Int J Game Theory (2006) 34:383–397
DOI 10.1007/s00182-006-0025-6

ORIGINAL ARTICLE

Diffusion and growth in an evolving network

George Ehrhardt · Matteo Marsili · Fernando Vega-Redondo

Received: 15 February 2005 / Published online: 8 September 2006

Abstract We study a simple model of a population of agents whose interaction network co-evolves with knowledge diffusion and accumulation. Diffusion takes place along the current network and, reciprocally, network formation depends on the knowledge profile. Diffusion makes neighboring agents tend to display similar knowledge levels. On the other hand, similarity in knowledge favors network formation. The cumulative nonlinear effects induced by this interplay produce sharp transitions, equilibrium co-existence, and hysteresis, which sheds some light on why multiplicity of outcomes and segmentation in performance may persist resiliently over time in knowledge-based processes.

Keywords Network formation · Diffusion · Transition · Hysteresis · Growth · Social norms

JEL Classification Numbers D83 · D85 · O33

1 Introduction

The generation and diffusion of new knowledge is typically viewed as a crucial ingredient in processes of economic growth. In the language of Quah (1996),

G. Ehrhardt · M. Marsili
The Abdus Salam International Centre for Theoretical Physics, Trieste, Italy

F. Vega-Redondo (✉)
Instituto Valenciano de Investigaciones Económicas, Universidad de Alicante, Alicante, Spain
e-mail: vega@merlin.fae.ua.es

F. Vega-Redondo
University of Essex, Colchester, UK

for example, the sustained growth of an economy involves the composition of two mechanisms: "pushing back" (say, the frontier of technological know-how) and "catching up" (approaching the leaders' performance). Indeed, one can argue that, in an abstract sense, most processes of knowledge buildup in social environments display these two sides to it: individually-fuelled advances, complemented by interaction-mediated diffusion.

Here we study a very stylized model of the interplay between those twin phenomena (knowledge generation and diffusion) in a context where the channels of dissemination of knowledge are given by a network that *co-evolves* endogenously over time. To fix ideas, let us think of knowledge as "technological know-how."[1] Then, the key feature that characterises network formation in our model is quite intuitive: new links are created only if the individuals/nodes involved are sufficiently similar or compatible, i.e. as long as they display technological levels that are not too far apart. Existing links, on the other hand, are assumed to vanish at a constant rate. In this latter respect, one possible interpretation is that the environment is subject to some background *volatility* that removes existing links due to, say, obsolescence considerations or capacity constraints.

In addition to such a network dynamics, we posit that there is a process of knowledge diffusion and/or growth taking place on the current network. Specifically, we suppose that "innovation" proceeds independently at each node, while there is a simultaneous process of diffusion among neighbours that tends to close their technological differences.

In the context just outlined, we raise a number of questions:

1. When will the co-evolutionary process succeed in building up a *dense* (i.e. highly connected) network?
2. How will the network emerge when the conditions (i.e. model's parameters) are adjusted so as to make this possible? In particular, do we expect *gradual* (i.e. continuous) or *sharp* (discontinuous) transitions?
3. What is the effect of the process of network formation on the technological *convergence* among the agents involved? How does it influence the overall rate of technological *advance*?

Admittedly, the model and questions we pose are too abstract to be directly useful in understanding particular instances of growth and technological change in socioeconomic environments. They highlight, however, what we believe is an important issue in these phenomena in the real world, namely, the interplay between the technological proficiency of individuals/nodes and their pattern of connectivity, which in turn affects their ensuing prospects of (technological) growth and thus their future connections as well.

To be more specific, and merely by way of illustration, the model might shed light on questions that have triggered an important and heated controversy in

[1] As explained below, however, only one of our two different formulations of diffusion is best conceived as reflecting technological diffusion. The alternative one suggests more naturally an interpretation where the social diffusion concerns either opinions or social norms.

the modern literature of growth and development. At least since the empirical work of Baumol (1986), researchers have debated whether there is an underlying force towards global growth convergence or, alternatively, the world is best viewed as segmented in disjoint "convergence clubs" with very different performance.[2] While the former view has been championed by researchers such as Barro and Sala-i-Martin (1992) and Mankiw et al. (1993), the latter has been stressed by others such as Durlauf and Johnson (1995) or Quah (1997).

Here, of course, we need not enter into the dispute, which is not only empirical but also theoretical. At a theoretical level, the discussion in the mainstream literature has focused on variations of the neoclassical growth model, suitably enriched with "complications" such as externalities, human capital, market imperfections, etc., see Galor (1996) for an excellent discussion of the alternative theoretical approaches. In contrast, our model directs attention to a mechanism that has been mostly ignored by this literature but has been instead underscored by authors in the historico-evolutionary tradition—see, e.g. Abramovitz (1986), Verspagen (1993), and Fagerberg (1994). The so-called *technological-gap approach* to economic growth casts the issue as one of *technological* convergence–divergence. And it tailors the growth or shrinkage of the technological divide between countries to the laggards' (in)ability to close the technological gap. In this somewhat eclectic literature, a wide array of different factors are singled out as the main culprits in the case of catch-up failure. But, formally or informally,[3] the existence of essential nonlinearities in the diffusion process always play a prominent role in the story. As we shall see, such nonlinearities are also central to the rich growth and diffusion performance exhibited by our model.[4]

More generally, a certain version of our model admits an interpretation of the agents' attributes as the behavior they choose under the local influence of neighbors. The model can then be brought to bear on the important issue of the evolution of social norms and the force towards conformity exerted in a context of local interaction—see, e.g. Ellison (1983). Recently, this issue has been studied in a dynamic context where, as here, the social network itself coevolves with the other dimensions of behavior. See, specifically, Jackson and Watts (2002) and Goyal and Vega-Redondo (2005) , in which agents decide both on whom to connect to as well as their behavior in an underlying coordination game.

2 For a brief but useful review of this literature, see the critical piece by Durlauf (2003).

3 See, for example, how Abramovitz's (1986) heuristic notion of *social capability* is modelled, explicitly as a nonlinearity, by Verspagen (1993).

4 Another important phenomenon to which our model may contribute some insights is the evolution of patterns of collaboration among academic scientists and industry researchers. In the academic realm, recent empirical work on collaboration networks has covered a wide variety of disciplines. For example, Newman (2001) has studied the fields of physics, biomedical research, and computer science, Grossman (2002) that of mathematics, and Goyal et al. (2003) economics. The latter authors report, for example, that in merely two decades the average degree of collaboration has doubled among economic researchers. Interestingly, such a sharp increase in the density of interactions is indeed one of the main consequences of the interplay between knowledge diffusion and network evolution in our model.

In contrast with this evolutionary literature, our model emphasizes the role of complexity, assuming that network formation is subject to persistent and sizable perturbations. Under these conditions, the analysis focuses on the interplay between the (in)ability to reach a dense social network and the rise (or not) of a coherent pattern of social behavior (say, in the form of a uniform convention). An important limitation of our approach is that we do not explicitly model individual incentives. These are accounted for only implicitly through our formulation of diffusion and some "compatibility requirements" on link creation. For a comprehensive survey of the strategic approach to network formation, with incentives playing a central role in the analysis, the reader is referred to Jackson (2005).

Finally, we mention that the present approach to the study of diffusion shares some of its key features with the booming recent field of complex networks [see Albert and Barabási (2002) or Newman (2003) for quite exhaustive surveys]. As in much of this literature, a central concern here is the induced dynamics of network formation. Our theoretical framework, nevertheless, is neither stationary nor growing, cf. Albert and Barabási and Albert (1999). We posit instead a context persistently affected by volatility, identified with a random process of link removal. The model, therefore, bears some similarities with that studied by Marsili et al. (2004). There, however, agents/nodes display no genuine interaction and the only dimension of change pertains to the network itself.

The rest of the paper is organised as follows. In Sect. 2, we present the general model as well as the different particular specifications consistent with it that will be studied in the paper. Section 3 undertakes the analysis of the model through extensive numerical simulations, while some theoretical insights on the resulting behavior are offered in the Appendix. The main body of the paper concludes in Sect. 4 with a summary and a short discussion.

2 The model

Consider a set $\mathcal{N} = \{1, \ldots, n\}$ of n agents who evolve in continuous time $t \geq 0$. Each agent i is characterised by an attribute $h_i(t)$—which we generically call knowledge—capturing her level of expertise or technological development. The agents' interaction network at time t is described by a non-directed graph $g(t) \subset \{ij,\ i,j \in \mathcal{N}\}$, where $ij \in g(t)$ if i and j interact at time t. The network $g(t)$ is the social backbone through which diffusion proceeds, as explained next.

We assume that each agent i receives an attribute update (or upgrade) possibility at a rate ν, meaning that the probability that h_i is updated in the time interval $[t, t + \mathrm{d}t)$ will be $\nu \mathrm{d}t$. If agent i receives such an opportunity at time t, we posit that

$$h_i(t^+) = D\{h_j,\ j \in \mathcal{N}_i(t)\} + \eta_i(t), \tag{1}$$

where t^+ is time immediately after the update and

- $\eta_i(t)$ is a random term capturing the idiosyncratic change of expertise due to i's own (say research) efforts;

- the function $D\{\cdot\}$ captures knowledge diffusion in the current neighborhood $\mathcal{N}_i(t) = \{j : ij \in g(t)\} \bigcup \{i\}$ of agent i (including i herself).

We will take $\eta_i(t)$ to be Gaussian i.i.d. random variables with zero average and variance Δ. Concerning diffusion, on the other hand, we will consider two alternative possibilities that, to fix ideas, we shall respectively label as Best-performance imitation and Merging behavior.

Best-performance imitation (BI): The revising player i attains a level $h_i(t^+)$ equal to the maximum available in her neighbourhood. Formally, this corresponds to the following formulation:

$$D\{h_j,\ j \in \mathcal{N}_i(t)\} = \max_{j\in\mathcal{N}_i(t)} h_j(t), \tag{2}$$

where $\mathcal{N}_i(t)$ includes i and $|\mathcal{N}_i(t)|$ is the number of agents in i's neighborhood.[5]

Merging behavior (MB): The revising player i "merges" her behavior with that displayed by her neighbors. This is formalized by making

$$D\{h_j,\ j \in \mathcal{N}_i(t)\} = \frac{1}{|\mathcal{N}_i|} \sum_{j\in\mathcal{N}_i(t)} h_j(t). \tag{3}$$

The first formulation, imitation of best-performance, is common in evolutionary literature—see, e.g. Vega-Redondo (1997) or Eshel et al. (1998). In our context, as suggested, it may be interpreted as modelling a process of technological growth under individual innovation and network-channeled diffusion. The second formulation, on the other hand, can be conceived as reflecting a process of opinion exchange (with no idea of relative "advance" in the levels displayed by different individuals), as in Weisbuch (2002) or DeMarzo et al. (2003). Alternatively, MB could be viewed as representing a context where the payoffs of interaction are enhanced by conformity (say, in social norms or technological standards) and, therefore, agents adjust their behavior towards that displayed by their neighbors.

The dynamics of Eq. (2) has been extensively studied on finite dimensional lattices and Cayley trees, because of its relation to a number of important physical processes such as turbulence and stochastic surface growth—see Halpin-Healy and Zhang (1995). The properties of (unidirectional) diffusion processes (Eq. 3) on fixed graphs have also been studied extensively (see the Appendix).

Here, in addition to the diffusion (sub)process, the interaction network of agents also evolves according to two further stochastic (sub)processes: link destruction and link creation, as we describe next.

First, each *existing link decays* (i.e. disappears) at a constant rate λ. This process models environmental volatility, e.g. link breakdown due to errors, or the fact that those conditions under which a particular link was formed might

5 Note that, if i has no neighbor, $\mathcal{N}_i = \{i\}$ and hence $D = h_i$ in both cases.

change over time, thus leading one of the two parties to discontinue the relationship.

Second, *links are formed* as follows. At a rate ξ, each agent i is given the opportunity to form a new link. In this event, a partner j is chosen at random in the population and the link ij is created (with probability one, if not already in place) provided $|h_i - h_j| \leq \bar{d}$, where $\bar{d}$ is a parameter of the model. Otherwise (if this inequality is violated), the link is assumed formed only with some probability ϵ, to be conceived as small. The differences in h might also be introduced as one of the factors affecting link decay. While this is certainly plausible for many applications, we prefer to keep our discussion as simple as possible. It is intuitive, however, that if the rate λ at which link ij disappears were an increasing function of $|h_i - h_j|$, the non-linear feedback effects discussed below will be enhanced, thus strengthening our results.

The rate ξ at which agents receive link-formation opportunities can be seen as a proxy for the intensity with which agents meet or, relatedly, the "networking" efforts exerted by agents. The parameters $\bar{d}$ and ϵ, on the other hand, determine how important it is that two agents be similar in order for the link to be formed. In general, we might expect that two agents with very different knowledge levels seldom form a link since it is unlikely that, in the foreseeable future, *both of them* may profit significantly from it. Heuristically, one could relate the magnitude of $\bar{d}$ and ϵ to the extent to which agents are forward looking and patient. To see this, suppose that (3) applies and identify the level h_i displayed by an agent i at some t with her instantaneous payoff. Then, if agents do pay sufficient attention to long-run payoffs, they will understand that, even if it might be unprofitable in the short run to establish a link with a less advanced individual, this may well induce higher payoffs in the future due to an enhanced potential of innovation.

Naturally, the *relative* magnitudes of ξ and λ will play a key role in the model.[6] In a nutshell, the basic intuition of the analysis is embodied by the following points.

1. A dense network promotes a uniform society, which in turn makes it easier to establish new links. Under these conditions, the decay of obsolete links can be efficiently balanced by a vigorous link formation process.
2. When a society has only few connections, agents tend to display very diverse attribute levels. This renders the creation of new relationships typically hard, which makes it impossible to overcome the link decay imposed by volatility.

The aforementioned points suggest that, within the same environment (i.e. identical underlying parameters), two polar configurations can materialise as stable states. In one of them, there is a dense network and an homogeneous society; in the other, the network is sparse and the society very heterogeneous.

[6] Notice that, a time rescaling $t \to t/\lambda$ does not affect the properties of the stationary state. This implies that results only depend on the ratio ξ/λ and ν/λ. Likewise, a rescaling $h_i \to h_i/\sqrt{\Delta}$ implies that results only depend on the ratio $\bar{d}/\sqrt{\Delta}$. Hence, for example, the dependence of the results on $\bar{d}$ can be inferred from their dependence on Δ in a trivial manner.

This is indeed what the numerical experiments in Sect. 3 will show. Furthermore, the contrast between (2) and (3) will provide a useful insight on how the directionality of the diffusion process affects the overall performance of the system. The specific implications of our conclusions for phenomena such as technological growth or social conformity will then be summarized in Sect. 4.

3 Analysis

The interesting region in parameter space is when the diffusion dynamics plays a key role in building the network. Consider the case where the distribution of h_i across the population is extremely broad with respect to $\bar{d}$, so that the probability p that $|h_i - h_j| < \bar{d}$ is negligible. Then link creation occurs at an effective rate $2\xi\epsilon$, the factor 2 coming from the fact that link creation can be initiated by both partners. The rate at which links are removed is instead $\lambda\langle k\rangle$, where $\langle k\rangle$ is the expected degree (i.e. number of neighbours) of an agent. In the stationary state, these processes must balance yielding an average degree $\langle k\rangle = 2\xi\epsilon/\lambda$. Our numerical experiments focus on the case where $\lambda \gg 2\xi\epsilon$, so that the emergence of a dense network is not possible in the absence of the diffusion process. In what follows, we take $\lambda = 1$ and $\epsilon = 0.001$ which ensures that the condition is met when ξ varies in the range 1–20.

Knowledge update is taken to occur at a rate $\nu = 10$, which is relatively fast compared to link decay process. In the polar case of slow knowledge update ($\nu \ll \lambda$) links typically disappear before they can have an impact on the diffusion process. We set $\bar{d} = 2$. The variance of the noise, Δ, is set differently for the two models: For BI (Eq. 2) we set $\Delta = 0.1$ whereas for MB (Eq. 3) we set $\Delta = 1$. The reason for this difference is that MB is much more effective than BI in achieving uniform knowledge levels, hence stronger shocks η_i are needed with the BI model to maintain a distribution of h_i with a spread of order $\bar{d}$.

System sizes of $n = 200, 500$ and $10,00$ agents are simulated. Each depicted point represents a system which was run up to $t = 1,000$ in order to reach a stationary state. Then averages were taken in the time interval $t \in [1,000, 1,100]$. In a first set of simulations, we took the empty network as the initial condition, whereas in a second set, simulations were started from a highly connected network. In an ideal experiment where, starting from low values, networking effort ξ increases very slowly to high values we would expect the system to go through the states visited in the first set of simulations. Likewise, the second data set describes what one may expect if, starting from high ξ, we decrease it very slowly. [7]

Results for the behaviour of network density are shown in Figs 1 and 2 for the cases of directional (BI) and unidirectional (MB) diffusion, respectively. These

[7] In the coexistence region, we found that rare fluctuation sometimes cause the system to jump from the low connected to the high connected state (this can be seen to occur occasionally in Fig. 1). Reverse jumps also may occur in a finite system, though they are much more rare. Hence in a real experiment where the parameter ξ is increased very slowly, the system may jump to the dense network phase earlier than in Figs. 1 and 2.

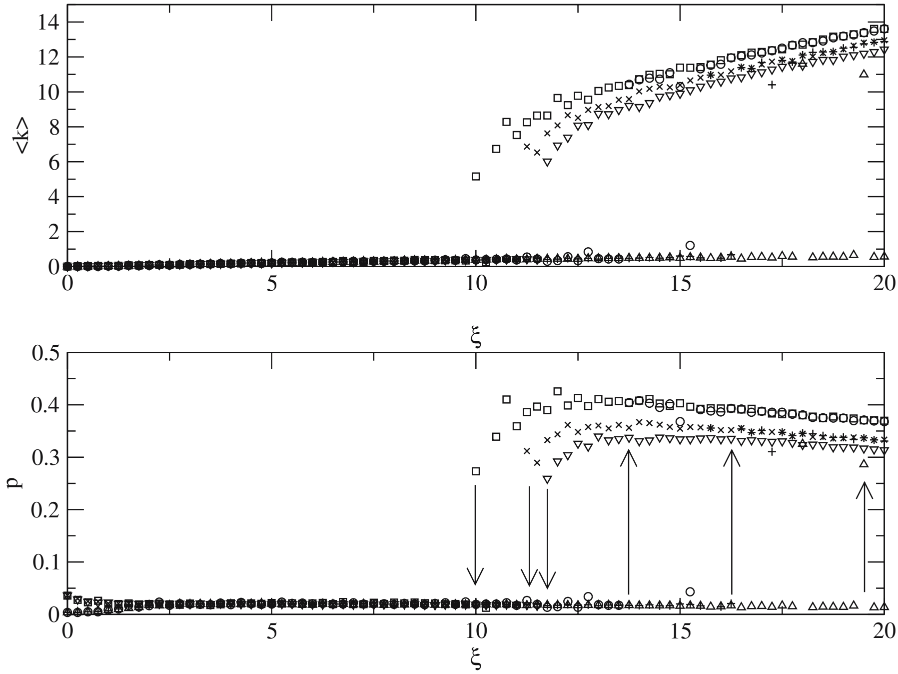

Fig. 1 Plot of the mean connectivity, $\langle k \rangle$, and probability that two nodes' attributes are within $\bar{d}$ of each other, p, as a function of ξ. The model is the BI model of Eq. (2). Plots are for agent numbers of $n = 200$ (*squares and circles*), $n = 500$ (*crosses and plusses*) and $n = 1{,}000$ (*triangles down and triangles up*). *Squares, crosses and triangles down* represent the initially coordinated states. *Circles, plusses and triangles up* represent the initially uncoordinated states. *Arrows* denote the approximate points ξ_1, ξ_2 at which the system jumps between states. The difference between the coordinated and uncoordinated data points at low ξ is due to the very low connectivity and low noise not equilibrating the h values within the runtime of $t = 1{,}100$

report the average degree $\langle k \rangle$ and the probability p that two randomly chosen agents have compatible knowledge levels, $|h_i - h_j| < \bar{d}$. In the stationary state, the rate at which links are lost should balance that at which links form, so that

$$\langle k \rangle = \frac{2\xi}{\lambda}\left[\epsilon + (1-\epsilon)p\right], \tag{4}$$

a relation which is fully confirmed by numerical simulations.

In both figures, we observe the same behavior: As long as ξ is small, only a stationary state with a sparse network (small $\langle k \rangle$) is possible. As the networking effort ξ increases, the global situation does not improve much, because still most attempts to create new links fail due to technological levels being too different (i.e. p small). When ξ gets large enough, however, this situation becomes unstable and a dense network forms in an abrupt manner. The network formation is a stochastic event and it may take place at slightly different values of ξ. Typically in societies with a larger number n of agents the sparse network is found to

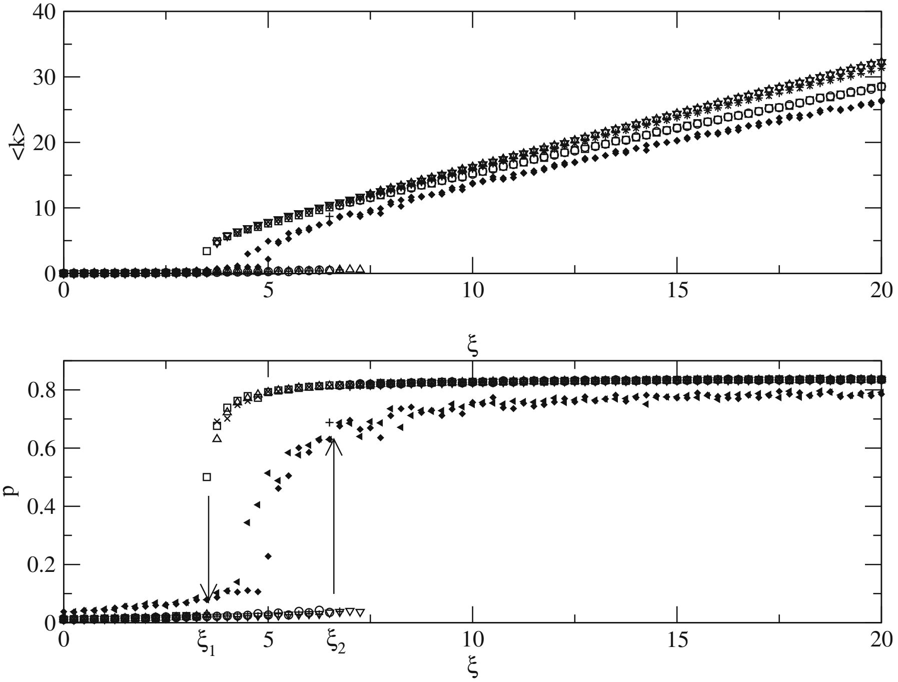

Fig. 2 Plot of the mean connectivity, $\langle k \rangle$, and probability that two nodes' attributes are within $\bar{d}$ of each other, p, as a function of ξ. The model is the MB model of Eq. (3). Plots are for agent numbers of $n = 200$ (*squares and circles*) and $n = 500$ (*crosses and plusses*). *Squares and crosses* represent the initially coordinated states. *Circles and plusses* represent the initially uncoordinated states. *Arrows* denote the approximate points ξ_1, ξ_2 at which the system jumps between states. Also shown are results of simulations started from dense (*filled triangle left*) or sparse (*filled diamond*) network for $\nu = 0.1$ (slow updating of the h_i) with $n = 200$

be more stable than in societies with small n. This is because the flip from the sparse to the dense network phases is triggered by large fluctuations, which are relatively more likely in small systems.

Once a dense network emerges, it remains stable both if the networking effect increases and if it decreases. Indeed, a dense network makes diffusion of knowledge (h_i) very effective, thus narrowing society's spread in h_i. This in turn makes the replacement of obsolete links fast (large p), thus sustaining a densely networked society. This points to the fact that such socio-economic networks not only exhibit a sharp transition but also a resilient one, i.e. after a dense network has formed it remains stable even when external conditions deteriorate. Figures 1, 2 show that the dense network reverts again to a sparse one when the networking effort ξ falls below a threshold ξ_1.

We remark that results depend on the population size n, we also found that, other things being equal, p decreases with Δ. On the other hand, increasing the rate ν of h_i updates makes the population more uniform, hence p increases.

Directionality of diffusion impinges on the results in several ways. First, MB is much more efficient than BI in achieving a uniform society. This is reflected in the values of Δ used in Fig. 2 for unidirectional diffusion, which is ten times larger than that used in Fig. 1. In both cases, the region of ξ where the transitions take place increases with increasing Δ.

Second, and most importantly, while for undirected diffusion (MB) the population average of h_i does not change significantly over time, for directed diffusion of knowledge (BI) we have:

$$\langle h(t) \rangle = \frac{1}{n} \sum_{i=1}^{n} h_i(t) = vt. \tag{5}$$

That is, $h_i(t)$ grows with time at a rate (velocity) v.[8] Figure 3 shows both the ξ dependence of the velocity and also some examples of the distribution of h.

We found that v is an increasing function of the average degree $\langle k \rangle$. One can think of this as agents having more choices of whom to copy "best practice" from and such practice diffusing faster in dense networks. The growth performance of the population then mirrors the behavior of network density, as the parameters change. The sparse network phase corresponds to a "stagnant" situation of slow growth whereas, as shown in Fig. 3, at the transition where a dense network emerges, growth accelerates abruptly. This shows that sharp transitions in growth rates can be related to the transition in the underlying socio-economic network. It is worth remarking that the individual "innovation process" is neutral, i.e. η_i can both be positive or negative and it has zero mean. A positive growth rate occurs solely because positive innovations are selectively transmitted through the network by the diffusion process.

The distributions of h are shown in Fig. 3. Below the transition, the agents are rarely connected and consequently do not share "best performance" so much. This is reflected in the broad distribution of h values seen and also by the fact that the peak of the distribution is far from the "best" agent. Above the transition, the agents are more highly connected and also have a higher v. The nodes in the giant component are much closer in h to the best agent and the distribution is narrower. Nodes not in the giant component get left behind (since their v is much lower) in the lagging smaller peak which can be seen far to the right of the main peak in Fig. 3. Such nodes only catch up again through the unconditional ϵ process, which is relatively slow.

By contrast, the distributions (not plotted) of h for the model based on MB are symmetric[9] and approximately Gaussian. The width of the distribution is much smaller in the highly connected phase, as expected.

[8] This can be understood by observing that Eq. (2) does not depend on absolute time and is invariant under translations in the h direction (i.e. $h_i \to h_i + c$ for all i with c constant). The only non-degenerate solution of the associated equation for the probability distribution of h_i's which is consistent with this property is a travelling wave solution $P(\{h_i\}, t) = f(\{h_i - vt\})$, for a suitably chosen value of v. It is clear from Eq. (2) that, on average, h_i increases at any update, hence $v \geq 0$.

[9] Indeed Eq. (3) enjoys the further symmetry $h_i \to -h_i$ for all i.

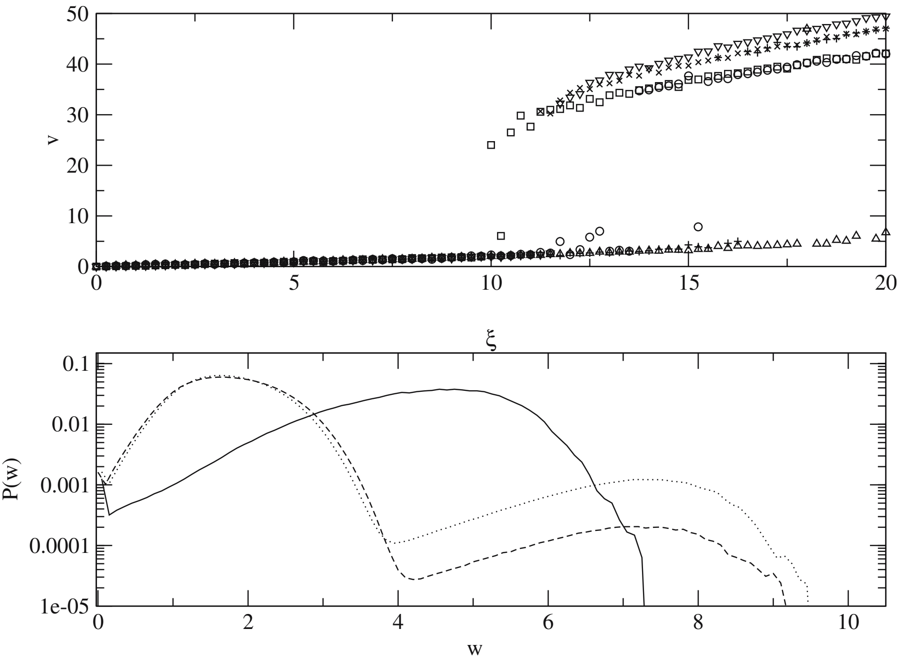

Fig. 3 *Upper* Plot of the mean velocity, v, as a function of ξ for the BI model of Eq. (2). Plots are for agent numbers of $n = 200$ (*squares and circles*), $n = 500$ (*crosses and plusses*) and $n = 1,000$ (*triangles down and triangles up*). *Squares, crosses and triangles* down represent the initially coordinated states. *Circles, plusses and triangles up* the initially uncoordinated states. *Lower* The distributions of the h_i for $n = 1,000$ and three values of ξ, $\xi = 10$ (*solid line*), $\xi = 15$ (*dotted line*) and $\xi = 20$ (*dashed line*), starting in the coordinated state. The histograms are taken from $t = 1,000$ to $t = 2,000$ and are of $w_i(t) = \ln(1 + \max_j(h_j(t)) - h_i(t))$. Note the second smaller peaks in the distributions for $\xi = 15$ and $\xi = 20$. The logarithmic scales were necessary in order to plot the histograms clearly on the same figure

Figure 3 shows that the $\xi = 20$ curve is slightly broader than the $\xi = 15$ curve. This is related to the decrease in p with ξ for large ξ which can be observed in Fig. 1, because p is a decreasing function of the spread of the distribution of h_i. There are two possible factors which act in the direction of broadening the distribution of h and hence may be responsible for this behavior. Firstly, the variance of the maximum of k random variables with unbounded domain increases with k, see Galambos (1987). Given that we have a giant component, this effect will act to broaden the distribution as ξ and hence $\langle k \rangle$ is increased.

Secondly, the velocity also increases with $\langle k \rangle$. Most progress is generated in the most highly connected part of the giant component. Nodes on the periphery of the giant component get "dragged along", their lag increasing with their distance from the core region and with the velocity, v. Thus when v increases these peripheral nodes lag farther behind. Both factors are not at work in the case of Eq. (2): the average of k random variables has a variance which decreases with k and there is no velocity ($v = 0$). Indeed, for Eq. (2), p is a monotonically non-decreasing function of ξ.

4 Summary and conclusions

The paper examines a simple model of network co-evolution in which the current network is the basis for a (relatively fast) diffusion process. The main insight that is gained from our analysis is that such interplay between network change and the diffusion of knowledge or behavior can generate the outcome multiplicity that is observed in many network-based phenomena in the real world, e.g. in technological change and development or the evolution of social norms. The entailed equilibrium multiplicity also displays many of the characteristics that are common in phase transitions of nonlinear dynamical systems. Especially important among them is hysteresis, which may be interpreted as a manifestation of the robustness and history-dependence that impinge on the system dynamics.

The general adjustment dynamics proposed here has been specialized in two different directions. One of them, labelled BI, seems more appropriate to study processes of diffusion of knowledge and growth such as those fuelled by technological change. The second one, referred to as MB, is better suited to model processes of opinion change and norm evolution. In both cases we find the same qualitative behavior, as it concerns sharp transitions, equilibrium multiplicity and hysteresis. But the implications and interpretation of those phenomena is different in each case. Under BI, the transition to a high-connectivity phase entails a marked jump in the overall rate of advance of the system, due to the large synergies then made available by much more intense imitation. Under MB, on the other hand, no such effect in the rate of advance occurs (nor would be really meaningful) and the transition entails a swift move towards social convergence and conformity (in opinions, social norm, etc.).

The theoretical framework studied in this paper is kept simple and stylised, in order to highlight the main forces involved in the resulting phenomenology. However, it can and should be extended in a number of directions. One of them is to enrich the network formation by introducing topology-based considerations of search in meeting new partners. For example, it could be posited that new potential partners are met through current neighbours, making link creation depend on the current network architecture. [In a different context, this approach has been explored by Marsili et al. (2004).] Another interesting extension of the model would be to couple the diffusion process with other phenomena that are economically relevant in growth and development processes. A natural one, of course, is capital accumulation, which then suggests the importance of introducing explicit economic incentives and, possibly, forward-looking considerations into the analysis. This, however, is bound to complicate the analysis very substantially. It also requires the combination of different methodologies and paradigms, an endeavour that usually proves markedly difficult.

Appendix: Some analytical insights for MB diffusion

We detail in what follows simple arguments which provide some insight for the behavior observed in numerical simulations and discussed in the text.

The MB model of Eqs. (1,3) can be analyzed analytically in the limit $\nu \gg 1$. We will assume that the network can be well approximated by an Erdos–Renyi random graph with average degree $\langle k \rangle$. Our strategy will be that of deriving the distribution of $h_i - h_j$ in the stationary state for a random Erdos–Renyi graph with average degree $\langle k \rangle$, and then using this to estimate the probability p that two randomly chosen nodes have $|h_i - h_j| < \bar{d}$. Then Eq. (4) can be used to obtain $\langle k \rangle$. In effect, the distribution of h_i will depend, in general, on the degree k_i of the node. Hence the probability a node gets a new neighbor will also depend on its degree and therefore the degree distribution will deviate from the Poisson law. We assume that these deviations are weak and do not affect the qualitative nature of the results.

In the limit $\nu \gg 1$ the dynamics is well described by a continuous Langevin equation[10]

$$\dot{h}_i = -\frac{\nu}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} (h_j - h_i) + \zeta_i \equiv -\nu \sum_j \mathcal{L}_{i,j} h_j + \zeta_i, \tag{6}$$

where $\langle \zeta_i(t) \zeta_j(t') \rangle = \nu^2 \Delta \delta_{i,j} \delta(t - t')$ and we have introduced the (normalized) Laplacian matrix of the graph $\mathcal{L}$. The dynamics of this model with a fixed network is well known in physics. We observe that (i) the distribution of h_i is Gaussian, as it is a combination of Gaussian variables; (ii) the dynamics is easily integrated in the normal modes of the diffusion operator. In other words, let $\mathbf{v}^\mu$ be the eigenvectors of $\mathcal{L}$, i.e. $\mathcal{L}\mathbf{v}^\mu = \mu \mathbf{v}^\mu$. Then the normal modes $h^\mu = \sum_i v_i^\mu h_i$ satisfy

$$\dot{h}^\mu = -\nu \mu h^\mu - \zeta^\mu, \tag{7}$$

where, in view of the orthogonality of the transformation $i \to \mu$, ζ^μ is again a white noise with the same statistical properties of ζ_i. The fluctuations of h^μ in the stationary state are $\langle (h^\mu - \langle h^\mu \rangle)^2 \rangle = \frac{\nu \Delta}{2\mu}$. Back transforming to the variables h_i one finds that

$$\langle (h_i - \langle h_i \rangle)^2 \rangle = \sum_{\mu > 0} \frac{\nu \Delta}{2\mu} = \frac{\nu \Delta}{2} \int \frac{\mathrm{d}\mu}{\mu} \rho(\mu). \tag{8}$$

Here $\rho(\mu)$ is the density of eigenvalues of the Laplacian matrix and it can be computed for a random graph in the limit $n \to \infty$. No simple closed form is

[10] In order to derive such an equation, fix a small time interval $\mathrm{d}t$. If $\nu \mathrm{d}t \gg 1$ the number of updates on each site will be large and hence, by the central limit theorem, the corresponding increments in the h_i's are well approximated by a deterministic term equal to the expected value of the r.h.s. of Eq. (1), times $\mathrm{d}t$, plus a random Gaussian contribution.

available—see e.g. Dorogovtsev (2003) for details. However we notice that the integral in μ, for Erdos–Renyi graphs, is a function $R(\langle k\rangle)$ of the average degree $\langle k\rangle$ alone. The function $R(c)$ (i) decreases monotonically and (ii) it diverges as $c \to 1^+$, when the giant component vanishes. At our level of approximation, we will assume that for $\langle k\rangle < 1$ the levels h_i are infinitely spread, i.e. $R(c) = \infty$ for $c < 1$. This allows us to write

$$\langle (h_i - \langle h_i\rangle)^2\rangle = \frac{\nu\Delta}{2} R\big(\langle k\rangle\big)$$

and to estimate the probability

$$p = P\{|h_i - h_j| < \bar{d}\} = \theta(\langle k\rangle - 1)\mathrm{erf}\left[\bar{d}/\sqrt{2\nu\Delta R(\langle k\rangle)}\right]. \tag{9}$$

Here $\theta(x) = 1$ if $x > 0$ and $\theta(x) = 0$ otherwise, thus ensuring that $p = 0$ in unconnected graphs ($\langle k\rangle \leq 1$).

This equation, combined with Eq. (4) captures the feedback effects which are responsible for the phenomenology observed in numerical simulations. A graphical solution of these equations shows that there is a value ξ_1 below which only one solution with $\langle k\rangle < 1$ exists and a value ξ_2 above which only a solution with $\langle k\rangle > 1$ exists. In the interval $[\xi_1, \xi_2]$ both solutions are possible.

The sharp dependence of $\langle (h_i - \langle h_i\rangle)^2\rangle$ on $\langle k\rangle$ for $\langle k\rangle \approx 1$, which generates the non-linearities which are responsible for the phase transition, gets smoothed as ν decreases. Indeed, in the opposite limit where the network evolves much faster than the dynamics of h_i, i.e. $\nu \ll \lambda$, the network is completely different each time an h_i is updated. The network is now no longer relevant—it is as if neighbours are drawn at random from the entire population, with a preference for neighbours whose h_j's are within $\bar{d}$. We do not expect a strong dependence of the distribution of h_i's on $\langle k\rangle$ for $\langle k\rangle \approx 1$.

Figure 2 shows results for $\nu = 0.1$—i.e. a slow update of h_i compared to the change of the network. For this plot, the hysteresis region is greatly reduced. Thus the hysteresis appears to rely on the fact that connections (links) last for several updates of h_i .

Acknowledgements We thank Antoni Calvó-Armengol and an anonymous referee for helpful comments. FVR acknowledges support by the Spanish MEC (SEJ2004-02170).

References

Abramovitz M (1986) Catching up, forging ahead, and falling behind. J Econ Hist 46:386–406

Albert R, Barabási A-L (2002) Statistical mechanics of complex networks. Rev Mod Phys 74:47–97

Barabási A-L, Albert R (1999) Emergence of scaling in random networks. Science 286:509–512

Barro R, Sala-i-Martin X (1992) Convergence. J Polit Econ 100:223–251

Baumol WJ (1986) Productivity, convergence, and welfare: what the long-run data show. Am Econ Rev 76:1072–1085

DeMarzo P, Vayanos D, Zwiebel J (2003) Persuasion bias, social influence, and unidimensional opinions. Q J Econ 118:909–968

Dorogovtsev SN, Goltsev AV, Mendes JF, Samukhin AN (2003) Spectra of complex networks. Phys Rev E 68:046109
Durlauf S (2003) The convergence hypothesis after 10 years. Research Paper 2003-06, University of Wisconsin, Madison
Durlauf S, Johnson P (2005) Multiple regimes and cross-country growth behavior. J Appl Econom 10:363–384
Ellison G (1983) Learning, local interaction and coordination. Econometrica 61:1047–1071
Eshel I, Samuelson L, Shaked A (1998) Altruists, egoists, and hooligans in a local interaction model. Am Econ Rev 88:157–179
Fagerberg J (1994) Technology differences in growth rates. J Econ Lit 32:1147–1175
Galambos J (1987) The asymptotic theory of extreme order statistics, 2nd edn. Krieger, Melbourne
Galor O (1996) Convergence? Inferences from theoretical models. Econ J 106:1056–1069
Goyal S, van der Leij M, Moraga-Gonzá lez JL (2003)'Economics: an emerging small world? Working Paper, University of Essex and Erasmus University
Goyal S, Vega-Redondo F (2005) Network formation and social coordination. Games Econ Behav 50:178–207
Grossman J (2002) The evolution of the mathematical research collaboration graph. Congressus Numerantium 158:201–212
Halpin-Healy T, Zhang Y-C (1995) Kinetic Roughening, stochastic growth, directed polymers & all that. Phys Rep 254:215–415
Jackson MO, Watts A (2002) On the formation of interaction networks in social coordination games. Games Econ Behav 41:265–291
Jackson M (2005) A survey of network formation models: stability and efficiency. In: Demange G, Wooders M (eds) Group formation in economics: networks, clubs, and coalitions. Cambridge University Press, Cambridge .
Quah DT (1996) Twin peaks: growth and convergence in models of distribution dynamics. Econ J 106:1045–1055
Quah DT (1997) Empirics of growth and distribution: stratification, polarisation, and convergence clubs. J Econ Growth 2:27–59
Marsili M, Vega-Redondo F, Slanina F (2004) The rise and fall of a networked society: a formal model. Proc Nat Acad Sci USA 101:1439–1442
Mankiw NG, Rimer D, Weill D (1992) A contribution to the empirics of economic growth. Q J Econ 107:407–437
Newman MEJ (2001) The structure of scientific collaboration networks. Proc Nat Acad Sci USA 98:404–409
Newman MEJ (2003) The structure and function of complex networks. SIAM Rev 45:167–256
Vega-Redondo F (1997) The evolution of Walrasian behavior. Econometrica 65:375–384
Verspagen B (1991) A new empirical approach to catching up or falling behind. Struct Change Econ Dyn 2:359–380
Weisbuch G, Deffuant G, Amblard F, Nadal J-P (2002) Meet, discuss and segregate! Complexity 7:55–63