tends to change very significantly over time. However, the overall structure of the network (e.g. its average density but *not* its specific links) often remains largely in place throughout these rapid changes. Staying with the case of the computer industry, a good illustration is the lucid account provided by Saxenian (1994) of the contrasting performance of two of its most emblematic industrial districts: the aforementioned Silicon Valley and the one that spread along Boston's route 128. The computer technology went through major innovations during the 1980's, which of course induced substantial changes on the industry and its underlying social networks. But, despite these network changes – in fact, *because* of them – Silicon Valley remained active and enterprising through these difficult times. This experience compares sharply with that of computer firms in the Boston area, which could tackle the crisis much less effectively because of their much less flexible structure of organization.

Combining the different points highlighted by the aforementioned discussion, three important features arise as natural implications to demand from our model of network dynamics: *fast transitions*, *resilience*, and *multiplicity* of long-run configurations. As we shall see, these are indeed characteristics delivered by our theoretical framework from the interplay (or coevolution) of two components:

(i) a fast dynamics, which tends to make nodes/agents that are close in the prevailing network (i.e. have a relatively short path between them) become similar in their behavior;

(ii) a slow dynamics that governs network adjustment and determines, based on nodes' affinity, the rate at which existing links are destroyed and new ones created.

More precisely, the fast dynamics posited in (i) is simply a *best-response* dynamics whereby agents tend to adjust their action in a myopically optimal fashion in a pure ($q \times q$) coordination game played locally with their current neighbors. On the other hand, the slower dynamics in (ii) embodies two opposing forces: volatility and link creation. Following Marsili *et al.* (2004), volatility is modelled as an exogenous force that destroys existing links at a given rate. Then, operating against the link destruction imposed by volatility, new links are taken to be created by every pair of unconnected agents at a rate that depends on whether or not their behavior is "coordinated" (and thus can profit from the link).

In mathematical terms, our main conclusion is that the asymptotic (long-run) behavior of the system exhibits a discontinuous transition in network connectivity as the rate of search – relative to volatility – exceeds a certain

[5] The network constructed by Castilla *et al.* (2000) has nodes be individual persons involved in the creation of Silicon Valley firms and a link between two nodes exists if the corresponding individuals were jointly involved in the establishment of at least one firm.

threshold. Moreover, as common in somewhat reminiscent models of statistical physics, this transition exhibits hysteresis (i.e. resilience). That is, once the transition has taken place, the high-connectivity phase persists as an equilibrium even when the search rate falls well below the aforementioned threshold.

The interplay between coordination games and networks has received ample attention in recent years. Initially, it was studied by Blume (1993), Ellison (1993), or Young (1998) in a context where players are involved in a fixed coordination game with their neighbors and the underlying network remains fixed. Subsequently, Ely (2003), Jackson and Watts (2002), and Goyal and Vega-Redondo (2005) have studied the issue in a setup where the network co-evolves with agents' behavior. In every case the main concern has been to understand how the pattern of local interaction (evolving or not) impinges on equilibrium selection – typically, on the tension between risk dominance and efficiency.

In this paper, since our primary concern is not one of equilibrium selection *per se*, the coordination game played by agents is assumed to be symmetric across different actions. Our focus, therefore, is not on the "convention" that might eventually be selected by society, but rather on how agents' struggle towards coordination may facilitate (or hinder) the build-up of the social network. As explained, we find that whether such a network buildup might indeed be achieved in the long run is not certain (i.e. depends on the parameters and, possibly, on the initial conditions). And, if it is achieved, the resulting network structure is always complex, in contrast with the simple "unrealistic" architectures (e.g. involving completely connected components) that typically arise in the aforementioned literature.

The modern theory of social networks has also studied models in which the formation of the network is the only strategic dimension in agents' behavior and therefore the network itself is the focus of the analysis. Seminal papers in this line research are Jackson and Wolinsky (1996) and Bala and Goyal (2000) – the reader can find in Jackson (2005) an exhaustive survey. In many of these models, suitable payoff structures (see e.g. different variations of the so-called "connections model") can generate nontrivial network architectures such as cycles, stars, linked stars, or flower networks. However insightful this is for the study of small or/and quite stable scenarios, its *equilibrium* approach seems hardly suitable for the study of large, complex, and volatile socio-economic environments. For these contexts (which arguably arise in many interesting applications), the traditional game-theoretic approach to studying network formation needs to be modified to obtain a better grasp of the problem. In particular, it is important to account for agents' incentives in ways that do not presume an excessive degree of rationality, information, or stationarity of the environment.

To do so is indeed one of the objectives of the present paper. Methodologically, we shall rely on some of the tools and insights developed in recent years by the recent literature on complex networks. The focus of this literature has been on understanding the large-scale (statistical) regularities induced by different network-formation processes. Salient representatives are the so-called *small-world networks* introduced by Watts and Strogatz (1999) and the *scale-*

free networks studied by Barabási and Albert (1999) – see Newman (2003) or Vega-Redondo (2006) for complementary surveys. This literature, however, typically relies on essentially *algorithmic* mechanisms of network construction that abstract from individual incentives. It misses, therefore, what are undoubtedly important considerations underlying the way in which networks are shaped in social contexts.[6] To bring them to bear crucially in the analysis is thus one additional motivation of the present research.

The rest of the paper is organized as follows. Section 2 presents the basic theoretical framework and describes in detail both the slow process of link formation and the fast process of action adjustment. This model is first studied in Section 3, where we characterize analytically the long-run behavior of the induced process and show that it produces the three key features mentioned above. Formally, the analysis is conducted for an infinite-population context, but the analysis is confirmed to match accurately numerical simulations conducted for large finite populations. Section 4 generalizes the basic scenario by allowing that agents' choices be subject to arbitrary noise (not necessarily small). We find that the gist of the conclusions derived from the basic model is maintained. The main body of the paper concludes in Section 5 with a summary of the model and a review of its main conclusions. The paper also includes two Appendices. In Appendix A, we provide the proofs of the results obtained for the basic model. In Appendix B, we explain the technical details underlying our analysis of the generalized setup.

2 Basic theoretical framework

Let there be a certain population of agents, $M = \{1, 2, ..., N\}$, who interact bilaterally over time as specified by the evolving social network. Time is modelled continuously, with $t \in [0, \infty)$. At any t, the state of the system $\omega(t)$ consists of two items: (i) the social network $g(t)$ that specifies the set of undirected links ij $(= ji)$ prevailing at t; (ii) the action profile $\alpha(t) \in A^M$, where $A = \{a_1, a_2, ..., a_q\}$ is the set of possible actions.

Players adjust both actions and links over time. The dynamics is described by a continuous Markov process for the state $\omega(t)$, and is therefore completely determined by the rates governing all possible transitions $\omega \to \omega'$. These transitions pertain to adjustments that involve (a) link creation, (b) link destruction, (c) action revision. We now describe each of these in turn.

Link creation: We posit that at a certain positive rate η each agent i receives a link creation opportunity. When such an opportunity arrives at some t, another agent j is randomly chosen in the population (all with the same probability). When no link exists between i and j (i.e. $ij \notin g(t)$), the link ij is formed if, and only if, $\alpha_i(t) = \alpha_j(t)$. This can be interpreted as a reflection of the fact that the establishment of a link is subject to a small

[6] Recently, however, Jackson and Rogers (2005) have studied a dynamic model of network formation where agents' incentives can play a significant role.

cost and thus will only be formed if the agents are "coordinated." On the other hand, if the link ij is already part of the network (i.e. $ij \in g(t)$), the link remains in place.

Link destruction: It is assumed that existing links decay at a rate λ, which for simplicity is taken to be constant and exogenous.[7] This component of the process may be provided with different (non-exclusive) interpretations. For example, it may be conceived as a reflection of unmodelled environmental *volatility* that affects the value or feasibility – and thus the persistence – of some of the existing links. Relatedly, it may also be understood as the result of capacity constraints in the ability of any particular agent to sustain many links. Note, in particular, that this formulation implies that the expected number of links that an agent can simultaneously sustain is finite. This simply follows from the fact that the rate of link creation is uniformly bounded while the removal of each existing link proceeds independently.

Action revision: Finally, concerning action adjustment, we posit that, at every t, each agent i independently receives at a rate ν the opportunity to revise her current action. If this revision opportunity materializes, we assume that the agent chooses any of the actions that is in the majority among her neighbors. More precisely, a particular action $a_r \in A$ is chosen with positive probability if, and only if,

$$|\{\alpha_j(t) = a_r : ij \in g(t)\}| \geq |\{\alpha_j(t) = a_{r'} : ij \in g(t)\}| \qquad r' = 1, 2, ..., q \tag{1}$$

where $|\{\cdot\}|$ stands for the cardinality of the set in question. For concreteness, we assume that every action a_r that satisfies (1) is chosen with the same probability.

To fix ideas, the above adjustment rule may be interpreted as reflecting a (myopic) best-response dynamics in a context where every pair of connected agents incur a linking cost $c > 0$ and interact according to a pure coordination game. In this game, the action space A defines the common strategy space and the payoff function $\pi : A \times A \to \mathbb{R}$ has the following form:[8]

$$\pi(a_r, a_{r'}) = \bar{\pi} > c \qquad \text{if } r = r' \tag{2}$$

$$\pi(a_r, a_{r'}) = 0 \qquad \text{if } r \neq r'. \tag{3}$$

Each $\pi(a_r, a_{r'})$ is interpreted as the gross payoff (*not* including the linking cost) of a player who plays a_r when faced with a player who chooses $a_{r'}$. Note that

[7] Alternatively, one could generalize matters and (in line with the considerations motivating link creation) posit that links connecting nodes i, j such that $\alpha_i(t) \neq \alpha_j(t)$ decay at a possibly larger rate $\lambda' \geq \lambda$. As it should be clear from Remark 1 below, the long run behavior for this general formulation would always coincide with the particular case where $\lambda' = \lambda$.

[8] In general, one could consider any bilateral coordination game. But then, in view of the link formation rule, we would have to make the implicit companion assumption that the cost of establishing a new link exceeds any payoff agents can earn when miscoordinated while, on the other hand, all (relevant) actions attain a *coordination* payoff higher than that cost.

the payoff resulting from any particular action grows linearly with the number of neighbors who choose it. Thus, if one makes the natural assumption that the *same* action in the set A must be chosen in *every* bilateral encounter, the number (or fraction) of neighbors who display each of them determines their relative performance. This is the observation that motivates the adjustment rule (1).

One of the three rates, η, λ or ν, can be normalized to one by a suitable choice of time unit. Hence we set $\lambda = 1$ hereafter, without loss of generality. The parameter ν, on the other hand, does not play a crucial role in the model, its relative magnitude having no important effect on the gist of our conclusions. Thus, in the interest of simplicity, we posit that $\nu \gg 1$, i.e. action adjustment is much faster than network change.[9] Our discussion, therefore, will focus on the effect of η (the rate of link creation), which will stand out as the key parameter of the analysis.

3 Long-run analysis

For any *finite* population, the Markov process $\{\omega(t)\}$ is ergodic and therefore displays a unique long-run (invariant) distribution. This is a direct consequence of the property that, for any two configurations ω and ω', it is possible to find a finite sequence of transitions (occurring with a positive joint probability) that leads from one to the other. Our interest, however, is on very large populations. Formally, we consider the limit $N \to \infty$, where analytical derivations are made simpler. But this limit case need no longer be ergodic – an instance of the so-called ergodicity breakdown for large systems. Indeed, such a phenomenon does arise in our (infinite) system, which will be seen to display multiple long-run outcomes within a certain parameter range. Conceptually, this must be interpreted as a situation where the expected waiting times across different basins of attraction are so long that, in effect, initial conditions do matter for long-run prediction when the population is large.

In order to characterize the long-run behavior of the system, it is useful to start from the following observation:

Remark 1 *The number of links $ij \in g(t)$ with $\alpha_i(t) \neq \alpha_j(t)$ vanishes almost surely as $t \to \infty$.*

The above statement simply derives from the fact that links between nodes displaying different actions are never created and every existing link decays at a constant rate. Furthermore, changes of action can never increase the total number of links that connect nodes with differing actions. So, at the beginning

[9] In effect, the assumption that $\nu \gg 1$ introduces two time scales and focuses on the scenario where the impact of network structure on agents' behavior is sharpest. Indeed, in the polar extreme situation where the adjustment of actions is much slower than that of the network ($1/\nu \gg 1$), matters are much less interesting – in particular, the network changes so fast that there is no clear sense in which the current specific architecture of the network has any effect on agents' behavior.

of the process, it is perfectly possible that nodes choosing some action a_r might be connected to other nodes displaying a different action. But if we focus on the long run, we can safely assume that only links between agents choosing the same actions will be present. This makes it convenient to partition nodes into $q+1$ subsets (or classes) $G_0(t),\ G_1(t), ...,\ G_q(t)$ as follows:

$$\begin{aligned} G_0(t) &= \{i \in M : z_i(t) = 0\} \\ G_r(t) &= \{i : z_i(t) > 0,\ \alpha_i(t) = a_r\} \qquad (r = 1, 2, ..., q) \end{aligned}$$

where $z_i(t) \equiv |\{j \in M : ij \in g(t)\}|$ is the degree of any node i. Such partition is useful because, as a consequence of Remark 1, only links between nodes in the same class G_r exist in the long run.

Concerning actions, on the other hand, the dynamic implications of the above partition are also straightforward. For each $r > 0$ and any node $i \in G_r(t)$ at some t, peer interaction will freeze this node to action a_r at all future time $t' > t$ as long as that node still keeps some link. In contrast, for isolated nodes ($i \in G_0(t)$) their actions are continuously updated at random at the rate ν. Thus, in view of the assumption that $\nu \gg 1$, whenever a node $i \in G_0(t)$ is involved in a link formation event, we can regard its current action as uniformly drawn at random from the set A.

As advanced, we want to focus on the limit case where $N \to \infty$. Then, the above considerations make it possible to derive the induced dynamics of the (relative) size distribution of each of the classes G_r as follows. Let

$$n_r(t) = \frac{|G_r(t)|}{N} \tag{4}$$

stand for the measure (or fraction) of nodes belonging to the class G_r at time t. Consider some "infinitesimal" time interval $[t, t + \mathrm{d}t)$ and denote by $\mathrm{d}N_r$ the corresponding change in the population of nodes $i \in G_r$. This change results from the sum of *independent* processes affecting each node (link creation) and each existing link (decay). This implies that $\mathrm{d}n_r = \mathrm{d}N_r/N$ satisfies the law of large numbers, i.e. it converges to $E[\mathrm{d}N_r]/N$.

Denote by $\mathrm{d}N_r^+$ the positive effect due to link creation and by $\mathrm{d}N_r^-$ the negative effect due to link destruction, i.e. $\mathrm{d}N_r = \mathrm{d}N_r^+ - \mathrm{d}N_r^-$. In view of Remark 1, their expected magnitudes can be respectively computed in the long run as follows.

i) ***Link creation:*** In the time interval of length $\mathrm{d}t$ new nodes can enter G_r from G_0 because of the formation of either a link between $i \in G_r$ and $j \in G_0$, or because of the creation of a link between $i, j \in G_0$ with $\alpha_i(t) = \alpha_j(t) = a_r$. Accounting for each of these two flows, we can write:

$$E[\mathrm{d}N_r^+] = \frac{2\eta \mathrm{d}t}{q}\frac{|G_0||G_r|}{N-1} + \frac{4\eta \mathrm{d}t}{q^2}\frac{|G_0|(|G_0|-1)}{2(N-1)}. \tag{5}$$

The right-hand side of the previous expression has two terms. The first one accounts for the rate at which new links are established between nodes in G_0 and the corresponding class G_r, $r > 0$. This can happen in two ways: either the node in G_0 receives the opportunity (at the rate η) and encounters a node in G_r (with probability $|G_r|/(N-1)$) or *vice versa*. Thus a factor of 2 must be included to reflect these two possibilities. In either case, the link is actually formed only if the node in G_0 displays action a_r – an event that has probability $1/q$. The second term in (5), on the other hand, reflects analogous considerations, but now concerning links that may be formed between two nodes in G_0. In this case, the rate must be proportional to the number of possible links among those nodes ($|G_0|(|G_0|-1)/2$) and the relevant probability (i.e. the probability that the two such nodes display action a_r) is $1/q^2$. Finally, note that the leading factor is twice as large as before because each link formed now contributes *two* new nodes to G_r.

ii) Link decay: Nodes can exit the class G_r and enter G_0 if they have $z_i = 1$ and their single link decays. Since this event occurs at the rate $\lambda = 1$ (thus with probability $\mathrm{d}t$), we have:

$$E[\mathrm{d}N_r^-] = -|G_r| \; P\{z_i(t) = 1 | i \in G_r(t)\} \; \mathrm{d}t \tag{6}$$

Dividing both (5) and (6) by N and using (4), the dynamics of $\mathbf{n}(t) \equiv (n_0(t), n_1(t), ..., n_q(t))$ may be described (when $\mathrm{d}t \to 0$ and $N \to \infty$) through the following set of differential equations:[10]

$$\dot{n}_r = \frac{n_0}{q} 2\eta n_r + \frac{n_0}{q} 2\eta \frac{n_0}{q} - n_r p_r(1), \qquad r = 1, \ldots, q \tag{7}$$

where $p_r(1) \equiv P\{z_i(t) = 1 | i \in G_r(t)\}$ is the probability that a randomly chosen node in $G_r(t)$ has $z_i(t) = 1$ (i.e. only one neighbor).

Our main concern is to characterize the locally stable equilibria of the dynamical system (7), i.e. the stationary size distributions $\mathbf{n} \equiv (n_0, n_1, ..., n_q)$ that are robust to small perturbations. Inspection of (7) shows that, in order to characterize those equilibria, it is necessary to obtain a corresponding estimation of the fraction of nodes in each class G_r that have one neighbor: these are precisely the nodes that either enter or abandon that class at any given point in time. In fact, one can go well beyond this and obtain a *full characterization* of the degree distribution prevailing at any equilibrium. This is the content of Proposition 2 below. The key intuition underlying this result is that, among the nodes of class G_r that continue belonging to it (i.e. remain connected), their degree evolves through pure random creation and deletion of links. Thus, in the long run, each of these classes is connected through a random network with a (conditional)[11] Poisson degree distribution.

---

[10] For notational simplicity, we dispense with the explicit specification of the time variable t.

[11] By definition, every node in G_r for $r > 0$ has at least one link. Thus, naturally, the corresponding Poisson distribution is conditioned on the event $\{i \in G_r : z_i \geq 1\}$.

Proposition 2 *Let* $\mathbf{n} \equiv (n_0, n_1, ..., n_q)$ *be the size distribution prevailing among the classes* G_r *in a stationary state when* $N \to \infty$. *For each* $r = 1, 2, ..., q$, *the nodes in* G_r *display a network structure given by a random undirected graph with degree distribution*

$$p_r(k) \equiv \frac{|\{i \in G_r,\ z_i = k\}|}{|G_r|} = \frac{c_r^k\ e^{-c_r}}{(1 - e^{-c_r})k!}, \qquad k \geq 1 \tag{8}$$

where

$$c_r \equiv 2\eta(\frac{n_0}{q} + n_r). \tag{9}$$

Proof. Appendix A. ■

The previous result allows us to regard each of the $q+1$ classes G_r as defining an independent subnetwork at equilibrium. First, there is the class G_0 of isolated nodes that trivially define an empty network. Then, for each $r = 1, 2, ..., q$, the corresponding class G_r defines a separate random network displaying a Poisson degree distribution (conditional on positive connectivity). Over time, the identity of the nodes belonging to each class changes. In equilibrium, however, their overall statistical structure remains unchanged.

Next, our objective is to characterize the scope of network structures that are possible at equilibrium. To this end, we first note that, using (9) to solve for every n_r ($r \geq 1$) and introducing them in (7), the stationarity condition characterizing equilibrium states $\mathbf{n} \equiv (n_0, n_1, ..., n_q)$ can be rewritten as follows:

$$\frac{c_r}{1 - e^{-c_r}} \left(\frac{n_0}{q} - \frac{c_r e^{-c_r}}{2\eta} \right) = 0 \qquad (r = 1, 2, ..., q). \tag{10}$$

This leads to the sharp characterization of equilibrium states given by the following result.

Proposition 3 *The stationary points satisfying (10) can be characterized by the measure* n_0 *of isolated nodes and the number* L_+ *of classes* G_r $(r \geq 1)$ *with a common* $c_r = c_+ > 1$ *while all other* $q - L_+$ *classes have a common* $c_r = c_- < 1$. *The values of* c_+ *and* c_- *satisfy*

$$\begin{aligned} c_+ e^{-c_+} &= \frac{2\eta}{q} n_0, \quad c_+ > 1 && (11)\\ c_- e^{-c_-} &= \frac{2\eta}{q} n_0, \quad c_- < 1 && (12)\\ L_+ c_+ + (q - L_+) c_- &= 2\eta. && (13) \end{aligned}$$

so that the degree distribution of a randomly chosen node is given by

$$p(k) = \frac{n_0}{q} \left(L_+ \frac{c_+^k}{k!} + (q - L_+) \frac{c_-^k}{k!} \right) \tag{14}$$

and the average degree is

$$z = \frac{n_0}{q} \left[L_+\ c_+\ e^{c_+} + (q - L_+)\ c_-\ e^{c_-} \right]. \tag{15}$$

Proof. Appendix A. ■

The equilibrium characterization provided by the former Proposition partitions the subnetworks associated to each G_r ($r > 0$) into two kinds: the L_+ subnetworks with a relatively high average degree $c_+/(1-e^{-c_+})$, and the remaining $(q - L_+)$ subnetworks with a lower average degree $c_-/(1-e^{-c_-})$. The former, almost surely displays a so-called giant component, i.e. a connected component spanning a finite fraction of all the nodes. Indeed the condition for the existence of a giant component is known to be:[12]

$$E[z_i^2] - 2E[z_i] = \frac{c_r(c_r-1)}{1-e^{-c_r}} > 0$$

which is satisfied if $c_r = c_+ > 1$. On the contrary, the subsets G_r with $c_r = c_- < 1$ are characterized by a sparse network of disjoint subgraphs of finite size.

The dichotomic equilibria characterized by Proposition 3 are uniquely described by two values: the number L_+ of high-connectivity classes and the fraction n_0 of isolated nodes. Given these two values, one can unambiguously determine the corresponding features of the equilibria (e.g. its degree distribution). Next, we study how the parameters of the model, η and q, shape the range of equilibrium possibilities for L_+ and n_0.

First, we focus on an equilibrium with $L_+ = 0$. Introducing this value of L_+ in (13) we find:

$$c_- = \frac{2\eta}{q}, \qquad (L_+ = 0) \tag{16}$$

which is compatible with the constraint $c_- < 1$ only for

$$\eta < \hat{\eta} \equiv \frac{q}{2}. \tag{17}$$

Thus only if η is low enough there exists an equilibrium with $L_+ = 0$. In this case, introducing (16) into (12), we find that the fraction of isolated nodes $n_0 = e^{-2\eta/q}$. Finally, in view of (14), let us note that the induced structure of the *whole* network is strictly that of a Poisson (or Erdös-Rényi) random graph, with

$$p(k) = \frac{(2\eta/q)^k}{k!} e^{-2\eta/q}, \qquad (L_+ = 0). \tag{18}$$

Next, consider the equilibrium with $L_+ = 1$. While the equilibrium with $L_+ = 0$ is unique, when it exists, because it is symmetric across the different classes G_r, this symmetry is broken in an equilibrium with $L_+ = 1$. Thus, associated to any such equilibrium, there are $q - 1$ alternative equilibria that differ only in the action a_r ($r = 1, 2, ..., q$) displayed in the single high-connectivity class. For simplicity, we shall refer to all these as a single equilibrium. To check

[12] The condition characterizing the existence of a giant component in (unconditional) Poisson random graphs was established, among others, by Erdös and Rényi (1959, 1960). For arbitrary random networks, this condition generalizes to $E[z_i^2] - 2E[z_i] > 0$, as determined by Molloy and Reed (1995) – see e.g. Bollobás (2001).

the conditions under which it exists, denote by $\Delta \equiv c_+ - c_-$ the difference between the two values of c_r characterizing the classes with high and low connectivity in the $L_+ = 1$ solution of (11)-(13). The relation $c_- e^{-c_-} = c_+ e^{-c_+}$ implies that $c_- = \Delta/(e^\Delta - 1)$ and (13) can be rewritten as $\Delta = 2\eta - qc_-$. Combining these considerations, we obtain

$$\frac{\Delta}{2}\left(1 + \frac{q}{e^\Delta - 1}\right) = \eta. \tag{19}$$

Let ψ stand for the function defined by the left hand-side of (19), i.e. $\psi(\Delta) \equiv \frac{\Delta}{2}\left(1 + \frac{q}{e^\Delta - 1}\right)$ for all $\Delta \geq 0$. This function is plotted in Figure 1 for two different values of q. It is a nonnegative and unbounded function that has a unique minimum and satisfies $\psi(0) = q/2 \equiv \hat{\eta}$ and $\psi'(0) = 1/2 - q/4$. Let $\check{\eta} \equiv \min_{\Delta \geq 0} \psi(\Delta)$. If $q > 2$, the function ψ has an interior minimum and therefore $\check{\eta} < \hat{\eta}$. Thus, it follows directly from (19) that an equilibria with $L_+ = 1$ exists if, and only if, $\eta \geq \check{\eta}$. In fact, this equilibrium is unique if $\eta > \hat{\eta}$. But if η lies in the interval $(\check{\eta}, \hat{\eta}]$ there are two values of Δ that satisfy (19) and hence two different equilibria with $L_+ = 1$. Finally, we note that in the border case where $q = 2$, the function $\psi(\Delta)$ is monotonically increasing for all $\Delta \geq 0$, as depicted by the dotted line in Figure 1. Therefore, we have that $\check{\eta} = \hat{\eta}$ and there is at most a single equilibrium with $L_+ = 1$ throughout.

For clarity, we summarize the most important points of the previous discussion in the following result.

Proposition 4 *An equilibrium exists for all values of η and q. Let $\hat{\eta} \equiv \frac{q}{2}$ and $\check{\eta} \equiv \min_{\Delta \geq 0} \psi(\Delta)$. Then, an equilibrium with $L_+ = 0$ exists if, and only if, $\eta \leq \hat{\eta}$. On the other hand, if $\eta > \check{\eta}$, there is an equilibrium with $L_+ = 1$. Thus, if $\check{\eta} < \hat{\eta}$ (which occurs iff $q > 2$), there is a range $(\check{\eta}, \hat{\eta})$ such that for all η belonging to it there are three equilibria, one with $L_+ = 0$ and two with $L_+ = 1$.*

The former proposition focuses on equilibria that either have no highly-connected class ($L_+ = 0$) or include just *one* of them ($L_+ = 1$). This leaves open the question of whether equilibria with $L_+ \geq 2$ might also exist under suitable circumstances. In fact they do, and one could conduct a similar characterization for such equilibria, addressing in turn the cases $L_+ = 2, 3, ..., q$. This, however, need not concern us here since, as established by Proposition 5 below, those equilibria are unstable and thus should not be conceived as robust predictions of the model. The following proposition also establishes, on the other hand, that the equilibrium with $L_+ = 0$ (when it exists) and *only one* of the two specified equilibria with $L_+ = 1$ (when they exist) are stable. As we shall see, the crucial importance of the stability requirement in weeding out equilibria in the long run is starkly confirmed by our numerical simulations.

Proposition 5 *The equilibrium with $L_+ = 0$ is stable for all $\eta < \hat{\eta} = q/2$. All equilibria with $L_+ > 1$ are unstable. Finally, only the equilibrium with $L_+ = 1$ that displays the largest Δ is stable.*

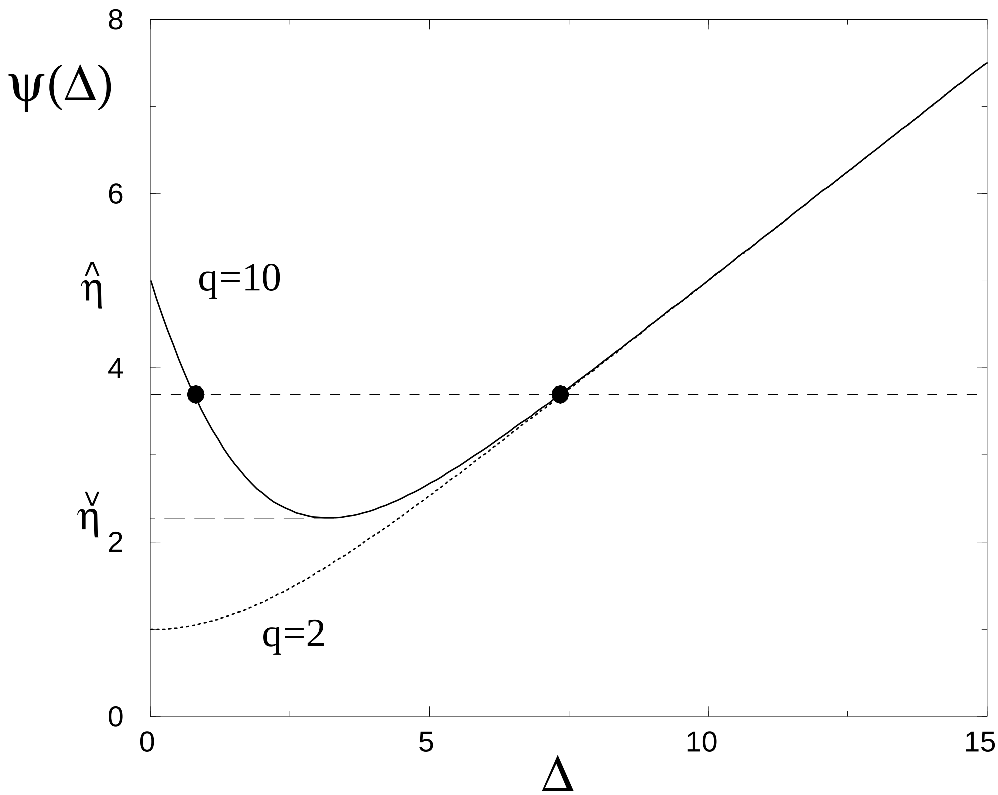

Figure 1: The solid line is the function $\psi(\Delta)$ for $q = 10$, while the dotted line is that for $q = 2$. The intersection with an horizontal line (light dashed) at height η yields the corresponding equilibrium points, marked as full dots.

Proof. Appendix A. ■

Intuitively, it is easy to understand why equilibria with more than one well-connected class are unstable. In order to render this situation stationary, a delicate balance/symmetry must be preserved between each of them. But, in the presence of fluctuations that make one of the c_+ classes get some more "mass" (i.e. nodes) than the others, the ensuing imbalance will be amplified exponentially by the dynamics, eventually leading to the shrinkage of all the high-connectivity classes except one. Also note that an ($L_+ = 1$) equilibrium with Δ which decreases with η would be rather unrealistic, as one expects that the network's density should increase with the networking effort η. It is then reassuring that such solutions are unstable.

Combining the two previous results, the predictions resulting from the model (for given $q > 2$) can be described as follows.

(1) If $\eta < \check{\eta}$, a single outcome can prevail in the long run in which all classes $r = 1, 2, ..., q$ display the low average degree $c_-/(1 - e^{-c_-})$ induced by the symmetric solution of (10).

(2) If $\check{\eta} < \eta < \hat{\eta}$, there is a long-run outcome as described in (1). In addition, there is another outcome where a single class r enjoys a high average degree $c_+/(1 - e^{-c_+})$, the $q - 1$ remaining classes $r' \neq r$ of connected nodes display a low average degree $c_-/(1 - e^{-c_-})$, and $\Delta = c_+ - c_-$ is given by the highest solution of (19).

(3) If $\eta > \hat{\eta}$, there is a single asymmetric outcome prevailing in the long run, which is as described in the latter part of (2).

The situation is graphically depicted in Figure 2, where we fix $q = 10$ and let η vary in the whole range $(0, \infty)$. The diagram traces the average degree induced by the stable solutions of the model. First, we find that if $\eta < \hat{\eta}(q) = 5$, there is an equilibrium with an average degree that is relatively low but increases slowly with η. This equilibrium corresponds to the case with $L_+ = 0$. On the other hand, if $\eta > \check{\eta}(q) \simeq 2.28$ there is an equilibrium with $L_+ = 1$ where the average degree is relatively high and increases quite steeply as η grows. Thus, in a middle range $(2.28, 5)$ the two kind of equilibria are possible − one with a low connectivity and large action heterogeneity, another with much higher degree and action homogeneity. It is worth stressing that, as shown in Figure 2 itself, this theoretical analysis provides an accurate prediction of the behavior observed in numerical simulations for large populations. Thus, in particular, within the range where the model displays equilibrium multiplicity, the two possible equilibria arise as long-run outcomes depending on initial conditions (i.e. on whether the starting configuration displays a high or low connectivity).

Finally, we argue that the model helps shed light on the qualitative empirical evidence discussed in the Introduction for a variety of social network phenomena. First, consider the occurrence of *sharp transitions* in response to slight changes

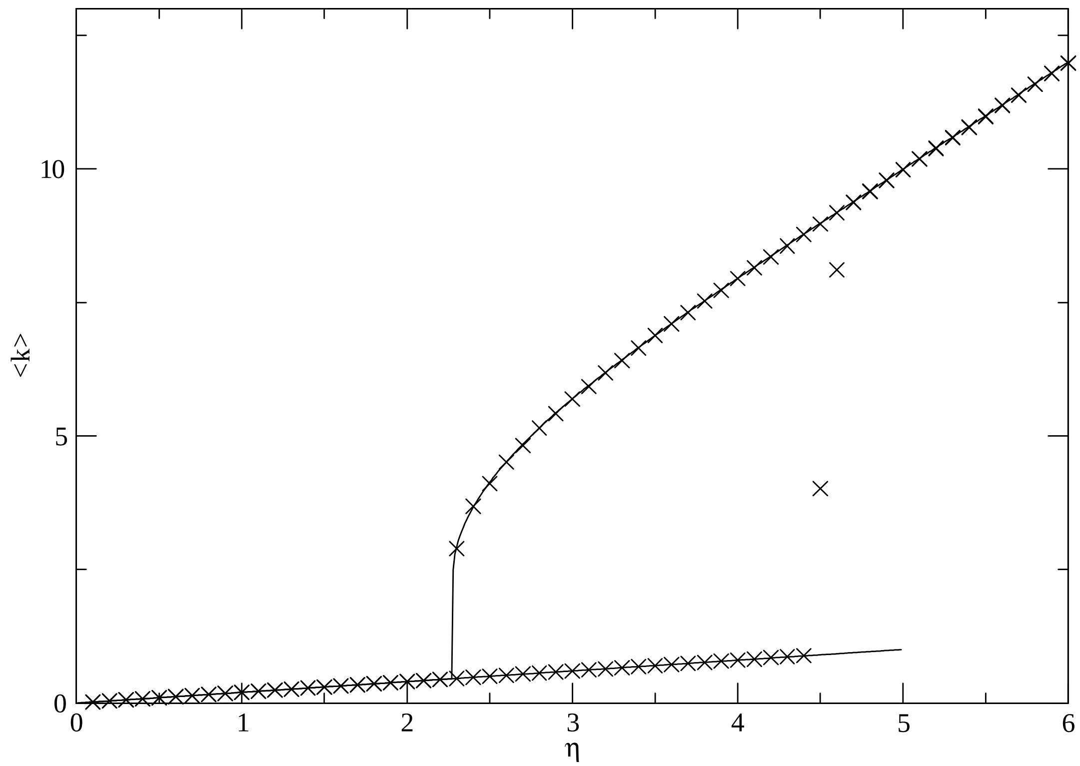

Figure 2: Plot of the mean connectivity induced by the model against the rate of link formation for $q = 10$, from equation 15. Crosses represent simulation points for a population size $n = 10000$. In the simulations, the low equilibrium becomes unstable below the predicted value of 5 because fluctuations have significant effects close to the transition point due to the finiteness of the population. The two crosses not on the theoretical curves are due to the system jumping from the low connectivity to the high connectivity state during the data taking time period.

in environmental conditions. To fix ideas, think of the parameter η as embodying an exogenous factor that affects the speed and effectiveness of link creation (e.g. the rate at which individuals "meet" or/and they find profitable opportunities for bilateral collaboration). Suppose that, starting from an originally stable configuration, η changes slightly from some η_0 to another $\eta_1 = \eta_0 + \mathrm{d}\eta$ for some small $\mathrm{d}\eta \neq 0$. Then, suppose that the system is given time to adjust from the "initial conditions" given by the previously stable configuration to another one that is stable under η_1. If η_0 is not close to one of the thresholds, $\check{\eta}(q)$ and $\hat{\eta}(q)$, the system will simply experience a small change in the situation, moving along the corresponding equilibrium branch on which it was originally placed. This would merely represent a slight adjustment in the phenomenon at hand. Now suppose that, instead, the initial η_0 lies, say, barely below $\hat{\eta}(q)$. Then, if $\eta_1 > \hat{\eta}(q)$, the state of the system can undergo a large "discontinuous" shift, moving from a state where connectivity and action homogeneity are relatively low to a situation where both are markedly higher. If this occurs relatively fast in terms of the time scale at which the changes in η unfold, one may interpret it as reflecting a sharp transition of the sort observed in the network examples discussed.

The model can also account for the indicated phenomenon of *resilience*, interpreted in the following sense. Suppose that, once the aforementioned transition to a dense network has materialized, the value of η changes back to the original $\eta_0 = \eta_1 - \mathrm{d}\eta$. Then, the system should not be predicted to return to the original situation with low connectivity. Rather, it must be expected to remain in the "upper equilibrium branch," thus exhibiting only a small downward adjustment in its average degree. After the first transition has taken place, in other words, a robust state of affairs is reached that tends to remain approximately in place even if the underlying environment is perturbed back towards the original situation.

Thirdly, note that the very phenomenon of resilience just explained entails *equilibrium multiplicity*, which was also mentioned as a salient feature of network phenomena in many social contexts. Indeed, it is such multiplicity that allows "history" to play a key role as a selection device in the model. In terms of our former heuristic discussion, this simply means that the outcome predicted at $\eta = \eta_0$ cannot abstract from whether the system arrives at this situation from an earlier state with high or low connectivity.

Overall, the above features of the model (i.e. sharp transitions, resilient transformations, and equilibrium multiplicity) delineate a subtle scenario of network coevolution where complex dynamics result from the interplay of three subprocesses: link creation, link destruction, and action adjustment. In the next section, we explore whether the same performance is obtained in contexts where, unlike what happens in our present benchmark model, the underlying rules of behavior are subject to (possibly sizable) noise.

4 A generalized model with behavioral noise

In principle, one could generalize the basic model studied above by introducing noise into both the link-creation mechanism and the rule for action adjustment. Since the former possibility does not add anything important to the analysis,[13] we choose to focus on the latter. So we now contemplate a model where, at the time of adjustment, agents revise their action in a way that embeds the previous formulation as follows:

Action revision, generalized: At every t, each agent i independently receives at a rate ν the opportunity to revise her current action. If this revision opportunity materializes, then she chooses every possible action $a_r \in A$ with probability

$$\Pr(\alpha_i(t) = a_r) = \frac{1}{H} \exp\left[\beta \left|\{\alpha_j(t) = a_r : ij \in g(t)\}\right|\right] \tag{20}$$

where $H \equiv \sum_{r'=1}^{q} \exp\left[\beta \left|\{\alpha_j(t) = a_{r'} : ij \in g(t)\}\right|\right]$ is a normalization factor and $\beta \geq 0$ is a parameter that modulates the sensitivity of agents' adjustment to their local environment.

The above formulation has the original one introduced in Section 2 as a particular case when $\beta \to \infty$. In general, even when β is finite, we can continue to motivate it as a best-response dynamics (now, a noisy one) for players who are involved in a pure-coordination game with payoffs as given in (2)-(3). The specific exponential form contemplated by (20) has been amply used in modern evolutionary literature to model gradual adjustment and learning in games (see, for example, Blume (1993), Durlauf (1997), or Young (1998)). It is in the spirit of the well-known formulation of *logistic quantal response equilibrium* proposed by McKelvey and Palfrey (1995), which has been provided with a natural bounded-rationality interpretation by Chen *et al.* (1997). As indicated, the parameter β modulates the noise impinging on agents' adjustment. If $\beta = 0$, noise is the overwhelming force and all actions are chosen with the same probability, irrespectively of local conditions and payoffs. In contrast, in the polar case where β is very large, only if a particular action is a genuine best response is it chosen with a sizable probability.

The adjustment rule given by (20) is also formally identical to the "spin dynamics" postulated by the so-called Potts model in statistical physics, itself a generalization of the canonical Ising model. In this context, $1/\beta$ plays the role of the temperature at which the particle interaction takes place and the q different actions are the possible spins. (See, for example, Vega-Redondo (2006) for a detailed explanation of these models and their relationship to the evolutionary literature.) The Potts model has been studied in detail by physicists and exact solutions for it exist for low-dimensional lattices as well as trees (cf. Baxter (1982)). Recently, the analysis has been extended to random networks by

[13] More specifically, we have studied a formulation where both action adjustment and link creation are subject to random "perturbations", the latter implemented by allowing some small probability $\varepsilon > 0$ that two agents form a link even when they display different actions.

Dorogvtsev *et al.* (2004) and Ehrhardt and Marsili (2004) for random networks. We crucially rely on the latter in our subsequent analysis.

For finite β, the present generalized model is substantially more complex than the particular version of it studied in Section 3. Consequently, we are unable to obtain an exact characterization of its stable equilibria and thus have to base our analysis on some simplifying assumptions. Naturally, this implies that the solution we arrive at can no longer be regarded as a fully accurate description of the long-run behavior of the model. The entailed approximation, however, turns out to be quite effective since, as we shall explain, it matches very closely the results obtained from numerical simulations for large populations.

The key simplifying assumption we make is that the network prevailing at any given point in time is a random network suitably characterized by a degree distribution $\mathbf{p} \equiv \{p(k)\}_{k=0}^{\infty}$ that specifies the fraction of nodes $p(k)$ that display each possible degree k. The defining property of a *random network* is the absence of statistical correlations. Thus, in particular, it is presumed that the degree of a node is stochastically independent of any of its neighbors. Such a property does not strictly hold in the generalized context – only approximately so.[14] This is why the random-network postulate must be viewed, in this case, as a convenient, but not fully accurate, description of the system at any point in time.

In its first step, the analysis requires the specification of the law of motion for the degree distribution $\mathbf{p}$. As before, it embodies two subprocesses: link creation and link destruction. The latter is just as before: every existing link disappears at a constant rate λ, which again we normalize to unity. Link creation, on the other hand, depends on the probability that any two agents who meet and have the potential of creating a new link happen to display the same action. For any given agent/node i, this *ex ante* probability depends on its degree $z_i = k$, so we denote it by $\pi(k)$. In essence, this probability results from the combination of the following three constituent probabilities:

(i) the (unconditional) probability ζ that, when node i selects another node at random, the latter belongs to the (unique)[15] giant component of the network;

(ii) the conditional probability $\gamma(k)$ that node i of degree k belongs to the giant component;

(iii) the conditional probability $\mu(k)$ that, if node i of degree k does belong to the giant component, its action coincides with that of a randomly selected node in that component.

The first probabilities, ζ and $\gamma(k)$ for each k, only depend on the underlying degree distribution $\mathbf{p}$. The probabilities $\mu(k)$, on the other hand, depends both

[14] By way of illustration, one of the features of the process that introduces internode correlations can be explained as follows. First note that, in accord with intuition, one finds that the postulated action dynamics leads high-degree nodes to exhibit, on average, stronger social "conformity" than lower-degree nodes. That is, they have a higher probability of choosing the action that is in the majority in the population. This in turn implies that links between high degree nodes will be formed with higher probability (i.e. at a higher rate) than between lower-degree nodes. In the end, therefore, positive degree correlations will tend to arise, high-degree nodes being more likely to be connected to other high-degree nodes than what is prescribed by the *unconditional* average.

[15] As explained in Section 3, if a giant component exists in the long run, it is unique.

on $\mathbf{p}$ and the value of β in (20) that modulates the pressure towards local conformity induced by the action dynamics. In Appendix B, we outline how such probabilities are obtained. Here it will suffice to advance that, in order to compute the probabilities in (i)-(ii), one can directly use the standard techniques of the modern theory of random networks. And concerning the probability in (iii), we rely on the aforementioned solution of the Potts model in random networks that has been developed by Ehrhardt and Marsili (2004).

Thus let ζ, $\gamma(k)$, and $\mu(k)$ for each k be the probabilities prevailing at some point in time when the underlying network is modelled as a random network with degree distribution $\mathbf{p}$. Then, the probability $\pi(k)$ that a randomly chosen node of degree k meets a node that displays its own action is simply given by:

$$\pi(k) = (1 - \zeta\gamma(k))\frac{1}{q} + \zeta\gamma(k)\mu(k). \tag{21}$$

The second term in the right-hand side of (21) corresponds to the event that some arbitrary node i of degree k happens to be part of the giant component and it meets another node j in that same component. In that case, the link between i and j is formed with probability $\mu(k)$, i.e. the probability that both display the same action.[16] The first term, on the other hand, contemplates what happens when either node i or/and the node j it meets are *not* in the giant component. Then, with probability essentially one in a large random network, both nodes are in different components. This implies that they will only display the same action (and thus form a link) "by chance," i.e. with probability $1/q$ since there are q possible actions.

Given the probabilities $\pi(k)$ specified in (21), the evolution of the degree distribution in (slow) time can be modelled through the following differential equation:

$$\begin{aligned} \dot{p}(k) &= (k+1)p(k+1) + 2\eta p(k-1)\pi(k-1) \\ &\quad -kp(k) - 2\eta p(k)\pi(k) \end{aligned} \tag{22}$$

where again we dispense with the time index for notational simplicity. The first two terms in the right-hand side of (22) reflect the inflow into the frequency of nodes of degree k. This inflow consists of those nodes that had degree $k+1$ and lost one of its links (which happens at a unit rate per link), combined with the rate at which nodes with degree $k-1$ form a new link. (In the latter respect, recall that the factor of 2 accounts for the fact that a link is created if either the node in question receives the initiating opportunity or some other node does.) On the other hand, the two last terms embody the opposite flow that decreases the frequency of nodes of degree k when these nodes either loose or create a link.

We are interested in characterizing the pair $\mathbf{p}^* = \{p^*(k)\}_{k=0}^{\infty}, \boldsymbol{\pi}^* = \{\pi^*(k)\}_{k=0}^{\infty}$ that defines a stationary point of the dynamical system. Such a stationarity embodies a twin requirement. First, given $\mathbf{p}^*$, the corresponding $\boldsymbol{\pi}^*$ must suitably

[16] Of course, this presumes that the link between i and j is not already in place, which is an event that can be essentially ignored in large populations.

characterize the long-run coordination probabilities induced by the fast dynamics governed by (20) – see above. Second, given $\boldsymbol{\pi}^*$, the degree distribution $\mathbf{p}^*$ must define a stationary point of (22). This latter requirement simply amounts to stating that, for all $k = 1, 2, ...,$

$$(k+1)p^*(k+1) + 2\eta p^*(k-1)\pi^*(k-1) = kp^*(k) + 2\eta p^*(k)\pi^*(k).$$

This defines a system of difference equations that can be solved recursively as follows:

$$p^*(1) = 2\eta p^*(0)\pi^*(0) \tag{23}$$

along with

$$p^*(k+1) = \frac{2\eta p^*(k)\pi^*(k) + kp^*(k) - 2\eta p^*(k-1)\pi^*(k-1)}{(k+1)} \qquad (k = 1, 2, ...), \tag{24}$$

once we impose the normalization $\sum_{k=0}^{\infty} p^*(k) = 1$.

As for the effect of η on the long-run behavior of the system, the analysis of the present generalized model delivers essentially the same performance as that obtained for the basic model. We choose, therefore, to focus our discussion on the role played by the new parameter β that marks the only difference with the previous framework (as explained, the original model obtains as a particular case when $\beta \to \infty$). Interestingly, we find that changes in β induce the same qualitative pattern of long-run behavior as observed before for changes in η. The conclusions are depicted in Figure 3.

Figure 3 shows that, as one implements gradual changes in the noise impinging on action adjustment (which can be suitably parametrized by $1/\beta$), the long-run behavior of the system displays the same three features that are obtained for changes in η. That is, both connectivity as well as social conformity exhibit sharp and resilient transitions across multiple equilibria as the noise level changes "slightly" around certain thresholds. And again, we find that the theoretical predictions are well supported by numerical simulations, even though in this case the analytically solved model can only be conceived as an approximate description of the system dynamics.

5 Summary and Conclusions

The main insight delivered by our model can be succinctly described as follows. In a scenario where the social network (co-)evolves alongside nodes' adjustment to coordinate with their partners, cross-reinforcing effects between these two dimensions can generate a rich dynamic performance. The focus of the model has not been on the coordination issue *per se* but on how the struggle of agents' to align their behavior affects the long-run network configuration. We have seen, specifically, that small changes in the underlying parameters can trigger sharp and resilient transitions between a sparse network and a connected one with complex architecture, in either direction.

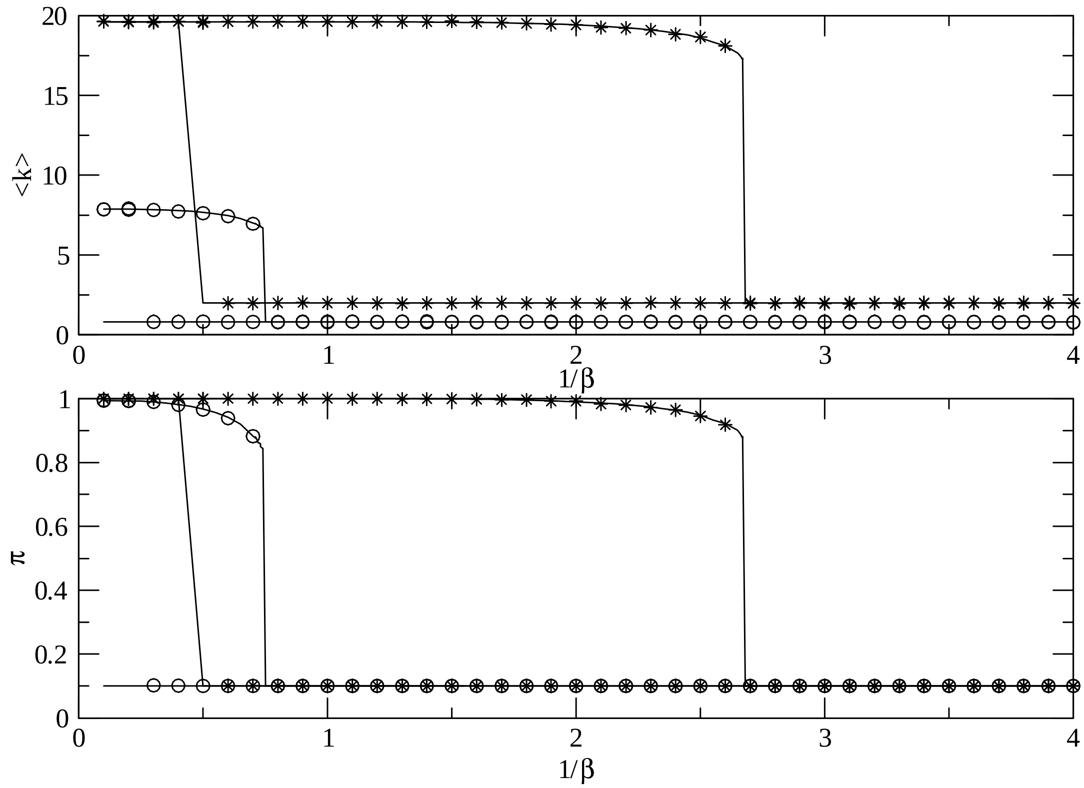

Figure 3: The upper panel plots the average connectivity $\langle k \rangle$ predicted by the model against the noise level $1/\beta$, for $q = 10$ and two different values of η, i.e. $\eta = 4$ (lower curves), $\eta = 10$ (higher curves). The solid lines trace the theoretical prediction while the points represent simulation results for $n = 1000$.The lower panel displays analogous results for the average probability $\langle \pi \rangle$ that two randomly chosen nodes display the same action.

The model is simple and stylized but starkly produces the performance that has been highlighted in a number of social contexts (recall the Introduction). Admittedly, some of the assumptions of the model are quite extreme and should be relaxed. We have already undertaken a first step in this direction when turning from the noiseless model of Section 2 to the context studied in Section 4 where adjustment is noisy and thus action convergence imperfect within each network component. Other interesting extensions (e.g. the relaxation of the asymmetry between the speeds of action and network adjustment contemplated here) are left for future research.

Appendices

A Proofs for the basic model

Proof of Proposition 2: Let us consider a typical node $i \in G_r$. The formation of link ij, where $j \neq i$ is chosen at random, occurs at a rate

$$\gamma_{ij} = 2\eta \frac{1}{N-1} P\{\alpha_j(t) = a_r\} \tag{25}$$

where 2η is the combined rate at which either i and j may initiate the link formation process and the factor $1/(N-1)$ is the probability that, in each respective case, j (or i) is chosen as the recipient of the link formation attempt initiated by i (or j). Finally note that, since the node j is chosen at random, we can write $P\{\alpha_j(t) = a_r\} = (|G_r|-1)/(N-1) + |G_0|/(Nq) \to n_r + n_0/q$ as $N \to \infty$.

Given any $t > 0$, let t_η be the latest previous time ($t_\eta < t$) at which an ij-link formation event occurred. Now given such t_η, let t' be the first subsequent time ($t' > t_\eta$) at which an event of decay for the ij-link occurs. Clearly, we can write

$$P\{ij \in g(t)\} = P\{t' > t\}.$$

Hence, since the ij-link creation events occur at the rate γ_{ij} given in (25) while its decay events occur at the rate $\lambda = 1$, we have:

$$P\{ij \in g(t)\} = \frac{\gamma_{ij}}{\gamma_{ij}+1} \simeq \frac{2\eta}{N}\left(n_r + \frac{n_0}{q}\right), \tag{26}$$

where the last expression retains the leading term in the limit $N \to \infty$. The events $\{ij \in g(t)\}$ for $j \neq i$ are independent, and each occurs with probability (26). Hence the conditional distribution of z_i is Poisson, as given by (8). $\square$

Proof of Proposition 3: The proof easily follows from (10). Since it is always the case that $c_r > 0$, stationarity requires that $c_r e^{-c_r} = 2\eta n_0/q$ for all

$r = 1, 2, ..., q$. The function $f(c) = ce^{-c}$ has a single maximum in $c = 1$ and it vanishes for $c \to 0$ and $c \to \infty$. Hence, for $2\eta n_0/q < 1/e$, each of the equations in (10) has two solutions in c_r, which we call $c_+ > 1$ and $c_- < 1$. Then, (13) results from taking the sum of (9) over $r = 1, \ldots, q$ and using $n_0 \equiv 1 - \sum_{r=1}^{q} n_r$. Finally, (14) derives from $p(k) = \sum_r n_r p_r(k)$ for $k > 0$ and $p(0) = n_0$, while the average degree (15) is obtained by direct calculation. □

Proof of Proposition 5: Let us write (7) as $\dot{n}_r = f(n_r, n_0)$, $r = 1, 2, ...q$, where each $p_r(1)$ is regarded as a function of n_r and n_0, through its dependence on c_r in (9). We focus on a stationary solution $\mathbf{n}^\star$ and consider small perturbations around this solution given by $n_r(t) = n_r^\star + \epsilon_r(t)$ with $\epsilon_r \ll 1$ for each $r = 1, 2, ...q$. Then, to linear order, ϵ_r satisfies

$$\dot{\epsilon}_r = \left.\frac{\partial f}{\partial n_r}\right|_{n_r^\star} \epsilon_r - \left.\frac{\partial f}{\partial n_0}\right|_{n_r^\star} \sum_{r'=1}^{q} \epsilon_{r'} \tag{27}$$

where we use $n_0 = 1 - \sum_{r'} n_{r'}$. The partial derivatives are computed at the stationary point $\mathbf{n}^\star$ using the chain rule, e.g.

$$\frac{\partial f}{\partial n_r} = \frac{2\eta}{q} n_0 - p_r(1) - n_r \frac{\partial p_r(1)}{\partial c_r} \frac{\partial c_r}{\partial n_r} \qquad (r = 1, 2, ...q).$$

A straightforward calculation, using the equilibrium condition at $n_r^\star$, allows us to write (27) as

$$\dot{\epsilon}_r = -\sum_{r'=1}^{q} T_{r,r'} \epsilon_{r'} \tag{28}$$

where the matrix $\mathbf{T}$ has elements

$$T_{r,r'} = \frac{c_r(1 - c_r)}{\exp(c_r) - 1} \delta_{r,r'} + \frac{c_r}{q} \left[1 + \frac{c_r}{\exp(c_r) - 1}\right] \tag{29}$$

and $\delta_{r,r'}$ is the Kronecker delta function (i.e. $\delta_{r,r'} = 1$ if $r = r'$ and $\delta_{r,r'} = 0$ otherwise). A solution $\mathbf{n}^\star$ is stable if all eigenvalues of $\mathbf{T}$ have a positive real part.

For the equilibrium with $L^+ = 0$, $c_r = c_- = 2\eta/q$ for all r. In this case, the matrix $\mathbf{T}$ has one eigenvalue $\mu_1^{(L^+=0)} = c_-/(1 - e^{-c_-})$ corresponding to the constant eigenvector $v_r^{(1)} = 1$. The $q - 1$ vectors $\mathbf{v}$ orthogonal to $v_r^{(1)}$ (i.e. with $\sum_r v_r = 0$) are also eigenvectors of $\mathbf{T}$, with eigenvalue equal to $\mu_2^{(L^+=0)} = c_-(1 - c_-)/(e^{c_-} - 1)$. Both eigenvalues are positive as long as $c_- < 1$, i.e. for $2\eta < q$. Hence the $L_+ = 0$ solution is always stable when it exists.

For equilibria with $L^+ \geq 2$ we find that perturbations of the form $\epsilon_r(t) = a(t)v_r$, with $v_r = 0$ for classes r with $c_r = c_-$ and $\sum_r v_r = 0$, grow exponentially in time. Indeed, $\mathbf{v}$ is an eigenvector of $\mathbf{T}$ with eigenvalue $\mu^{(L_+>1)} = -c_+(c_+ - 1)/(e^{c_+} - 1) < 0$. Hence we have that $a(t) = a(0)e^{-\mu t} \to \infty$ as $t \to \infty$. There are $L^+ - 1$ independent vectors $\mathbf{v}$ of this type, which means there are $L_+ - 1$ such unstable directions. These vectors describe instabilities where an imbalance in the distribution of nodes between classes with $c_r = c_+$ increases exponentially.

In order to discuss the case $L^+ = 1$, let the class $r = 1$ be one with $c_r = c_+$ while for $r > 1$ we have $c_r = c_-$. There are $q - 2$ eigenvectors of $\mathbf{T}$ with $v_1 = 0$ and $\sum_{r>1} v_r = 0$. Their eigenvalue is $\mu^{(L^+=1)} = c_-(1-c_-)/(e^{c_-} - 1) > 0$, which means that all these are stable modes. The remaining two eigenvalues have the form $v_1 = u$ and $v_r = w$ for $r > 1$. This reduces the problem to that of computing the eigenvalues of a 2×2 matrix

$$\tilde{\mathbf{T}} = \begin{pmatrix} \frac{c_+}{q}\left(1 + \frac{c_+}{e^{c_+}-1}\right) - \frac{c_+(c_+-1)}{e^{c_+}-1} & \frac{(q-1)c_+}{q}\left(1 + \frac{c_+}{e^{c_+}-1}\right) \\ \frac{c_-}{q}\left(1 + \frac{c_-}{e^{c_-}-1}\right) & \frac{(q-1)c_-}{q}\left(1 + \frac{c_-}{e^{c_-}-1}\right) + \frac{c_-(1-c_-)}{e^{c_-}-1} \end{pmatrix}$$

As a a preliminary step, we need to confirm that the eigenvalues of $\tilde{\mathbf{T}}$ are real. To see this, denote by a_{ij} $(i,j = 1,2)$ the elements of the matrix $\tilde{\mathbf{T}}$. Its eigenvalues are the solution of the equation $(a_{1,1} - \mu)(a_{2,2} - \mu) - a_{1,2}a_{2,1} = 0$, i.e.

$$\mu = \frac{a_{1,1} + a_{2,2} \pm \sqrt{(a_{1,1} - a_{2,2})^2 + 4a_{1,2}a_{2,1}}}{2}.$$

Since $a_{1,2}a_{2,1} \geq 0$, the desired claim follows.

First, we consider the case when η is large. Then, the trace $\mathrm{Tr}\tilde{\mathbf{T}} = \mu_1 + \mu_2$ of $\tilde{\mathbf{T}}$ is positive since $\mathrm{Tr}\tilde{\mathbf{T}} \to c_+/q$ when $\eta \to \infty$. Thus if the determinant $\mathrm{Det}\tilde{\mathbf{T}} = \mu_1\mu_2$ is also positive, then we may conclude that both eigenvalues of $\tilde{\mathbf{T}}$ are positive in this case, i.e. the corresponding solution is stable. A straightforward calculation leads to

$$\begin{aligned} \det\tilde{\mathbf{T}} &= \frac{c_+c_-}{q(1-e^{-c_+})(1-e^{-c_-})}\left[(1-c_-)e^{-c_-} - (q-1)(c_+-1)e^{-c_+}\right] \\ &= \frac{2\eta n_0(2\eta - qc_+c_-)}{q^2(1-e^{-c_+})(1-e^{-c_-})}. \end{aligned} \tag{30}$$

Here the second expression (30) is derived by using the equilibrium conditions $c_\pm e^{-c_\pm} = 2\eta n_0/q$ and $c_+ + (q-1)c_- = 2\eta$. The above expression can be rewritten as follows:

$$\det\tilde{\mathbf{T}} = \frac{8\eta n_0(c_+ - c_-)}{q^2(1-e^{-c_+})(1-e^{-c_-})}\frac{d\psi}{d\Delta} \tag{31}$$

where recall that ψ stands for the function given by the left hand-side of (19) and we use the fact that $\frac{d\psi}{d\Delta} = (2\eta - qc_+c_-)/(4\Delta)$. Since at equilibrium $\psi(\Delta) = \eta$, we have that $\frac{d\psi}{d\Delta} = \frac{d\eta}{d\Delta}$ and thus it follows that the solutions where the difference $\Delta = c_+ - c_-$ increases with η are stable whereas any solution where Δ decreases with η is unstable. It is easy to check that, for large η (and thus large Δ), $\frac{d\psi}{d\Delta} > 0$ and therefore the $L_+ = 1$ solution is stable in this case.

Now consider the situation for general η. As η falls from high levels (the formerly considered case), one of the two eigenvalues vanishes and changes sign at the value

$$\check{\eta} = \min_{\Delta} \psi(\Delta)$$

i.e. when $d\psi/d\Delta = 0$. Thus, for all $\eta > \check{\eta}$ the equilibrium with $L_+ = 1$ and the largest Δ is stable. Indeed, it is the unique stable one since the alternative equilibrium with the lowest Δ has $d\psi/d\Delta < 0$ and therefore one of its eigenvalues has a negative real part in view of (31). $\square$

B Technical discussion for the generalized model

In this Appendix, we undertake two tasks. First, we explain the way in which, given the degree distribution $\mathbf{p} = \{p(k)\}_{k=0}^{\infty}$ characterizing a random network, one can compute the probabilities ζ and $\gamma(k)$ used in (21). Second, we outline the considerations that underlie the additional probabilities $\mu(k)$ also used in that expression.

1. The determination of ζ and each $\gamma(k)$ starts with the "instrumental" consideration of a related probability applied to links, that is denoted by θ This is the probability that, if a particular *link* is chosen at random, this link does *not* belong to (i.e. does not connect to a node in) the giant component of the network. Of course, if a link does not belong to the giant component, the same must apply to all of the links that spring from either of the two nodes that can be reached from that link. This point can be used to posit a self-consistency condition that determines θ uniquely. To formulate this condition, we need to consider what is the degree distribution of the nodes that can be reached by following a randomly selected link. This degree distribution is generally not equal to $\mathbf{p}$ because those nodes are not just arbitrarily selected nodes but are reached through one of their links. Thus, for the correct *conditional* distribution, denoted by $\tilde{\mathbf{p}} = \{\tilde{p}(k)\}_{k=0}^{\infty}$, the probability $\tilde{p}(k)$ of reaching (or linking to) a node with degree k must be proportional to the number of links that can be used to access it, i.e. proportional to k. Hence we must have $\tilde{p}(k) \propto kp(k)$ so that, after normalization, we obtain:

$$\tilde{p}(k) = \frac{kp(k)}{\sum_{k'=1}^{\infty} k'p(k')} \qquad (k = 1, 2, ...).$$

With this distribution in hand, the indicated self-consistency condition can be written as follows:

$$\theta = \sum_{k=1}^{\infty} \tilde{p}(k)\theta^{k-1}.$$

This condition simply expresses the requirement that the fraction of links that do *not* belong to the giant component is equal to the fraction of nodes that have *none* of its links part of the giant component. Having computed θ in this fashion, the probabilities of interest, ζ and $\gamma(k)$ for each k, can be simply computed as follows:

$$\gamma(k) = 1 - \theta^k \tag{32}$$

$$\zeta = 1 - \sum_{k=1}^{\infty} p(k)\theta^k \tag{33}$$

where note that, in computing ζ, we rely on the unconditional degree distribution $\mathbf{p}$ since the node under consideration is chosen fully randomly, i.e. not by following a particular link.

2. Next, we verbally summarize the considerations that underlie the determination of the conditional probabilities $\mu(k)$ in (21). Given the degree distribution $\mathbf{p}$, recall that each $\mu(k)$ is the probability that, if a node with degree k belongs to the giant component, its action coincides with that of a randomly selected node in that component. Its determination requires solving for the (long-run) equilibrium of the fast action dynamics given by (20) on a given random network. As explained in the text, the problem is isomorphic to finding the solution (also in a random-network context) of the model in statistical physics known as the Potts model.

The Potts model (or its simplified version with two spins/actions known as the Ising model) has been solved for lattice networks of dimension 1 or 2. For higher dimensions, however, no analytical solution is currently known and only approximate methods (such as mean-field analysis) have been developed. The alternative kind of regular networks for which an exact solution is available are the so-called Bethe lattices – essentially, they are infinite trees with a uniform branching factor. The key feature of Bethe lattices allowing one to obtain an analytical solution is that, just as it happens in a one-dimensional lattice, a tree network displays no closed loops. This avoids the severe complications that make it so difficult to reach an exact solution of the model as the number of dimensions increase. The interested reader can refer to Baxter (1982) for a detailed discussion of these models.

For our present purposes, the key point to note is that, locally, random networks are essentially trees with common (albeit random) branching conditions at every node. Heuristically, therefore, one may think of a random network as a "stochastic Bethe lattice" and obtain the solution of the Potts model in this context by an appropriate generalization of the techniques used for (deterministic) Bethe lattices. This is indeed the idea exploited by Ehrhardt and Marsili (2004) to obtain an analytical solution of the Potts model for an arbitrary random graph. A key component of the solution is what is known as the *magnetization* of the system – essentially, a measure of spin (or "action") conformity among the whole "population" of nodes. In this paper we rely on it to compute the key probabilities $\mu(k)$ for each k. These are then combined with the other probabilities prevailing at every point in slow time – i.e. η and $\gamma(k)$ for every k – in order to define the network dynamics and then, as described in the text, eventually solve for the long-run equilibria.

References

[1] van Alstyne, M. (1997): "The state of network organization: a survey in three frameworks," *Journal of Organizational Computing and Electronic Commerce* 7, 83-151.

[2] Bala, V. and S. Goyal (1999): "A non-cooperative model of network formation," *Econometrica* **68**, 1181-1230.

[3] Barabási, A.-L. and R. Albert (1999): "Emergence of scaling in random networks", *Science* **286**, 509-12.

[4] Baxter, R.J. (1982): *Exactly Solved Models in Statistical Mechanics*, London: Academic Press.

[5] Blume, L. (1993): "The statistical mechanics of strategic interaction", *Games and Economic Behavior* **4**, 387-424.

[6] Bollobás, B. (2001): *Random Graphs* (2nd. Edition; First Edition 1985), Cambridge (U.K.): Cambridge University Press.

[7] Burt, R. S. (1992): *Structural Holes: The Social Structure of Competition*, Cambridge: Harvard University Press.

[8] Calvó-Armengol, A. and M. O. Jackson (2004): "The effects of social networks on employment and inequality," *American Economic Review* **94**, 426-454.

[9] Castells, M. (1996): *The Information Age: Economy, Society, and Culture, Volume I: The Rise of the Network Society*, Massachusetts: Blackwell Publishers.

[10] E. J. Castilla, H. Hwang, E. Granovetter, and M. Granovetter (2000): "Social Networks in Silicon Valley," in C.-M. Lee, W. F. Miller, M. G. Hancock, and H. S. Rowen, editors, *The Silicon Valley Edge*, Stanford: Stanford University Press.

[11] Chen, H.-C., J. W. Friedman, and J.-F.Thisse (1997): "Boundedly rational Nash Equilibrium: a probabilistic choice approach," *Games and Economic Behavior* **18**, 32-54.

[12] Clear, T. R., D. R. Rose, and J. A. Ryder (2001): "Incarceration and the community: the problem of removing and returning offenders," *Journal of Political Economy* **47**, 335-51.

[13] Crane, J. (1991): "The epidemic theory of ghettos and neighborhood effects on dropping out and teenage childbearing", *American Journal of Sociology* **96**, 1226-1259.

[14] Delapierre M. and L. Mytelka (1998): "Blurring boundaries: new interfirm relationships and the emergence of networked, knowledge-based oligopolies," in M. G. Colombo (ed.), *The changing boundaries of the firm*, Routledge Press, London.

[15] Dorogovtsev, S.N., A.V. Goltsev, and J.F.F. Mendes (2004): "Potts model on complex networks, *European Physical Journal B* **38**, 177.

[16] Durlauf, S. (1997): "Statistical approaches to socioeconomic behavior," in Arthur, W. B., S.N. Durlauf, and D. A. Lane (eds.), *The Economy as an Evolving Complex System II*, Reading, Ma; Addison-Wesley, 81-104.

[17] Ellison, G. (1993): "Learning, local interaction, and coordination," *Econometrica* **61**, 1047-1071.

[18] Ely, J. (2003): "Local Conventions," *Berkeley Electronic Press Journals, Advances in Theoretical Economics* **2**,1,1.

[19] Erdös, P. and A. Rényi (1959): "On random graphs I," *Publicationes Mathematicae Debrecen* **6**, 290-297.

[20] Erdös, P. and A. Rényi (1960): "On the evolution of random graphs," *Publ. Math. Inst. Hung. Acad. Sci.* **5**, 17-61.

[21] Ehrhardt, G. C. M. A. and M. Marsili (2004): "Potts model on random trees," preprint cond-mat/041126v1.

[22] Fafchamps, M. and S. Lund (2003): "Risk-sharing networks in rural Philippines," *Journal of Development Economics* **71**, 261-87.

[23] C.W. Gardiner (2004): *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences*, Berlin: Springer-Verlag.

[24] Glaeser, E. L., H. D. Kallal, J. A. Scheinkman, A. Shleifer (1992): "Growth in Cities," *The Journal of Political Economy* **100**, 1126-1152.

[25] Goyal S., M. van der Leijy, J. L. Moraga-González (2003): "Economics: an emerging small world?," forthcoming in *Journal of Political Economy.*

[26] Granovetter, M. (1974): *Getting a Job: a Study on Contacts and Careers*, Chicago: Chicago University Press.

[27] Granovetter (1985): "Economic action and social structure: the problem of embeddedness," *American Journal of Sociology* **91**, 481-510.

[28] Grossman, J. (2002): "The evolution of the mathematical research collaboration graph," *Congressus Numerantium* **158**, 201-212.

[29] Hagedoorn, J. (2002): "Inter-firm R&D partnerships: an overview of major trends and patterns since 1960", *Research Policy* **31**, 477-492.

[30] Hagedoorn J., A. Link, and N. S. Vonortas (2000): "Research partnerships," *Research Policy* **29**, 567-586.

[31] Jackson, M. O. (2005): "A survey of models of network formation: stability and efficiency," in Demange G. and M. Wooders (eds.), *Group Formation in Economics; Networks, Clubs and Coalitions*, Cambridge: Cambridge University Press.

[32] Jackson, M. O. and A. Wolinsky (1996): "A Strategic model of economic and social networks," *Journal of Economic Theory* **71**, 44-74.

[33] Jackson, M. O. and B. Rogers (2005): "Search in the Formation of Large Networks: How Random are Socially Generated Networks?" California Institute of Technology, HSS Working Paper no. 1216.

[34] Jacobs, J. (1984): *Cities and the Wealth of Nations: Principles of Economic Life*, New York : Random House.

[35] Kirman, A., D. K. Herreiner, and G. Weisbuch (2000): "Market organization and trading relationships," *Economic Journal* **110**, 411-436.

[36] Krackhardt, D., and J. R. Hanson (1993): "Informal networks: the company behind the chart," *Harvard Business Review* **71**, 104-11.

[37] Kranton, R. and D. Minehart (2001): "A theory of buyer-seller networks," *American Economic Review* **91**, 485-508.

[38] Marsili, M., F. Vega-Redondo, and F. Slanina (2004): "The rise and fall of a networked society: a formal model," *Proceedings of the National Academy of Sciences, USA* **101**, 1439-43.

[39] McKelvey, R. D. and T. R. Palfrey (1995): "Quantal response equilibria for normal form games," *Games and Economic Behavior* **10**, 6-38.

[40] Molloy M. and B. Reed (1995): "A critical point for random graphs with a given degree sequence", *Random Structures and Algorithms* **6**, 161-79.

[41] Montgomery, J. (1991): "Social networks and labor market outcomes: toward an economic analysis," *American Economic Review* **81**, 1408-1418.

[42] Murgai, R., P. Winters, E. Sadoulet, and A. de Janvry (2002): "Localized and incomplete mutual insurance," *Journal of Development Economics* **67**, 245-274.

[43] Newman, M. E. J. (2001): "The structure of scientific collaboration networks," *Proceedings of the National Academy of Sciences USA* **98**, 404-409.

[44] Newman, M. E. J. (2003): "The structure and function of complex networks", *SIAM Review* **45**, 167-256.

[45] Newman, M.E., S.H. Strogatz, and D.J. Watts (2001): "Random graphs with arbitrary degree distributions and their applications," *Physical Review E* **64**, 026118.

[46] Orsenigo L., F. Pammolli, M. Riccaboni (2001): "Technological change and network dynamics: lessons from the pharmaceutical industry," *Research Policy* **30**, 485-508.

[47] Saxenian, A. (1994): *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Cambridge, Mass.: Harvard University Press.

[48] Vega-Redondo, F. (2006): *Complex Social Networks*, forthcoming in the *Econometric Society Monograph Series*, Cambridge University Press.

[49] Watts D.J. and S. H. Strogatz (1998): "Collective dynamics of 'small-world' networks", *Nature* **393**, 440-42**.**

[50] Young, P. (1998): *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton, NJ: Princeton University Press.