EUI Working Papers

Network Organizations

Fernando Vega-Redondo

Network Organizations

FERNANDO VEGA-REDONDO

Network organizations

Fernando Vega-Redondo[*][†]
European University Institute, Florence

January 21, 2008

Abstract

It is common to define a network organization as one that is fast and flexible in adapting to changes in the underlying environment. But besides the short-run advantages of adaptability, fast changes in the structure of the organization can also be detrimental in the longer run. This happens, in particular, because agents need to depend widely on that structure to channel appropriately (and thus speed up) search.

I discuss the trade-off between adaptability and structural stability in a changing environment where, if the structure of the organization adjusts, information on the exact nature of the change becomes known only with some lag. The main conclusion obtained is that, as environment becomes more volatile, the optimal mode of the organization sharply switches from being totally flexible to being completely rigid, i.e. no intermediate configurations are essentially ever optimal. This has stark implications on the dichothomy of stability versus change that has been highlighted by recent organization literature.

JEL Classification nos.: D20, D83, D85.

1   Introduction

A "network organization" is usually conceived as an organization that is quick and flexible in adapting to changes in its environment. But changes in the structure of the organization can also be detrimental in the medium run, since it is partly the knowledge of the organization's structure that mediates (and thus speeds up) search. Here I discuss the tension between these two considerations. That is, I study the trade-off between adaptability and structural stability in a (network) organization that confronts a changing environment.

1

The model proposed to study this trade-off is particularly simple and stylized. The organization consists of an underlying backbone structure (a one-dimensional lattice network) that remains fixed and a limited number of links that can be "rewired" over time (for simplicity, just one per agent). In the background, there is an environment that changes at a given pace, i.e. the volatility rate. More specifically, it is assumed that every node/agent has a target node it has to reach, whose identity independently changes at a rate p. The dilemma faced by a node whose target has been reassigned is the following: should I redraw my flexible link to the new target? If this is done, direct access to that node (possibly the target as well in the immediate future) is secured. But, on the other hand, under the assumption that freshly rewired links take some time to become widely available to the organization at large, such an adaptation also imposes a global cost . Namely, it removes from the immediate operational structure of the organization some links that can be particularly valuable for overall search.

So, in a nutshell, the problem we pose can be formulated as follows. What is the optimal speed at which the organization should adapt to the changing environment? To cast this question sharply, the adaptability of the organization is supposed embodied by a single parameter q, the probability with which a node will redraw its flexible link to a new target. In this setup, the answer provided by the paper is a drastic one: depending on whether the value of p (volatility) is high or low, the optimal q (adaptability) should essentially be, respectively, either zero or one. Thus, in this sense, one finds that an optimal organization is either totally rigid or totally flexible, rarely something in between. As we shall explain, this has a sharp bearing on the dichotomy stability versus change that has recently motivated much of the organization literature.

The rest of the paper is organized as follows. Next, Section 2 provides a brief discussion of related literature. Then, Section 3 presents the model. The analysis is undertaken in Section 4, while Section 5 concludes.

2 Related literature

There are three distinct branches of literature quite related to our present concerns. In decreasing affinity, these are: (a) the economic theory of organizations, (b) models of search in complex networks, (c) the transactive-memory theory of organizations. I briefly discuss each of them in turn.

The economic theory of organizations has produced a huge body of research whose focus has been both on incentive issues and/or the way in which organizations can effectively handle diverse and decentralized information. In the latter vein, the work of Radner (1993) was a seminal contribution that (abstracting from incentive considerations) first modelled explicitly the organization as a network of informal flows. Other subsequent researchers have followed his lead (see e.g. Bolton and Dewatripoint (1994), van Zandt (1999), and Garicano (2000)), all aiming at characterizing the optimal network structure that, under varying conditions and in different senses, best pools the information disseminated

throughout the organization.

In recent years, and partly motivated by the rise to prominence of internet, there has been a surge of interest on the problem of how to conduct effective search in large and complex social networks. Building upon the early experimental work of Milgram (1967) on "small worlds" and the subsequent theoretical developments of Watts and Strogatz (1998), a key issue has been that of searchability. More specifically, the question is how to find short paths joining the nodes of large networks which involve (as indeed happens in the real world) a significant random component. Kleinberg (2000) provided key insights on the problem, leading an approach that has been adapted by several authors – see e.g. Guimerà *et al.* (2002) or Dodds *et al.* (2003) – to address the problem of how to design the (fixed) network of large organizations.

Finally, I refer to the so-called transactive-memory theory of organizations. This theory originates in the work of Wegner (1986). He stressed the importance of the process by which, as new information arrives to an organization, it is first allocated to individuals, then registered in the "organizational directory," and later retrieved in the most efficient manner. This three-fold mechanism is what has been called the organization's transactive memory system. A large body of theoretical and empirical literature has followed suit (see e.g. Moreland and Argote (2003) for a survey). In much of it, researchers have emphasized the importance of informal (and thus flexible) links in the successful implementation of an organization's transactive memory.

A good case in point is provided by the empirical work of Hansen (2003).[1] He studied 120 product development projects in a large electronics company, where each project was separately undertaken by one of the 41 business units of the firm. Hensen started by constructing a knowledge network, on the basis of the informal contacts identified among the members of the different units. Then, much in line with the key assumptions of our model, he found that the overall performance of each unit (specifically, the fraction of projects completed and the speed of completion) was highly dependent on the existence of short network paths to other units possessing relevant knowledge. Indirect connections, in other words, were crucial for good results, but their value was found to decay significantly with distance.[2] This, indeed, is consistent with the central measure of performance contemplated in our model, which in turn underlies the problem of network design addressed by our theoretical analysis.

3 The model

Our theoretical approach is purposefully simple and stylized, in order to capture the main ideas in the starkest fashion. Consider a large set of nodes $N = \{1, 2, ..., n\}$ arranged along an organizational backbone, which is here as-

---

[1] See also Hansen (1999) and Schulz (2003).

[2] These considerations were most decisive for the transfer of knowledge that could be largely codified. Instead, for hardly codifiable knowledge, direct and close contact between the source and the target played a primary role.

sumed to be a one-dimensional lattice without boundaries, i.e. a ring. Each node i is connected to both of its direct neighbors in the backbone, $i - 1$ and $i + 1$ (where the index here is interpreted as "modulo n"). These links are interpreted as formal and rigid ones, possibly reflecting the *formal chart* of the organization. In addition to such formal links, every i also has one additional link to some $\alpha(i)$, which will generally be "long-range" (i.e. far away from i on the underlying backbone). These long-range links are flexible, possibly informal, and can be adjusted over time, as determined by the plasticity/adaptability of the organization (see below for details). For the moment, let start by assuming that each $\alpha(i)$ is randomly selected from $N\backslash\{i\}$ with uniform probability. The resulting (undirected) network – consisting of the backbone plus the long-range links – will be denoted by Γ.

Let us further postulate that each node $i \in N$ has a target $\tau(i) \in N\backslash\{i\}$, whom i has to reach in order to address a specific demand or tackle a particular problem. Again, let us start by assuming that $\tau(i)$ is randomly selected from $N\backslash\{i\}$ with uniform probability. Then, given the prevailing network Γ and an array of targets $\boldsymbol{\tau} = [\tau(i)]_{i=1}^{n}$ the aim of the organization is to minimize the average path length (along the network) between every node and its target, so that its objective function can be identified with

$$\rho = - \left\langle d(i, \tau(i)) \mid \Gamma \right\rangle_{i \in N}.$$

But, as advanced, the focus of the paper is on the tension between adaptability and structure in a dynamic context where the environment changes over time. So let us introduce time into the model, indexing it by $t = 0, 1, 2, ...$ and dating the prevailing states $\omega_t \equiv [\boldsymbol{\tau}_t(i), \Gamma_t]$ accordingly. Suppose that the initial state $\omega_0 \equiv [\boldsymbol{\tau}_0(i), \Gamma_0]$ is constructed randomly, as above. Then, as time proceeds, the law of motion that governs the change in ω_t is decomposed into the following components (assumed, for simplicity, to be implemented sequentially):

1. Independently for each node i, its previous target $\tau_{t-1}(i)$ is redrawn afresh with probability $p \in [0, 1]$. (Thus, with probability $(1 - p)$, we have $\tau_t(i) = \tau_{t-1}(i)$.) When a new target for i is redrawn, each $j \in N\backslash\{i\}$ is selected with uniform probability.

2. Independently for each node i, its previous long-range link to agent $\alpha_{t-1}(i)$ is rewired with probability q to the current target $\tau_t(i)$. (Thus, with probability q, the long-range link of i connects to $\tau_{t-1}(i)$.)

3. For each i, the shortest average distance is computed, but using only the links in $\Gamma_{t-1} \cap \Gamma_t$. That is, only the informal links that have remained in place for at least one period form part of the operational structure of the organization. (At the beginning of the process, we posit that $\Gamma_0 = \Gamma_{-1}$, so that all initial links form part of the organizational structure.) Formally, therefore, we simply measure organizational performance at each t through

$$\rho_t = - \left\langle d(i, \tau_t(i)) \mid \Gamma_{t-1} \cap \Gamma_t \right\rangle_{i \in N}. \tag{1}$$

The above three items have the following interpretation. The *first* one embodies the idea of *volatility*: over time, the environment evolves and the needs/tasks/objectives of individual nodes are affected by it. The parameter $p \in [0, 1]$ modulates the rate at which the environment changes, thus leading to pressure for some adjustment to take place.

The *second* item specifies how and when, in response to such an adjustment pressure, actual changes in the network structure indeed occur. The parameter $q \in [0, 1]$ is a measure of organizational *plasticity*. It can be conceived as an attribute of organizations – say, a part their "culture" – and will generally differ across them. Sometimes, it may also be regarded as an outcome of design, at least partially, as when the declared aim of a certain policy is to reshape the culture of the organization in order to improve its performance.[3]

Finally, concerning performance, the *third* item proposes a precise way of measuring for it that is identified with the average network distance between every node and its corresponding target. This is motivated by the idea that the network distance separating an agent from a valuable partner (e.g. one that helps undertake current tasks) should have an important bearing on the speed and success of job completion. As explained above – recall the Introduction – this idea is not only intuitive but also enjoys some significant empirical support.

4 Analysis

Our main objective is to shed light on the interplay between the *plasticity* of the organization (as given by q) and the *volatility* of its environment (as captured by p). To fix ideas, a useful way to grasp this relationship is to address an optimal-design problem in which p is the exogenous parameter and q is the decision/design variable. Naturally, this problem must be formulated in a long intertemporal framework, where volatility and adjustment have a full chance to unfold. Thus, let us take a truly long-run perspective and identify the overall performance of the organization with

$$\rho \equiv \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \rho_t.$$

Since the underlying stochastic process is ergodic, ρ is independent of initial conditions and can be conceived as a function of p and q alone. Let us write $\rho(p, q)$ to reflect such dependence. Then, our main theoretical concern is well captured by an optimization problem formulated as follows: given any $p \in [0, 1]$, find $q^*(p)$ such that

$$q^*(p) \in \arg\max_{q \in [0,1]} \rho(p, q). \tag{2}$$

---

[3] See, for example, the influential work of Schein (2002, 2004) who conceives culture as the background for change in any organization. In fact, somewhat in line with the role of q in our model, he suggests that the "shared assumptions and beliefs about the stability of human relationships" is a key cultural dimension that differentiates organizations.

In essence, this optimization problem reflects a trade-off between two opposing objectives:

- *adaptability* – swiftness in responding to a mismatch between links and targets;

- *structure* – preservation of the network connectivity required to conduct search effectively.

To gain an analytical understanding of the essential implications resulting from this trade-off, we study the problem through an idealization of our framework in which the dynamics of the system is identified with its expected motion. As customary, we call such an idealization the Mean-Field Model (MFM). Given the stochastic independence displayed by the forces impinging on each node (both concerning volatility and link rewiring), it is natural to conjecture that the MFM should capture the essential behavior displayed by large finite systems. Indeed, this will be confirmed by numerical simulations, whose performance are found to match with great accuracy the theoretical predictions.

The MFM is defined by a dynamical system formulated on a population-wide (anonymous) description of the evolving situation. Then, under the implicit assumption that stochastic correlations can be ignored, a sufficient specification (or state) of the system is given by the fraction of nodes that are currently on target – i.e. those connected to their respective target through their long-range links. Let $\mu(t)$ stand for the fraction of such nodes prevailing at some t. Then, its law of motion is given by the following simple difference equation:

$$\mu(t+1) = (1-p)\mu(t) + q\left[\mu(t)p + (1-\mu(t))\right].$$

Thus the system globally converges to a unique positive fraction of nodes on target given by

$$\mu^* = \frac{q}{p + q(1-p)} < 1.$$

This implies that, in the long run, the total number of long-range links that are fully operational (i.e. have been in place for at least one period) is given by:

$$
\begin{aligned}
\lambda^* &\equiv \left[(1-p)\mu^* + (1-q)(1-(1-p)\mu^*)\right]n \\
&= \frac{p(1-q) + q(1-p)}{p + q(1-p)}n.
\end{aligned}
$$

Our next step is to compute the average distance between a node and its target, when a direct link does not exist between them. To this end, notice that the "operational network" prevailing at t (i.e. $\Gamma_{t-1} \cap \Gamma_t$) can be conceived as a small-world network of the sort studied by Newman *et al.* (2000), itself a variation of the original setup proposed by Watts and Strogatz (1998). Specifically, the number of shortcuts in the small-world network can be identified with the number λ of operational long-range links in our setup. This then allows one to rely directly on the expression derived by Newman *et al.* (2000) to approximate

the average network distance in their small-world setup. In terms of our present notation, they found it to be proportional to a function $F(\lambda)$ given by

$$F(\lambda) = \frac{2}{\sqrt{\lambda^2 + 2\lambda}} \tanh^{-1}\left(\frac{\lambda}{\sqrt{\lambda^2 + 2\lambda}}\right).$$

The function above, of course, only applies to the nodes that are not on target. Since the fraction of these in the long run is $[1-(1-p)\mu^*]$, the objective function ψ to be maximized is

$$\psi(p,q) = -[1 - (1-p)\mu^*(p,q)]F(\lambda^*(p,q))$$

where the notation $\mu^*(p,q)$ and $\lambda^*(p,q)$ highlights the fact that both parameters of the model, p and q, affect the long-run values. Thus, in sum, the optimization problem that concerns us here can be formulated as follows: given any $p \in [0,1]$, find a solution $q^*(p)$ of the optimization problem

$$\max_{q \in [0,1]} -[1 - (1-p)\mu^*(p,q)]F(\lambda^*(p,q)). \tag{3}$$

Once the full dependence on p and q is made explicit, the function $\psi(p,q)$ is found to display a complicated form:

$$\psi(p,q) = -\frac{4\left(1 - \frac{(1-p)q}{p+q(1-p)}\right)}{\sqrt{8n\left[\frac{q(1-p)}{p+q(1-p)} + (1-q)\left(1 - \frac{q(1-p)}{p+q(1-p)}\right)\right] + 4n^2\left[\frac{q(1-p)}{p+q(1-p)} + (1-q)\left(1 - \frac{q(1-p)}{p+q(1-p)}\right)\right]^2}} \times$$

$$\tanh^{-1}\left\{\frac{2n\left[\frac{q(1-p)}{p+q(1-p)} + (1-q)\left(1 - \frac{q(1-p)}{p+q(1-p)}\right)\right]}{\sqrt{8n\left[\frac{q(1-p)}{p+q(1-p)} + (1-q)\left(1 - \frac{q(1-p)}{p+q(1-p)}\right)\right] + 4n^2\left[\frac{q(1-p)}{p+q(1-p)} + (1-q)\left(1 - \frac{q(1-p)}{p+q(1-p)}\right)\right]^2}}\right\}$$

This leads us to approaching the optimization problem (3) through numerical methods, which can be used to deliver a specific solution to any desired level of accuracy (as a function of the volatility rate p). Figures 1-3 describe the solution $q^*(p)$ for different values of population size and show that it is qualitatively the same across a wide range in orders of magnitude.

The analytical solutions depicted in Figures 1-3 closely match the behavior observed in numerical simulations of the model, even if the population is relatively small. By way of illustration, Figure 4 shows the corresponding simulation results for $n = 100$. Specifically, one finds that the only two values of p for which the highest $\bar{\rho}$ was achieved at an interior q (i.e. one satisfying $0 < q < 1$) were $p = 0.55,\ 0.575$. Below these values, the optimal q is 0, and above them it is 1. These results are closely in line with the theoretical analysis – in particular, that described in Figure 1, which also applies to a context with $n = 100$.

The results reported above provide a stark picture of the way in which the tension between structure and adaptability is resolved in a network organization
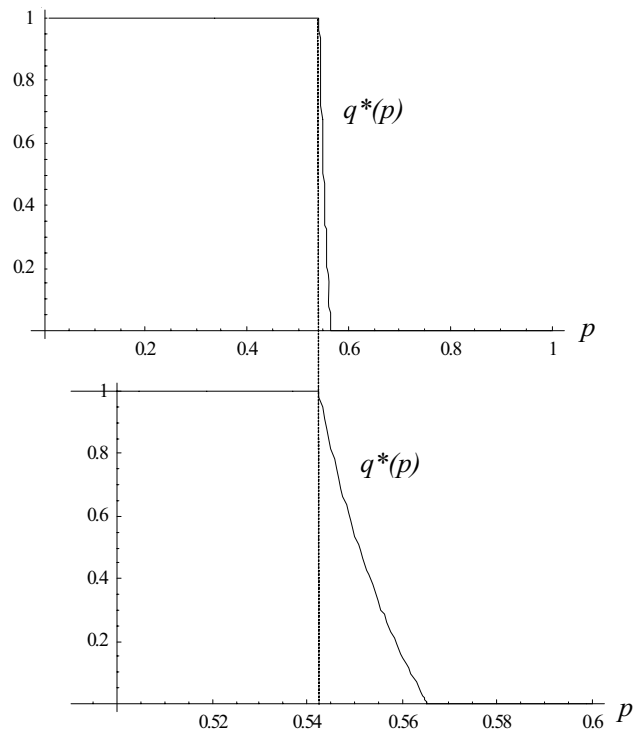
Figure 1: Optimal plasticity $q^*(p)$ as a function of volatility p for a population size $n = 10^2$. The function is shown both for the whole domain $p \in [0, 1]$ as well as for a scaled-up version that is "zoomed in" on the region where the transition from high to low optimal values takes place.
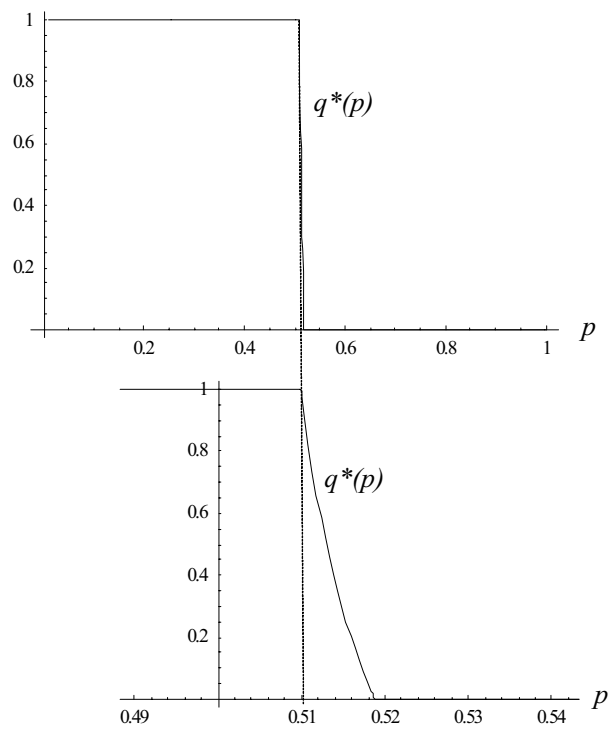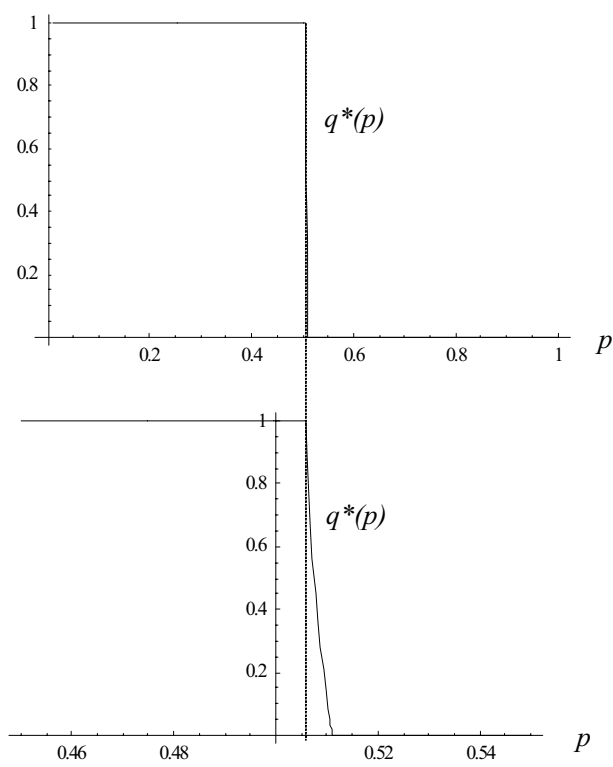
Figure 2: As in Figure 1, for $n = 10^6$.

9

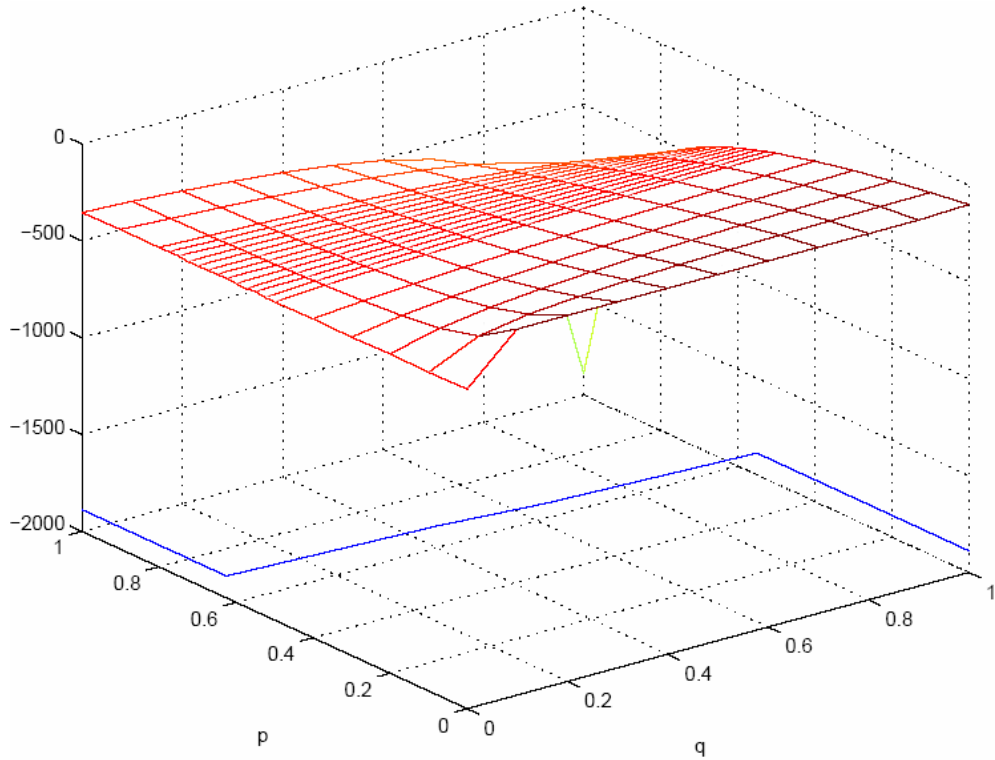Figure 3: As in Figure 1, for $n = 10^{10}$.

Figure 4: The upper surface depicts the average performance $\bar{\rho} \equiv \frac{1}{T} \sum_{t=1}^{T} \rho_t$ over $T = 20000$ rounds in a context consisting of $n = 100$ agents where each ρ_t is computed as in (1). The lower line on the p-q plane represents the plasticity q for which $\bar{\rho}$ is maximized at each volatility rate p considered. Specifically, p was changed along a grid with a step value $\Delta = 0.025$. The only two values of p for which the highest $\bar{\rho}$ was achieved for an interior q (i.e. one satisfying $0 < q < 1$) were $p = 0.55$, 0.575. Below these values, the optimal q was 0, and above them it was 1.

that is suitably described by our model. For low levels of volatility, the rate of change in the organization matches that of the environment since the plasticity of the organization is maximal. Thus, as the environment gets more volatile, the organization undergoes a fully parallel increase in link adjustment. This state of affairs, however, ends abruptly at levels of volatility well below complete target turnover. At a value of p sizably below 1, the optimal plasticity of the organization falls steeply to zero. There is, therefore, a wide range for p in which the best performance is achieved by freezing the network of the organization at its original random configuration.

These conclusions shed light on points made, in diverse forms, by the recent organization literature. For example, Schein (2002, 2004) argues that stability and change are "two sides of the same coin," and that both are part of any successful adaptation to an environment in perpetual flux. Moreland and Argote (2003), on the other hand, elaborate on this idea by emphasizing that too much flexibility may deteriorate the so-called "intellectual capital" of the organization (i.e. the knowledge available to an organization through its workers). This capital is accessed by the organization's transactive-memory system – recall Section 2 – whose operation is crucially facilitated by "a shared awareness among workers of who knows what (...)".

Our analysis may be regarded as having descriptive and normative implications. On the descriptive side, a prediction of the model is that, to the extent that organizations can be taken to operate efficiently, the most rigid ones should be those operating in the most volatile environments. This, however, raises normative issues as well, bearing on the likely conflict between individual incentives to adjust and the effects of such an adjustment on the overall performance of the organization. Even though the model does not take individual payoffs into account, an extended approach that would do so should naturally posit that they are inversely related to individual node-target distances. Then, it is clear that, if one abstracts from adjustment costs, maximum plasticity would be optimal from a purely individual perspective. But, as we have understood, this may be suboptimal for the organization as a whole if the volatility of the environment is relatively high. In essence, the problem at stake is a classical one of externalities – in this case, externalities on the search effort by others. And, as usual, what the problem may require is simply an appropriate intervention that, by impinging on individuals ability or/and payoffs to adjust, lead to a socially optimal outcome. To formulate and analyze this "implementation problem" in any detail is outside the scope of the present paper,

5   Summary and future research

The paper studies a model of a network organization that lives in a volatile environment and must therefore face the trade-off between the adaptability to changing circumstances and the preservation of an operational network structure. The theoretical framework is very stylized, which allowed us to obtain rather clear-cut conclusions. Specifically, we found that the positive effects of

adaptability completely dominate for low levels of volatility but are sharply offset when higher levels are achieved (well below, however, the situation of maximum turnover). This raises positive and normative issues on the "dynamic design" of organizations, which are left for subsequent research.

Additional issues to be explored in the future concern the sensitivity of the conclusions to some of the simplifying features of the approach. By way of illustration, I close the paper with the outline of some variations of the model that seem worth considering.

1. The model assumes that the rewiring of a long range link occurs with the same probability, independently of the distance closed by the adjustment. In the spirit of our approach, one could propose instead that, for example, rewiring occurs (again with some probability q, a design parameter) only if that distance exceeds a certain threshold. Naturally, this could just improve the overall performance of the organization. But a trade-off akin to that of the original model would still arise and, correspondingly, similar qualitative conclusions may be expected.

2. The framework posits no limit on the number of links an agent can have (either supported by herself or others). This, of course, is unrealistic for any reasonable interpretation of what links entail. Indeed, if no such constraints were binding, an organization that is structured as a star (i.e. all agents connected to a central agent) would achieve the best performance in many cases. The implicit assumption of the model is that such a star configuration is not feasible and, therefore, the organization is generally better off displaying some adaptability as the environment changes. A model where the degree of nodes is subject to some relatively tight upper bound would make that implicit feature of the model fully explicit.(For example, it could be assumed that any contemplated rewiring cannot be completed if it violates that capacity constraint.) It seems reasonable to conjecture that no essential insight would be afected by this variation.

3. The model has assumed that the backbone of the organization is a regular boundariless lattice. Often, however, the formal and stable network of an organization is best conceived as displaying a less symmetric form. A natural (but still stylized) alternative is given by a hierarchical tree structure, where each individual – except for the single apex – is connected to one "supervisor." Such a hierarchy is descriptive of many of the real-world structures observed in organizations, and probably this is partly due to the advantages it allows in the routing and processing of information (cf. Radner (2003)). Recently, however, it has been argued (see e.g. Dodds *et al.* (2003)) that the addition of long-range links connecting distant parts of an underlying hierarchic structure can greatly improve its overall performance. Indeed, this is supported by a large body of empirical research which finds that "(...) much of the real work in any company gets done through an informal organization, with complex networks of relationships that cross functions and divisions." (Cf. Krackhardt and

Hanson (1993).) The models that have been proposed in the theoretical literature to understand this phenomenon, however, have been mostly static. They conceive the organization network as fixed, even if it consists of a complex blend of hierarchic and transversal links. To enrich that approach with a genuinely dynamic model of the organization seems a welcome development that could be carried out along the lines suggested in this paper.

References

[1] Bolton, P. and M. Dewatripont (1994): "The Firm as a Communication Network," *Quarterly Journal of Economics* **109**, 809-839.

[2] Dodds, P. S., D. J. Watts, and C. F. Sabel (2003): "Information exchange and the robustness of organizational networks," *Proceedings of the National Academy of Sciences* **100**, 12516-12521.

[3] Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy* **108**, 874-904.

[4] Guimerà, R., A. Díaz-Guilera, F. Vega-Redondo, A. Cabrales, and A. Arenas (2002): "Optimal network topologies for local search with congestion," *Physical Review Letters* **89**, 248701.

[5] Hansen, M.T. (1999): "The search-transfer problem: the role of weak ties in sharing knowledge across organization subunits." *Administrative Science Quarterly* **44**, 82-111.

[6] Hansen, M.T. (2002): "Knowledge networks: explaining effective knowledge sharing in multiunit companies." *Organization Science* **13**, 232-48.

[7] Kleinberg, J. (2000): "Navigation in a Small World," *Nature* **406**, 845.

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

[9] Milgram, S. (1967): "The small-world problem," *Psychology Today* **2**, 60-67.

[10] Moreland, R. L. and L. Argote (2003): "Transactive memory in dynamic organizations," in R.S. Peterson and E.A. Mannix (Eds.), *Leading and Managing People in the Dynamic Organization*, New Jersey: Lawrence Erlbaum Associates, Publishers.

[11] Newman, M. E. J., C. Moore, and D. J. Watts (2000): "Mean-field solution of the small-world network model," *Physical Review Letters* **84**, 3201-04.

[12] Radner, R. (1993): "The Organization of Decentralized Information Processing," *Econometrica* **61**, 1109-1146.

[13] Schein, E.H. (2002): "Models and tools for stability and change in human systems," *Reflections* 5, 34-45.

[14] Schein, E.H. (2004): *Organizational Culture and Leadership*, Third Edition. New York: Wiley Publishers.

[15] Schulz M. (2003): "Pathways of relevance: exploring inflows of knowledge into subunits of multinational corporations," *Organization Science* **14**, 440-59.

[16] van Zandt T. (1999): "Real-Time Decentralized Information Processing as a Model of Organizations with Boundedly Rational Agents," *Review of Economic Studies* **66**, 633-658.

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

[18] Wegner, D. M. (1986): "Transactive memory: A contemporary analysis of the group mind," in B. Mullen and G.R. Goethals (Eds.), *Theories of Group Behaviorocial Cognition,* 253-76, New York: Springer Verlag.

[19] Wegner, D. M. (1995): "A computer network model of human transactive memory," *Social Cognition* **14**, 319-39.