The Complex Dynamics of Innovation Diffusion and Social Structure: A Simulation Study

Michele Simoni
University of Naples "Parthenope"
Department of Business Studies
E-mail: mic.simoni@gmail.com

Adam Tatarynowicz
University of St. Gallen
Centre for Business Metrics
E-mail: a-tatarynowicz@kellogg.northwestern.edu

Gianluca Vagnani
University of Rome "La Sapienza"
Department of Management Science
E-mail: gianluca.vagnani@uniroma1.it

Economic and sociological theories increasingly highlight the importance of social networks for the diffusion of technological innovations. However, while it is generally agreed that network structure affects the speed and scale of technology diffusion by offering efficient communication channels for the transfer of information, it seems that current literature does not say whether also the process of diffusion can affect the very structure of the networks on which it occurs. We attempt to address this gap by examining how the spread of a technological innovation influences patterns of dyadic relationships within a large social network and, in turn, how those changes interact with the diffusion process. In this paper we develop a computer simulation applying agent based modeling techniques and use it to inquire about the feedback loop between the number of adopters and diffusion as well as about the extent to which local bandwagon forces reshape entire networks that generate them. Our initial results show, first, the emergence of an interesting phase transition in the speed of technology diffusion and, second, the acquisition of a complex property by the network structure when the stability of social ties, the average number of connections per agent, and certain agent internal characteristics are manipulated.

1. INTRODUCTION

We are witnesses of a growing consensus in the scientific community that social networks play an important role in the diffusion of major technological innovations. The common theme in this diverse body of economic, management, and sociological literature (e.g. Bala & Goyal, 1998; Barley, 1990; Katz & Shapiro, 1992; Rogers, 1995) is that network structure affects both the speed and scale of diffusions by providing social actors with a powerful communication channel via which information about new innovations can spread widely and fast prompting adoptions. Still, even though the link from network structure to diffusion dynamics seems well established theoretically and empirically, the overall picture remains incomplete. First, social networks are viewed in most accounts as predominantly stable while stability in real life is anything but typical in social settings. Second, it is evident that current literature implies unidirectionality seeing diffusion as guided by the underlying network structure and not the other way around. Thus, questions whether and how the structure of the social network becomes affected over time by the diffusing technology remain largely unanswered.

It is our interest in this paper to address this gap. We examine how the spread of an innovation is influenced by actor relationships within a large social network and, in turn, how the diffusion process shapes those very relationships. To achieve this goal, we develop an agent based model through which we control the microfoundations and track joint outcomes of the two phenomena. First, our model draws on the growing body of knowledge about complex adaptive systems in general in which local interaction coupled with changing agent behavior is known to produce different outcomes as a result of small modifications in the initial conditions. Second, we acknowledge that such an approach has recently gained prominence in the management literature as it gives the researcher an opportunity to bypass the requirement of strict mathematical formalization in place of a more contextual treatment of the underlying process (Lane & Maxfield, 1997). It is therefore particularly well suited for the study of complex social processes that succumb to no other but non-linear analyses.

The agents we model are both heterogeneous and adaptive. They assess individually their expected returns from the new technology based on their own internal beliefs about it as well as a general level of ambiguity with which the technology enters the market. They also bring in varying attitudes towards learning. Some agents may be more interested in exploiting the technology they already use, others in exploring new possibilities. The trade-off between these characteristics gives every agent a unique behavioral trait that governs its networking activities but it is not fixed throughout the model: by a strict rule certain individuals are allowed to adapt once they adopt the innovation by changing their learning focus while others never adapt. Controlling statistically which agent type will prevail we give the agent community a global learning perspective equivalent with that of real networks without sacrificing abstraction, parsimony and ease of implementation.

That community, whose size is kept constant, exists on an undirected graph of low average degree in which every agent is connected only to a small number of other agents. While the immediate environment of such an agent is its local ego-network, the sum of all local neighborhoods generates a large social network that can act as a powerful tool to channel information about the diffusing technology. Thus, by being embedded within a macro scale network, agents are exposed to micro scale pressures from their own neighborhoods to adopt. Structurally, network evolution is governed by a random search process in which the probability of a new tie being created and an existing tie dissolved depends both on individual agent characteristics and on parallel diffusion dynamics.

2. COMBINED MODEL OF TECHNOLOGY DIFFUSION AND NETWORK CHANGE

2.1. Threshold Model of Innovation Diffusion

We use a threshold model because traditional rate-oriented diffusion models are primarily designed to answer the question *how fast*, rather than *when*, diffusion occurs. In that sense, they do little to explain the emergence and extent of bandwagons. We propose a mathematical model of diffusion which is very much in line with the original model advanced by Abrahamson & Rosenkopf (1993; 1997). Models like that assume that potential adopters may have varying attitudes towards the new technology and that, therefore, the timing of adoptions is likely to be different across different adopters. At any moment in time, the comparison between costs and benefits of the new technology is affected by the number of individuals who have already given in to that technology. Still, while pressures stemming from earlier adopters may at each diffusion cycle cause a wave of subsequent adoptions, an individual does not have to jump on the bandwagon automatically. Bandwagons add significantly to the individual's perception of the new technology but are not its only determinant. Only after the bandwagon pressure exceeds a certain, individually derived threshold will an adoption occur. Thus, an equally important factor is the potential adopter's subjective idea about the innovation resulting from probabilistic expectations of whether using the new technology will generate a profit or a loss: the greater the expected loss the higher the threshold. Formally, we express the adoption decision of a single agent as:

$$D_i = I_i + A \cdot P_i \tag{2.1}$$

where D_i stands for agent *i*'s decision to accept the innovation meaning "adopt" if $D_i > 0$ and "reject" if $D_i \leq 0$. This decision is the sum of *i*'s subjective assessment of the new technology I_i drawn from a normal distribution with a negative mean and the product $A \cdot P_i$ which stands for the actual bandwagon pressure experienced by agent *i*. *A* is a measure of ambiguity of the new technology and indicates the extent to which an individual values information represented by P_i defined as the proportion of adopters in *i*'s ego-network:

$$P_i = \frac{\sum_{j=1}^{d} a_j}{d} \tag{2.2}$$

where d is the focal agent's degree and a_j equals 1 if the neighboring agent *j* is an adopter or 0 otherwise.

P_i used here is a modified version of the original construct proposed in Abrahamson & Rosenkopf's second paper (1997). While in their seminal work the number of adopters in the focal agent's ego-network is divided by the number of all potential adopters in the network (i.e. by total network size), here the denominator is the *degree* of an agent. That produces a more intuitive value of P_i calculated with respect to the agent's local environment and not the entire social network structure of which that agent may not be aware.

2.2. Evolution of the Social Network

Two processes make up topological change in any social network: the dissolution of existing links and the formation of new ties. To model these dynamics, we advance a procedure we call "guided random rewiring". Our technique has the feature common to all known rewiring procedures in that it keeps the total number of connections and actors constant and thus significantly reduces the analytic effort to trace changes in network structure over time. But it is also different from the more stylized stochastic rewiring techniques known from statistical physics (Watts, 1999; Watts & Strogatz, 1998) in that the algorithm takes account of node attributes when deciding which ties are to be disrupted and which to be formed.

At the outset, the agents we consider are endowed with a random structure. Random networks have been well studied analytically but are topologically quite unlike most empirically observed social systems. Therefore, they constitute a generic class of artificial graphs that can be used as a good baseline for deploying processes likely to result in some quantifiable structural change. A random network is typically a graph in which either (a) a given probability is associated with the formation of an edge between any pair of vertices, or (b) a given number of edges are allocated among randomly chosen pairs of vertices (Cowan & Jonard, 2004; Cowan *et al.*, 2002). The initial network modeled here belongs to the first category.

Having assumed a random network of size *N*, in each step our model selects with uniform probability a sub-set of agents of size less than *N*. For each of the selected agents a set of probabilities is derived with which the links connecting those agents with their neighbors can be removed. For a focal agent, this results in a ranking of existing ties from the most likely to the least likely to dissolve. With that probability, then, the selected agent gives up one tie. Once tie dissolution is completed, a different agent set of the same size is chosen for tie formation. Similarly, for each of the newly chosen agents a list of potential partners is compiled spanning over the entire community except that agent's ego-network. Then, every potential link is assigned a probability ranking and the selected agents each create a single new link. Ranking criteria used here are analogous to those responsible for tie removal and will be explained in due depth shortly. How big are the selected sub-sets is a key exogenous parameter in our model. In fact, we see it as a simple way to control network dynamics since the total number of nodes allowed to have their ego-networks reshaped in each simulation step translates directly into how much structural change the network will experience in total over time.

Based on what criteria do real people formulate their networking behavior? Much of the existing literature on learning prescribes that we are characterized by an internal trade-off between on one hand willing to explore new possibilities using our available contacts and, on the other hand, exploiting what we are already familiar with (Lee *et al.*, 2003; March, 1991; Nooteboom, 2000). Unless there is an extreme case of a fully exploitative or explorative individual, it is likely that in most of us both qualities appear in combination. The model we develop here has this duality built into every agent. Consider *i* as uniquely characterized by a propensity to explore $0 \le P_i(Exploration) \le 1$ and a propensity to exploit $P_i(Exploitation) = 1 - P_i(Exploration)$. Given this specification, an agent viewed as explorative will have $P_i(Exploration) > 0.5$ while an exploitative agent will come with $P_i(Exploitation) > 0.5$. The individual propensities are drawn at simulation begin from the statistical beta distribution. This distribution offers a simple mathematical way to cover the entire breadth of an agent's exploration preference by manipulating only two shape parameters α and β as in:

$$f(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1} \tag{2.3}$$

where $0 \le x \le 1$, α and β are positive integers, and $\Gamma(\alpha)$ is the Euler gamma function.

The idea is to view the trade-off between explorative and exploitative behavior as the key mechanism by which agents rank their subsequent networking decisions and the social structure evolves over time. We design what we believe is a highly realistic framework in which exploitative agents are driven by the principle of homophily when connecting with similar others with a high propensity to exploit while, at the same time, explorative non-adopters tend rather to link with adopters of the innovation. A non-adopter *i* whose $P_i(Exploitation) > 0.5$ thus forms with a high propensity a new tie with *j* who is using the same old technology. We calculate that propensity as $P_i(Exploitation) \cdot P_j(Exploitation)$ considering that *i* and *j*'s preferences to exploit are independent of each other. Similarly, if *i* is a non-adopter with $P_i(Exploration) > 0.5$ it will be more likely to establish new relations with adopters. That joint propensity to cooperate will be given by $P_i(Exploration) \cdot P_j(Exploration)$. Finally, the extent to which an adopter whose $P_i(Exploitation) > 0.5$ will be interested in pairing with other exploitative adopters equals $P_i(Exploitation) \cdot P_j(Exploitation)$. It is possible to regard these values as measures of mutual fit between two agents. The higher the fit the higher also the likelihood that the model will select them as neighbors.

Purposefully, we reserve for last the discussion of an explorative adopter. We define explorative non-adopters as adaptive and thus avoid the paradoxical situation of agents who are already in possession of the new technology but are still treating it as unknown. Adaptation occurs when an explorative agent adopts the technology and simultaneously adjusts to the new learning target by becoming exploitative. In this way, we guarantee that the following is true: while the relative proportion of explorative and exploitative agents among non-adopters is controlled at simulation begin by the beta distribution, post adoption all explorative adopters decide to learn about the innovation that they have just accepted to use. The only exception are early adopters who use the innovation as a result of subjective assessment rather than bandwagons.

Finally, the paradigm of mutual partner fit applies also to existing ties that face deletion. In this case, however, we are interested not in the joint propensity of two agents to remain connected $P_i \cdot P_j$ but in them becoming disconnected due to mutual misfit, given by $1 - P_i \cdot P_j$. Again, several scenarios are thinkable depending on the technology that both agents use and their idiosyncratic learning preferences. We summarize those in Table 2.1.

Table 2.1. Ranking criteria for tie deletion with respect to agent's learning preference and technology used.

Agent Type	**Non-Adopter**$_j$	**Adopter**$_j$
Non-Adopter$_i$	$1 - P_i(Exploitation) \cdot P_j(Exploitation)$	$1 - P_i(Exploration) \cdot P_j(Exploration)$
Adopter$_i$	$1 - P_i(Exploration) \cdot P_j(Exploration)$	$1 - P_i(Exploitation) \cdot P_j(Exploitation)$

One last issue remains to address, namely how the model guarantees that ties are really deleted and formed with the likelihood prescribed by their rankings. The easiest way is to consider the propensities calculated for agent pairs on the basis of their joint characteristics as weights fed into a stochastic procedure that subsequently picks ties either for removal or creation depending on where the model currently is. Guided by these weights, the procedure decides probabilistically which tie it will choose, with the most heavily weighted tie being also the most likely to get picked. To achieve this goal, we simply divide each resultant weight by the total sum of all weights in order to turn it into values normalized between 0 and 1. These values then indicate final probabilities with which each tie in the agent sub-set can be selected by the random procedure.

3. RESULTS

3.1. Network structure and technology diffusion

At this point, we are interested in the impact of changing network structure on technology diffusion. Hence, the two exogenous parameters we manipulate are structure-relevant: the mean degree of the network and network dynamics. Holding the size constant at 500 agents we formulate 35 different initial conditions by moving from a sparsely linked random graph with an average degree $k = 2$ to a dense graph with $k = 10$ (by steps of 2), and simultaneously varying network dynamics D between a low 2% and a high 11% (by steps of 1.5%). We also assume that the network consists to 50% of explorative agents with $P(Exploration) > 0.5$ while the other half show an exploitative attitude. To obtain this distribution we set both beta parameters to 5.0. Technological ambiguity A is set to a low 2 in all scenarios. Finally, into all of these networks we introduce a single early adopter being the only agent with a positive I_i. We replicate every scenario 100 times for a total of 3500 runs. In each run we first measure the cumulative share of adopters over time. Then, to compare the speed of technology diffusion across all scenarios, we look at how many steps it takes for 60% of agents to adopt. The simulation lasts until that cumulative share of adopters is reached or, otherwise, it ends after 10000 steps.

Figure 3.1a below depicts the typical progress of technology diffusion in a situation where structural change is slow with just 2% of individuals redefining their social relationships within the network in each step. The striking result is that connectivity seems to have a negative impact on adoptions: the higher the average degree the lower the speed of diffusion. This effect is particularly strong early on, roughly during the first 200 steps of the simulation. After that, technology adoption slows down in the sparse network and begins to show patterns similar to those in denser topologies. It is important to note that the critical point marking 60% of adopters is reached much faster by the sparse network. The case of high connectivity but varying dynamics (Figure 3.1b) reveals quite opposite patterns: here the higher is the network dynamics the more there are adoptions.

Figure 3.1a. Diffusion under low network dynamics.

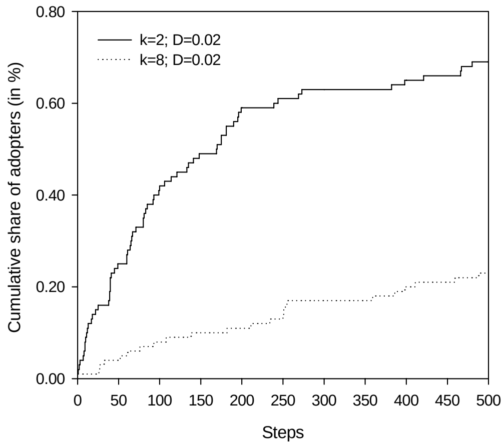

Figure 3.1b. Diffusion under high connectivity.

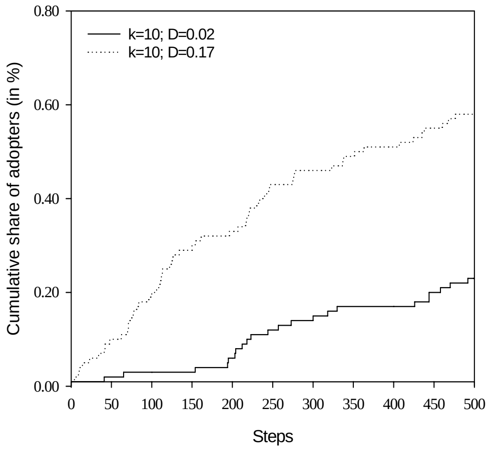

Is it because dense networks require more structural change to overcome inertia? Even more so, is some critical mass of change necessary to trigger diffusion the way it sets off here only after the first 200 steps? To enhance these findings formally, in Figure 3.2 we depict the combined effect of connectivity and dynamics on diffusion.

Figure 3.2. Combined effect of network structure on adoption (in steps necessary to reach a 60% adoption rate).

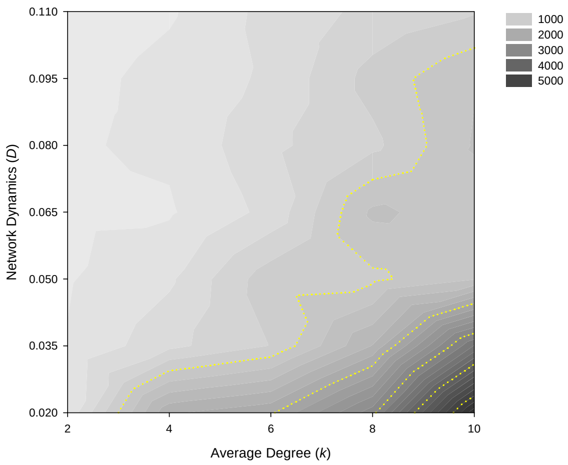

The contour plot shows that the time needed to reach the benchmark adoption rate increases dramatically as connectivity goes up but network dynamics remain low. At the opposite end lies the fast diffusion scenario representing a sparse network. Interestingly, the rate of topological change plays little role in fostering diffusion if connectivity is low. In fact, the progression from light grey to black in the chart above is so rapid that we can think of it as a phase transition between two distinctive network states: a fast and a slow diffusion state. The transition occurs somewhere between the average degree $6 < k < 8$, and dynamics $0.035 < D < 0.05$. Points lying on the "edge" of phase transition between the second and the third yellow contour line are those combinations of low dynamics and high connectivity that designate the frontier between two very different efficiency regions in this network.

We argue that the effect observed here stems largely from the agents' adoptive behavior. Since the agents experience social pressures to adopt from their ego-networks only, the higher the degree of an agent the lower also the level of bandwagon pressures placed on it by others. Thus, higher connectivity necessarily results in stronger lock-in effects of non-adopters in the old technology. On the other hand, high network dynamics bring more structural change into local neighborhoods. Under certain conditions, this change may absorb the lock-in created by the old technology as explorative agents will automatically tend to reconnect from users of the old technology to users of the innovation. Thus, the combination of low average degree and high dynamics increases the probability of an agent to adopt whereas high connectivity coupled with low dynamics dramatically lowers that probability regardless of whether the agent exploits or explores. We believe that this finding may be of particular importance for real world diffusion processes. It clearly shows that even in very innovative communities it is difficult to escape rigidity and conservatism over the long run as social ties grow in number and initially fluid structures stabilize.

3.2. Technology diffusion and network structure

We turn our attention to the emergent network structure now as a product of the interaction between the network and technology diffusion. At this point, we offer an initial perspective using graph visualizations only. Starting from a random graph of 500 agents we consider two networks that evolve under two different parameter configurations, one for a low $k = 2$ and a high $D = 0.11$ (Figures 3.3a and 3.3b) and the other one for a high $k = 10$ and low $D = 0.02$ (Figures 3.3c and 3.3d, respectively). Through multiple runs we track both networks until no more technology adoptions occur, and choose the most typical graph in each series of runs for visual presentation. To indicate the passage of time, the graph is shown at an early stage after the first 5% of diffusion cycles and a later stage after the completion of 20% of cycles. Blue color indicates adopters, red – non-adopters.

Figure 3.3a. Network evolved under low average degree and high dynamics after 5% of steps.

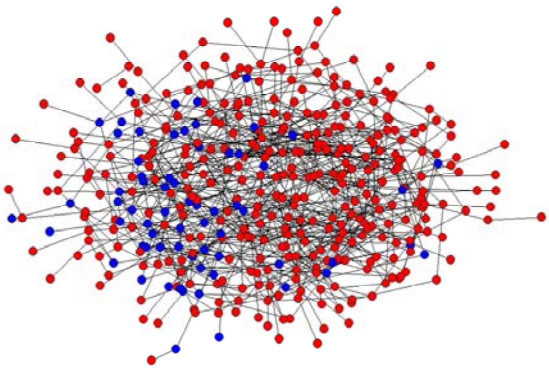

Figure 3.3b. Network evolved under low average degree and high dynamics after 20% of steps.

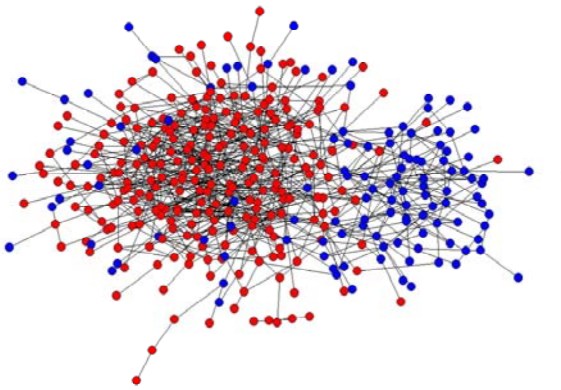

Figure 3.3c. Network evolved under high average degree and low dynamics after 5% of steps.

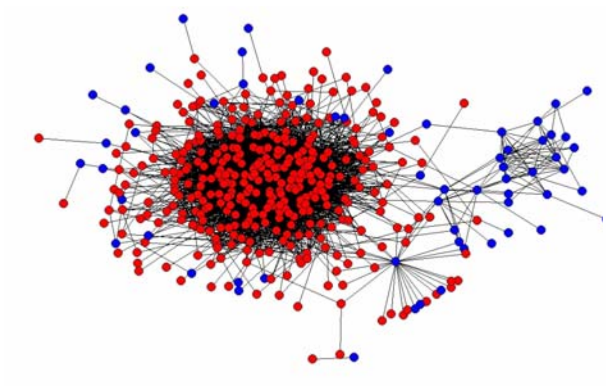

Figure 3.3d. Network evolved under high average degree and low dynamics after 20% of steps.

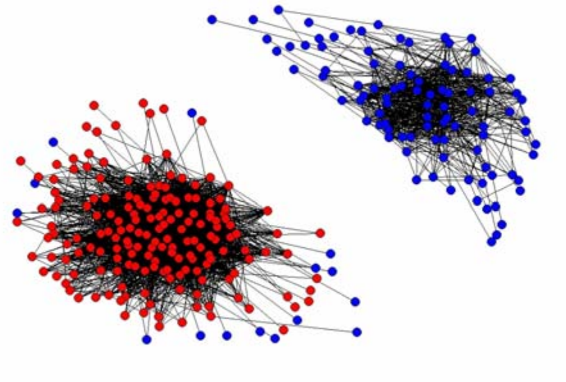

These images seem to contradict all our intuition. Should low connectivity and high dynamics (Figures 3.3a and 3.3b) not be more likely to cause a disruption in the social structure as opposed to a network that is sparsely interconnected and evolves fast (Figures 3.3c and 3.3d)? What we see here is really the opposite. It is in the well connected network of low dynamics that most adopters separate from the non-adopter cluster and create a dense component of their own. Our observations indicate that these two groups go apart early in the model and remain disjoint for the rest of the simulation. Clearly, the disconnected structure inhibits further bandwagon adoptions that would otherwise continue to have a strong effect on the topology of this network.

It seems to be an emergent property of our complex system that the structural dynamics observed here follow closely the logic of social tie strength without it being an explicit component of agent behavior. If network change is slow, dense ties connecting exploitative agents persist long enough to push out the formation of mixed explorative pairs if the average degree remains constant. Simultaneously, even though these homophilous links have a much higher probability of being selected for deletion as there are simply more of them, they also have a much higher likelihood of being formed again than the adopter–non-adopter links increasingly declining due to agent adaptation. Tie strength can thus be perceived as an emergent feature of the entire social community rather than of a single social tie. In our system, ties exist only in the context of the entire network and only in the context of that network at some point they become either strong or weak.

An even more important finding, though, is that the community which develops this internal property performs worse. Note that the duo of low dynamics and high connectivity ranks slowest on the speed of diffusion chart in Figure 3.2 whereas the dynamic sparse graph is also the best diffusing one. Long lasting exploitative ties increase agent lock-in in the given technology type and hamper further spread of the innovation. At the same time, explorative ties are no longer there to create bandwagons that would keep the network together. In this way, one of the communities shown above splits up in perfect accordance with the learning attributes and adoption status of its agents so that now no further diffusion is possible. On the contrary, a more egalitarian network where no strong within-group or weak between-group ties emerge manages to preserve a substantial number of explorative links from one cluster to another which then foster bandwagons and enable vast diffusion.

4. CONCLUSION

Our model is a dynamic set-up of high internal activity in which the progressive diffusion of a new technology and continuous evolution of network ties are mutually intertwined to an extent allowing for an observation of complex systems outcomes. Diffusion involves passing information about the innovation from one agent to another. If we assume that such information can be transmitted via dyadic interaction channels only, then the overall topology of the interaction framework becomes critical for the diffusion process. In this model, the topology of the network changes permanently creating at each subsequent step an entirely new social architecture which is not only structurally different from its previous state but which also performs differently.

While we ascertain that by this mechanism network evolution gives direction to diffusion, we find that also the opposite is true. Singular adoption events induce the dissolution of some ties and cause others to be preserved. Yet, by the time existing ego-networks are disrupted, adoptions generate bandwagon pressures that prompt even more adoptions and even more shifts in the underlying network topology. This complex interplay is the engine that drives our model yielding interesting observations. First, we see that there is a strongly manifested phase transition in social network efficiency and, second, we discover that that phase transition is related to a complex property acquired by some social networks and not acquired by others depending on their initial conditions.

Is it possible to specify formally the set of such conditions necessary for networks to remain solely within the region of high efficiency to the left of the phase transition and thus to make an attempt at engineering social networks? If so, what implications can we derive from this knowledge for social systems in general? Are most real networks that typically develop a mixture of strong and weak ties always doomed to fail because inefficiency becomes their signature when structural change and economic processes come together?

While with this model we can undoubtedly raise a whole spectrum of other explorative issues, we take a shot at addressing the first question in the full paper where we derive analytically the set of points on the edge of phase transition in network efficiency. At the same time, we leave other issues, in particular those pertaining to real social networks, for a future study in which we will empirically assess the validity of this agent based framework and make more generic statements about the nature of technology diffusion in multiple network scenarios.

References

Abrahamson, E. , and L. Rosenkopf. 1993. Institutional and competitive bandwagons: Using mathematical modeling as a tool to explore innovation diffusion. *Academy of Management Review* 18:487-517.

Abrahamson, E., and L. Rosenkopf. 1997. Social network effects on the extent of innovation diffusion: A computer simulation. *Organization Science* 8:289-309.

Bala, V., and S. Goyal. 1998. Learning from neighbours. *Review of Economic Studies* 65:595-621.

Barley, S. R. 1990. The alignment of technology and structure through roles and networks. *Administrative Science Quarterly* 35:61-103.

Cowan, R., and N. Jonard. 2004. Network structure and the diffusion of knowledge. *Journal of Economic Dynamics and Control* 28:1557-1575.

Cowan, R., N. Jonard, and J. B. Zimmermann. 2002. The Joint Dynamics of Networks and Knowledge, 155-174. In: *Heterogeneous Agents, Interactions and Economic Performance. Lecture Notes in Economics and Mathematical Systems*, edited by Robin Cowan and Nicolas Jonard. Berlin: Springer.

Katz, M., and C. Shapiro. 1992. Product introduction with network externalities. *Journal of Industrial Economics* 40:55-84.

Lane, D. A., and R. R. Maxfield. 1997. Foresight, complexity, and strategy, 169-198. In: *The Economy as an evolving complex system II*, edited by W. B. Arthur, S. N. Durlauf, and D. A. Lane. Boulder, CO: Westview Press.

Lee, J., J. Lee, and H. Lee. 2003. Exploration and exploitation in the presence of network externalities. *Management Science* 49:553-570.

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* 2:71-87.

Nooteboom, B. 2000. *Learning and Innovation in Organizations and Economics*. Oxford: Oxford University Press.

Rogers, E. M. 1995. *Diffusion of innovations*. New York: The Free Press.

Watts, D. J. 1999. Networks, dynamics, and the small-world phenomenon. *American Journal of Sociology* 105:493–528.

Watts, D. J., and S. H. Strogatz. 1998. Collective dynamics of small-world networks. *Nature* 393:440-442.