

Constraints and Freedom of Action: a fitness trade-off

Massimo Ricottilli

Department of Economics and
Centro Interdipartimentale 'Luigi Galvani', Università di Bologna
Piazza Scaravilli 2, 40126 Bologna

August 27, 2006

Abstract

This paper attempts to discuss the role of constraints in network formation and in the exploration process of complex spaces of fitness possibilities. It is held that similarities appear in both social and biological systems. The paper further argues that constraints are instrumental in setting a fitness trade-off between specialised and across-the-board searching.

1 Introduction

It is a well established fact that social agents carry out what can be defined as an active search for improvement normally assisted by bounded rationality. It would seem, at first sight, that no relationship would hold between this social behaviour and the process of evolutionary change that applies to natural and, more specifically, to biological systems. Yet, this is not the case. Social systems can be viewed as sets of agents that interact within the bounds of locally defined contexts, or, using current terminology, neighbourhoods. Searching is an activity that requires skills and capabilities. The latter define, quite in general, the body of knowledge that makes searching possible. In turn, both the success and failure of this activity, achievement of desired goals and missing the expected target lead to learning. This is the fundamental process through which experience builds up to allow it to continue and eventually be temporarily successful. Furthermore, if it is recognised, as it should be on the grounds of a mounting body of evidence, that rationality is indeed applied but that it is bounded. This fact owes to the paucity of collected and available information, to the limits to the capacity to calculate, to the actual system of organising the stock of knowledge and, more generally, to experience and modes of learning that thwart and bend the basic tenets of pure rationality. If this is the case then this problem solving process amounts to a search in a space of possibilities that

obeys some specific topological properties. Social systems are in this sense complex adaptive systems. Its members must adapt to the given context, to the relationships tying them to other members, to the change that occurs to each of them and to the system as a whole. The logic behind this adaptive behaviour can best be rendered in terms of an 'if..then' procedure that scans the space of possibilities to change and improve. The latter can be formalised as a landscape comprising all the conceivable options to which some specification of fitness has been attached. Fitness is, of course, a biological concept that, however, translates well to a social and economic framework. To grasp this point, it suffices to employ, *in lieu* of fitness, the more *social* term of performance. When discussing technological progress, for instance, we can indeed speak of a more performing technology as one that is more productive or just more profitable for those that undertake to use it. It is, in this sense, fitter. The concept that designates all the conceivable options defining the cardinality of the space of possibilities is a more delicate one. When dealing with a process of change, it is clear that outcomes are uncertain and, indeed, not foreseeable, at least in the more radical cases. Thus, the landscape portraying novel and as yet untried system layouts is conjectural and its function serves a useful but merely heuristic function. The crux of the matter is that of investigating the direction and the implications of change. The latter comes about in consequence of a more or less intricate pattern of interaction that makes of a social system a complex network.

These characteristics of a social system dynamics are not far removed from those of natural, certainly biological ones. Mutation and recombination are the relevant processes of change in this context and can be analytically investigated as walks on fitness landscapes. Selection rather than rationally bounded judgement decides on new configurations but the similarity, if not brought to unlikely extremes, still holds. What rationality, albeit bounded, actually does is to produce searching patterns that occur on a 'fast' (human) time scale whilst biological ones walk on their landscapes in epochal time. Likewise, the physics of networks yields results that appear to apply to both scientific domains.

This paper attempts to review some basic characteristics of searching for improvement of systems nested and interacting in networks of both a social and biological kind, as they are analyzed in models that seem to easily straddle both without much loss of generality. More specifically, it is intended to focus on the role of constraints, that is to say of devices that appear to limit the extent and breadth of procedures leading to evolution and change.

2 Constraints to network evolution

It seems quite appropriate to begin this discussion models that deal with the very formation of a network and with its spacial characteristics (see Albert and Barabasi (2002), Boccaletti et alii (2006)). Connectivity is a very important network feature since its architecture may work to either enhance or lower the probability of finding performing solutions. In what follows networks will be modeled as graphs. Given a population N of nodes in a graph representing a

network F , each carrying a specific fitness, say a technology, a set of useful information, chemical stimuli and the such like and that is evenly if not randomly distributed over a landscape, an efficient connectivity enables to establish relationships of either directed or undirected exchanges with nodes that are highly performing with as low a number of costly links (in terms of energy, expenditure etc.) as possible. An entirely linked network is normally not efficient. A network is totally linked when, given N nodes, the number L of undirected links is $L = \frac{(N-1)N}{2}$ such that the average degree is $\hat{k} = \frac{N-1}{2}$; in the case of directed links, $\hat{k} = N - 1$. Having such a high number of links carries a great dispersion of *energy* : this is the case of nodes, or agents, possessing idiosyncratic properties, for instance technological knowledge that must be spread over a large number of heterogeneous partners; it is often a redundant architecture. Specialised relationships, confined to a few but useful nodes, normally carries better fitness. This is the case of clusters measured by a coefficient that counts all the links between nodes that are a node's neighbours, i.e. that are linked to it, in relation to those that would exist if they were all linked. Thus, if k_i is node i 's degree, then its neighbours would be a complete cluster if the number of links between them were $\frac{(k_i-1)k_i}{2}$. If the actual case counts only a given E_i then node i 's coefficient would be $c_i = \frac{2E_i}{(k_i-1)k_i}$ and the system's average $\hat{c} = \sum_{i=1}^N \frac{2E_i}{(k_i-1)k_i}$. The reason why clusters seem to possess efficiency rests with their specialisation. The case of economic clusters may be construed to shed some light on this issue. It has often been observed that many economies feature firms being highly concentrated in a given spatial context, in what has come to be known as an industrial district. In these districts, closely knit firms gather and exploit positive externalities, exchange technical information, foster the supply of a specialised work force and the services of consultancies, they easily reach out to know-how and skills that develop and thrive because of the district's existence, in other words on account of the links between the various firms. Yet, the nature of these firms and the pattern of relations that they hold is homogeneous in various respects: they produce similar output, they exploit similar technologies and are often of similar size. The implication is that they are constrained to operate within the bounds of the cluster they belong to, the information that flows through its links being, in general, quite idiosyncratic with the result that firms tend to ignore what actually happens in other, different, clusters. They, in this case, act as constraints in the sense that they enhance positive externalities but at the same time tend to isolate their members from other contexts. Clustered connectivity, therefore, is both a drawback and a focusing device. Watts and Strogatz (1998) have shown that randomly rewiring between clusters enhances a network performance if measured by its average path length. It is befitting to consider the latter as the main ingredient of a measure of efficiency since it is an index of how much an information must travel to reach out to the whole system. An useful index exploits this feature by employing its reciprocal in the following manner: $\delta = \frac{1}{N(N-1)} \sum_{i,j \in N} \frac{1}{d_{ij}}$, where d_{ij} is the path length between nodes i and j ; the system's average path length

being $\hat{l} = \sum d_{ij}$. These authors' result is quite telling since it shows that by a relatively few rewired links between nodes belonging to different clusters, while the average clustering coefficient, \hat{c} , remains rather high, the average path length sharply falls and, conversely, δ significantly increases. This result is obtained, as is widely known, by randomly subjecting nodes to rewiring with some probability: the latter need not be at all high, in any case quite far from $p = 1$ that would generate a random connectivity pattern. It follows that links straddling different clusters need not be many. The advantage of clustering is the smooth flowing of specialised information and of its positive externalities, yet it is absolutely clear that if left in isolation, no matter how easily and fluently the former went about and the latter became available among cluster members, fitness would be limited if not hampered by the paucity of outside sources of information, knowledge and innovative capabilities developed elsewhere within the system. On the other hand, excessive and redundant connectivity, as in the random case, would void the advantage of specialisation. It is indeed the intermediate case, the *Small World*, that preserves the advantages of the latter while eschewing the redundancy of the former that fosters greater efficiency and fitness.

What is crucial in this rewired layout is the role of some nodes as informational hubs redistributing information all over the network: what counts is their place between clusters of other nodes. A betweenness measure serves the purpose of stressing this strategic position: the betweenness of node i can in fact be rendered by $b_i = \sum_{j,k \in N, j \neq k} \frac{n_{jk}(i)}{n_{jk}}$ where $n_{jk}(i)$ is the number of shortest paths connecting j with k and going through i while n_{jk} is the number of shortest paths linking them. Nodes with a high b_i fulfill this important role. Trading off the binding nature of constraints specialising and focusing activity with the *freedom* of across-the-board connectivity resulting from random linking that enhances heterogeneity but indiscriminate informative content appears to be the key to efficiency and therefore to fitness.

Fine-tuning these two network properties leading to a network architecture is, therefore, of the utmost importance for both social and biological systems. The question that arises at this point concerns the way systems organised in networks evolve their connectivity. Recent work on these matters may be helpful to shed some light on node evolution (Intrator and Cooper (1992), Castellani et alii (1999), Andergassen et alii (2006)). For instance, recent work on biological network evolution has shown this process as one of adaptation grounded on *experience*. In these models, *experience* is meant to denote the record kept by a node by resorting to its specific means (chemical or logical), a neuron but the term could equally apply to a social agent, and encoded in a memory device. Action towards setting up an edge with other nodes is then assumed to be taken based on this record and a model can therefore be construed by specifying an adaptive, learning algorithm. There is a compelling reason that supports this procedure. Setting up a link is an energy consuming and costly undertaking that gradually lends strength to a potential edge until it finally becomes a functioning link but that may also meet with failure if not efficiently

carried out. Its outcome can accordingly be thought of as being subject to some loss if not well suited for the purpose: the strength cumulated thanks to past action can vanish thus frustrating the link building effort.

Let $\Phi(u_{ij}, \Theta_x)$ be a function that drives u_{ij} , an index of the strength building process for an ij edge. An edge can be considered as established when $u_{ij} = 1$ and definitely refused when $u_{ij} = 0$. In turn, Θ_x designates the *experience* record functioning as an adaptive benchmark for node i . The loss that a node is likely to incur after some time from the beginning of the process can be measured by :

$$L_i(\Theta_x) = \mu \int_0^{u_{ij}} \Phi(s, \Theta_x) ds$$

$s \in (0, u_{ij})$, μ being a rate of decay, from which a measure of risk straightforwardly follows as its expected value, i.e. $R_{u_{ij}} = E(L_i(\Theta_x))$, defined over the entire domain of u_{ij} . The role of function $\Phi(s, \Theta_x)$ must, therefore, enact an efficient process capable of minimizing this risk of loss. A suitable function has been shown to be:

$$\dot{u}_{ij} = u_{ij}(u_{ij} - \Theta_x)$$

a simple Volterra-Lotka-like differential equation. Θ_x , the record or memory, acts as a threshold and can be made explicit by $\Theta_x = \sum_{j \in N_i} u_{ij}^2$ or $\Theta_x = \sum_{j \in N_i} u_{ji}^2$, the first applying to outgoing links, the latter to incoming ones. Crucially different results are obtained according to how N_i is defined and the reader can easily check that stable solutions are obtained for either $u_{ij} = 1$ or $u_{ij} = 0$. The width of N_i is indeed fundamental: the model shows that a range of connectivity patterns emerge from a completely random graph when only immediate neighbours of each node are taken, for all i 's, to entirely specialised connections when the whole system of nodes is included within the threshold. In the latter case each i either receives just one input from a specific node, randomly providing output to the system, or *vice versa* is a specialised supplier randomly receiving from the network. Intermediate cases are very interesting. It can, indeed, be shown that for a combination of input providers and output receivers entering the threshold Θ_x the connectivity pattern exhibits a near power law distribution. It can be argued, but the actual proof will have to be left to empirical verification, that it is the latter case to be the most likely since it is the expression of an adaptive search that is sufficiently local, i.e. bounded by a well defined neighbourhood, without being restricted to very few nodes.

3 Interdependence between nodes

As it has been seen, connectivity affects fitness. It is, now, interesting to investigate how a specific form of connectivity establishing a link of interdependence between nodes influences fitness by means of limiting constraints. A family of

biological models that have found interesting applications in the social sciences lend themselves quite well to analyse this role (Auerswald et alii(2000), Kauffman (1989), Kauffman and Johnson (1990), Kauffman and Mcready(1995)) . An apparently baffling result seems to emerge from these models: in this context, limiting constraints need not be elements necessarily dampening fitness but may rather enhance it. Formally, the reason again lies in the statistical properties of a focused search compared to a random, across-the-board exploration of the whole landscape, the former acting as an efficient device that leads to a high fitness part of it allowing to retrieve highly performing configurations. Take a fitness landscape depicting the likely state of a network of nodes. Each point is then a network state associated to an index of average fitness. The network is, in turn, made up by N nodes taking a characteristic chosen from a list of A possible ones. There is no loss of generality by setting $A = 2$. Each network state or configuration can then be symbolically represented by a string of N nodes named $(c_{i1}, c_{i2} \dots c_{iN})$ where i is the state or configuration and each $c_{ij} \in (0, 1)$ is a specific characteristic. It can further be assumed that each such characteristic randomly carries a fitness such that the average for the whole network is the simple arithmetic mean of the N nodes. There are, accordingly, 2^N points (states or configurations) in the landscape. The scope of this exercise is to show a search for an improvement motion over this landscape as a node attempts to change its characteristic. Let a random search be discussed to begin with. If the N nodes were totally independent, i.e. not linked, a change occurring in any of them would have no effect on the remainder and any fitness improvement would be consolidated given the assumption that a worsening change would be adaptively turned down, for instance by a selection mechanism, and any better one accepted. In these circumstances, the landscape would feature just one clearly identifiable maximum that the network would gradually climb. This would no longer be the case if each node were connected to nodes by a link of interdependence (or epistasis): in this case, a change occurring in one would bear consequences on all those that are so connected. If the degree of interdependence were such that each were influenced by all the remaining $N - 1$, then a newly randomly determined fitness acquired by a node by changing its characteristic would, in turn, randomly change the fitness value of every other node. In this case, a network would stand a probability of having acquired a state carrying the maximum fitness among its neighbouring states, those differing by just one characteristic and hence by just one possible mutation of any of its nodes, that is equal to $\frac{1}{N+1}$. Thus, the expected number of local maxima would be $\frac{2^N}{N+1}$. The implication is that if by happenstance a network hits a local maximum in the above stated sense, it gets locked-in there since any change by one of its node would yield an inferior fitness. Let the degree of interdependence be designated by K . The case that has just been discussed, of total interdependence, corresponds to $K = N - 1$ while $K = 0$ indicates the one of complete independence. The number of local maxima, then, roughly scales according to $\frac{2^{K+1}}{K+2}$ and, accordingly, the probability of locking-in in a local maximum being also the absolute one equals $\frac{K+2}{2^{K+1}}$, falling continuously as $K \rightarrow N - 1$. Fur-

thermore, for $N \rightarrow \infty$ and by the Central Limit Theorem the expected fitness of each local maximum when $K = N - 1$ tends to the average: in case of a uniform fitness distribution over $(0, 1)$ the average would just be .5. The landscape is a flattened one and mediocrity is pervasive. The implication is that as K increases, the probability of getting locked in a poor maximum increases and as N becomes very large and all local maxima would feature roughly the same fitness close to the distribution average. Limiting constraints are here the links of interdependence causing a connectivity and a very rugged fitness landscape to impose a change that can go either way as long as it is random: it can signify a betterment as well as a worsening of the network state. In this sense, constraints play a negative role in searching-for-improvement random walks. Yet, node interdependence within a network, or inner interdependence, may lend a positive influence when a co-evolutionary movement of many networks is involved. The classical example, in this case, is that of co-evolving species in biology and of competing as well as co-operating but heterogeneous firms in an economic system or, more generally, of social organisations. If different networks co-evolve connections are necessarily established between at least two of their nodes. The connectivity that is involved can, for instance, take the form of competition for the same food niche, prey-predator co-evolutionary adaptation, technological competition or just imitation. What is implied by this further complication is that the evolutionary rugged landscapes owing to inner interdependence are dynamically deformed as each network walks on its own. This is due to the fact that as a node carrying a link not only with other members of its own network but also with at least one of another changes its characteristic and thus its fitness, it generates a change not only in the fitness of its own network but also in that of the other network with the possibility of a further feedback that continuously reshapes both landscapes until a lock-in peak is eventually achieved by both. This is an important point. Suppose, for instance, that any of several connected structures reaches what appears to be a local maximum and locks therein. If, however, any of the connected ones adaptively changes its own fitness in consequence, this will in due time produce a response that moves the original one away from its temporary equilibrium on a local maximum. Consider, for simplicity's sake, just two co-evolving networks with coupled landscapes. The ruggedness of both depends, as it has been seen, on the degree K of interdependence. Co-evolutionary equilibrium occurring when both networks get locked-in on a local maximum depends, from a probability viewpoint, on how many of the latter exist, a magnitude that scales as $\frac{2^{K+1}}{K+2}$. If they are not both at rest, their fitness keeps oscillating and it will do so by threading a path that lies on the landscape 'valley', somewhere between this low and the high of local maxima and occasionally on the latter from which they are, however, budged out as long as any of the two keeps changing its fitness. It is, then, clear that the expected average fitness of out-of-equilibrium networks, that is while they keep on oscillating, is lower than when both rest on local peaks. Clearly, the higher is the inter-network connectivity, i.e. the number C of edges straddling the two networks and connecting as many nodes, the higher

is the likelihood that they be out of equilibrium. Likewise, the time required to simultaneously reach equilibrium is accordingly higher. These results apply straightforwardly if the number of co-evolving networks is larger than two and the higher is the number of nodes involved in inter-network evolution. It follows, in this more general case, that the higher is C the lower is the percentage of networks that at any time have reached a *stasis*, given, it is important to stress, the average K . It also follows that, given the latter, the expected average fitness of out-of-equilibrium networks is lower than that of those that are in equilibrium.

In this framework, K , measuring the inner limiting constraints, has a very interesting role to play. As noted before, increasing it means augmenting each landscape ruggedness. When $K = 0$, given N , the network cardinality, there is just one maximum and, furthermore, the fitness of neighbouring points (networks differing by just one node characteristic) is highly correlated since the difference must be ascribed to just one, and only one node. For $K = 2$ and 3 , there are roughly two local maxima and the correlation is still high, albeit somewhat smaller. It follows that the two maxima are quite near each other and the expected average fitness somewhat lower than in the case of just one maximum. As K rises, fitness becomes less correlated and the expected average fitness of local maxima lower. When cast in a co-evolutionary framework, landscapes are, in a manner of speaking, pitted against each other: a higher K insures a higher probability of locking-in into equilibrium which in itself implies a higher expected average fitness than if out of it but at the same time landscapes become flatter. There clearly is a trade-off. For a given C , there is likely to exist a value $K = K^*$ that insures the highest achievable fitness: before and past K^* average fitness is lower. Thus, if on the one hand limiting constraints render the search for improvement less enticing, on the other they increase the probability of reaching an equilibrium the fitness of which is higher than if oscillations occur.

These results hold when walks are random, clearly the case when changes are mutations of biological entities. Within limits, however, they can be shown to hold when randomness is paired with deterministic interdependence (Ricottilli (1999)). The latter, for instance, is certainly the rule in most social networks. Firms and organisations can be regarded as networks of elements or nodes of given cardinality N that are mutually interdependent. Links in this case establish a functional relationship such that the fitness of each can be posited as a function of part or all the remaining ones. Networks are specific but so are the single nodes that make them up since each of them serves a specific purpose and therefore possesses its own adaptation mechanism. This implies that each node generally links up with other nodes thanks to its own specific function. Consider the following example. Let V_i be the fitness of the i -th landscape configuration and let V_{li} be that of the l -th node: $i = 1, 2, \dots, 2^N$ and $l = 1, 2, \dots, K$. The system

$$V_{li} = F_{li}(V_{1i}, \dots, V_{Ki}; \bar{V}_{K+1,i} \dots \bar{V}_{Ni})$$

$\forall l$, describes the links of deterministic interdependence tying K nodes in the i -th configuration whilst $N - K$ are independent of any other node and whose fitness $(\bar{V}_{K+1,i}.. \bar{V}_{N,i})$, is historically determined. In a likewise manner to the case discussed above, we would therefore say that the degree of interdependence is K . Each V_{ji} is the fitness of an organisational *locus*, member of a network composed of a fitness string $(V_{1i}, ... V_{ki}... V_{Ni})$ while the total V_i is, as usual, the simple average. Each of these loci carries on a search for betterment ruled by an adaptation mechanism that eventually leads to innovation. Therefore, each locus exhibits a fitness that is actually a time variable whose scale is the innovation events that occur anywhere in the string. There is no loss of generality, for the purpose of this paper, to render the above indicated system in a very simple linear and dynamic form:

$$V(t) = AV(t) + B\bar{V}(t)$$

Where $V(t) = (V_1(t), ... V_K(t))$ and $\bar{V}(t) = (\bar{V}_{K+1}(t).. \bar{V}_N(t))$ are column vectors. Note that t is a time variable and such that a change from t to $t + 1$ marks a change in a configuration thanks to one node changing its characteristic; hence, $t + 1 \in (t, t + 2, ... t + N)$, i.e. it belongs to a local neighbourhood and such that $(t, t + 2, ... t + N) \subseteq (1, 2... 2^N)$, that is a subset of the entire landscape. A is a K^2 square matrix where each component a_{ik} denotes the index of dependence of node i from the K that compose the strictly interdependent part of the network and B is a $K \times (N - K)$ rectangular matrix linking the remaining $(N - K)$ nodes to the first K according to coefficients b_{ir} . For the network to be viable for a current configuration t , this simple system must admit solutions:

$$V(t) = (I - A)^{-1} B \bar{V}(t)$$

and the usual conditions for this to hold must be satisfied. While this system describes the deterministic part, it can reasonably be assumed that an innovation, for instance a change in any $\bar{V}_r(t) \in \bar{V}(t)$ be a random event. Suppose that this is actually the case and that $\bar{V}_r(t) \rightarrow \bar{V}_r(t + 1)$ all other independent nodes staying the same. The system now solves for $V(t + 1) = (I - A)^{-1} B \bar{V}(t + 1)$ where $\bar{V}(t + 1) = (\bar{V}_{K+1}(t + 1) = \bar{V}_{K+1}(t); ..; \bar{V}_r(t + 1); ... \bar{V}_N(t + 1) = \bar{V}_N(t))$. Since the node that changes does so randomly, $\bar{V}_r(t + 1)$ is a random variable but so are all the $V(t + 1)$ that depend on it. In this sense all the $\bar{V}(t + 1)$ are historically but randomly determined. Note that in this example $K + 1$ fitness variables change value, one autonomously the other ones in consequence. It is also interesting to note that if one of the $V(t)$ were to change autonomously the system would simply exhibit a lower rank, the number of independent variables rising to $N - K + 1$. Let it be now assumed, on the contrary, that the number of fully dependent variables rose to $K + 1$. If one of the $N - K - 1$ independent ones were to randomly change, there would obviously be $K + 1$ changing in consequence. If the fully dependent were $K + 2$ these would also change for any autonomous shift in the independent variable. If all nodes were interdependent, the system would reduce to an homogeneous one and for any random change in one node all the remaining $N - 1$ would change accordingly.

The fitness correlation would clearly fall the greater is the number of variables involved in the change and the least correlation would be exhibited by the full interdependence case. Thus, ruggedness increases as correlation falls. The case discussed is simply an extension of the main results of random walks on rugged landscapes when a random change occurs as a consequence, say, of a search for innovation. It is straightforward, therefore, that the main results discussed in the co-evolutionary case of many networks being involved in the dynamics of change hold *a-fortiori*.

These general results can now be put to use to make a last point. The inner and outer connectivities measured by average K and average C respectively are crucial to determine the likelihood that a system reach equilibrium, that is a general lock-in into local maxima. If there are many networks in a given system whose nodes are outward connected by an average C , it is the $\frac{K}{C}$ ratio that rules how many of them reach a *stasis* and instead how many keep on oscillating. The greater are the outer links, C , the more likely are oscillations which are, however, fettered by K . These facts lend themselves to an interesting interpretation. Let f^* denote the share of all networks that over a conventionally long period of time have reached equilibrium in the sense employed above. It is clear that $f^* = f(\frac{K}{C})$ and $\frac{K}{C}$ can be viewed as the index setting the probability that any two networks be simultaneously in equilibrium. Simulations carried out on a lattice in which a given number of networks is placed and such that adjacent ones interact show that as $\frac{K}{C}$ rises, f^* rises very slowly and finally tends rapidly to 1 as $\frac{K}{C} \rightarrow (\frac{K}{C})^*$, a result that can be interpreted in terms of a percolation process. Thus, the balance of inner versus outer constraints while leading searching to parts of the space with relatively high fitness decides also of the system eventual equilibrium or persistent chaos.

4 Constraints to patterns of searching

Spatial connectivity and its intensity within and among networks underscore the utmost importance of constraints in establishing the quality and dynamics of the network configuration and eventual equilibrium. This importance lies in the role they play as devices that capture searching efforts towards efficient sections of the space of mutational possibilities. Constraints, however, have an equally crucial impact on the outcomes of searching patterns by shaping how the process unfolds. This is a very relevant issue in social networks, and more broadly in social systems, in which this process is performed according to rational, albeit bounded, procedures. Yet, procedures of an adaptive kind fashioned on a highly different time scale are relevant in biological systems as well. The following discussion will briefly dwell upon some recent results concerning searching in complex technological spaces and on the spreading of technological information. Recent work, (Silverberg and Vespagen, 2005a, 2005b), has shown a trade-off to exist between search difficulty and the relative freedom of agents, in their case firms, to move on a space of technological possibilities. By relative freedom

it is meant a behavioural routine that makes agents capable of changing their technological specificity rather than remaining bounded on a neighbourhood of their achieved position. It is interesting to report the basic features of this model since it is a good illustration of the role of constraints. The space of technological possibilities is defined as a Manhattan grid lattice, unbounded from above, at the baseline of which firms are situated, one in correspondence of each cell. Cells are randomly seeded with technological opportunities of varying difficulty: no technology is impossible but all have a discovery cost; some have one that tends to infinity. A distribution of difficulties is thus assumed: from very easy, nearly costless to discover, to exceedingly difficult and hence hugely expensive. Different technological environments can then be fashioned by assuming initial distributions with different mean, \bar{q} , and variance, σ_q . Firms carry out a local search by exploring near-by cells according to a diamond shaped radius of a given magnitude m . To achieve this goal they invest an R&D budget evenly spread on the cells within the radius with the result of lowering the difficulty level. Once discovered a technology need not be immediately implemented since to become a viable innovation it must be linked all the way to the baseline with already viable technologies. The purpose of this effort is to rise as high as possible on the lattice, the height achieved and connected to the baseline measuring a success carrying a pay-off that then enriches the R&D budget: a success breeds success feed-back mechanism. The following are the essential equations that are then simulated.

The difficulty level: $q_{ij,t+1} = q_{ij,t} - b\omega$

ij define the ij -th cell given column i and row j ; b is the amount spent on searching any of the nearby lattice cells that yields results with probability ω . In turn, $b = \frac{B_t}{2m(m+1)}$, where $2m(m+1)$ is the number of cells that can be explored in a Manhattan grid neighbourhood with radius m .

The firm's budget : $B_t = \beta + \sum_k s_{k,t-1}\pi$

β is a given budget level equally assigned to all firms at the outset of the period, $s_{k,t-1}$ measures the gain in technology levels accrued in the previous period thanks to the k active cells discovered because of R&D and to each of which a pay-off π is associated.

The definition of the best practise frontier:

$$BPF(t) = \{i, j(i); i = 1, N_c; j(i) = (\max j \mid a_{ij} = 2)\}$$

where a_{ij} defines the technology state of each cell: $a_{ij} = 0$ is an undiscovered technology, $a_{ij} = 1$ is a discovered but as yet not viable one, i.e. not being fully connected to the baseline, $a_{ij} = 2$ is a discovered and viable one.

$$p_{ij} = \frac{u_j}{U}; \quad U = \sum_j u_j; \quad u_j = e^{h_j - h_i}$$

states the probability of a free-to-move firm to go from the BPF in column i where it is located, when given a chance to move, to column j if it pays to do so, namely if the difference in height $h_j - h_i \geq 0$.

Simulations are run by seeding the lattice at the outset with a distribution of q , $N(\bar{q}, \sigma_q)$ and by testing the performance for parameters $\bar{q}, \sigma_q, m, \pi$ in the

two different behavioural protocols.

The simulations based on this model show quite rich and interesting results but for the purpose of this paper it is fitting to concentrate on the following ones. For any given \bar{q} as σ_q rises, free to move firms do progressively better than the ones that are constrained on their technological area. This holds mildly for a given, low value of \bar{q} but it becomes distinctly so as the latter rises, that is as the average level of difficulty to discover rises. Yet, for any given level of σ_q free to move firms exhibit higher innovation rates for low levels of \bar{q} but progressively worse than the constrained firms. It is interesting to note that this result holds for any level of the searching radius m but becomes more pronounced the greater it is. These apparently puzzling results are explained by the underlying trade off between firms' relative freedom to shift from their technological position and the level of search difficulty. The reason lies with the behavioural routines assigned to firms in relation to the statistical properties given to the lattice but are quite illuminating of actual situations. Free-to-move firms tend to orient their search in sections of the space where technological opportunities seem to be higher, namely by shifting to higher points on the extant technological frontier from whence they carry out their searching. By so doing they all tend to cluster in a few and specific sections of the search space. On the contrary, firms constrained to narrowly remain in the position they have acquired rest widely distributed on the latter. When \bar{q} is high but smoothly distributed, constrained firms have a nearly equal chance of dealing with very difficult and not so difficult technological opportunities whereas the non-constrained ones stand a higher probability of locking in very difficult sections of this space and hence perform less well. This situation is reversed when σ_q is high and especially when \bar{q} is also high: in this case, when the variance is high, there are likely to be few sections of the lattice with very high q 's that mostly concern the very distributed agents (firms), those that are constrained, whilst the non constrained and free-to-move ones tend to go on specific parts of the space that have a lower probability of being difficult to explore. It is now them that do better. The tendency to specialise on a few sections by clustering seems to cut in more than one way: it yields better results when searching difficulties feature a high standard deviation but worse ones when difficulties appear high and evenly distributed. In the latter case, the behavioural routine that constrains agents to be kept active all over the lattice insures a better performance.

In foregoing paragraphs a discussion has been made on the process leading to technological discovery or invention and then to innovation. In this context, firms act alone but it must be recognised that the innovative process is very much the upshot of interaction tying a whole network of firms. Recent models have highlighted the role of routines, featuring constraints in varying degree, both in determining fitness, or performance in economic matters, and the emergence of linkage patterns between agents. Reference can usefully be made to Andergassen, Nardini and Ricottilli (2005a, 2005b, 2006). The issues set forth in these papers deal with the development of innovative capabilities, or of technological fitness, by firms that actively carry out a searching process for this very purpose. It is precisely the procedures employed to unravel its

difficulties, the actual routines followed by these firms that make a remarkable difference on results that are ultimately achieved. In these models firms owe their innovative capabilities to two activities that although separate add up to a functional whole. The first is an in-house R&D activity leading to an autonomously generated capability whilst the second is an activity aimed at observing other firms' capabilities in view of enhancing their own because of a spillover effect. On account of a broadly defined bounded rationality, firms acting as nodes of a network are assumed able to glean the information broadcast by other firms only within a local, mainly cognitive, neighbourhood made up by members that can actually be reached out. They are further posited to evaluate the contributions the latter supply and update this membership by eschewing the least performing neighbour to choose a new one conditionally to the achievement of a higher capability. It is precisely how this process of sorting out the least performing and filling in one that it is more so, i.e. what is the routine that is implemented, that decides of final performance as well as of the emergence of technological paradigm setters. The latter are considered to be those firms that become the main source of information for the rest of the network or the whole economy. The model set out to describe the network evolution in this context shows firms to avail themselves of two broadly defined routines to proceed to the updating task involved in this process. The first consists, once the worst performing neighbour has been singled out, in randomly drawing a new one from members of the entire network (the economy): an across-the-board, unconstrained search. The second, on the contrary, constrains firms to draw a new neighbour from the forsaken neighbour's neighbourhood, a procedure that rests on the idea of seeking out a new contributor through a very local and hence constrained search.

The basic features of this simple model are the following. Let N be the number of firms in an economy. $C_i(t) \in (0, 1)$ defines for any node i , $i = 1, 2, \dots, N$, a level of in-house, autonomously developed technological capability or fitness; $V_i(t)$ is the total one obtained by adding to the former the spillovers broadcast by other firms and actually retrieved by node i :

$$V_i(t) = \sum_{j=1}^N a_{ij} b_{ij}(t) V_j(t) + C_i(t); \quad i = 1, 2, \dots, N$$

where a_{ij} indicates the constant part of each firm j 's total capability that can cognitively be passed on to firm i , $a_{ii} = 0$, $\forall i$. $b_{ij}(t)$ is an element of the adjacency matrix $B(t) = [b_{ij}(t)]$ and such that $b_{ij}(t) = 1$ or $= 0$ according to whether neighbour j has or has not been identified as a useful contributor. $B(t)$ evolves as a result of searching for a better neighbour subject to the condition $V_i(t) > V_i(t-1)$.

These equations define a system of evolving capabilities or fitnesses:

$$V(t) = [I - M(t)]^{-1} C(t)$$

where $M(t) = [a_{ij} b_{ij}(t)]$ and $V(t)$ and $C(t)$ are column vectors. This system

is tested by simulations that assume a constant and given number of neighbours, $k_{in} = \bar{k}$, that every single firm observes and that make up what can more formally be termed its *inward* neighbourhood. This contrasts with the definition of its *outward* neighbourhood. The latter is simply made up by the firms that observe it and it actually evolves as firms within this virtual economy change their neighbours. This definition allows to define more formally the notion of technological paradigm setters. *Outward* neighbourhoods are ranked according to the number of firms that the simulations classify into given cohorts (1-2, 3-4, etc.). Firms that possess a positive probability of having an *outward* neighbourhood comprising most other firms are defined as paradigm setters. If no perturbations took place, this system would converge to the discovery of the best performing neighbours and all would, in time, achieve the same fitness. The system, however, is shocked by randomly allowing the in-house capability to change by drawing a new $C_i(t) \in (0, 1)$ every μ periods of time on average. This mean waiting time μ is then a simulation tuning parameter. The other important parameter that is tuned is τ , the one measuring the routine that implements the neighbour switching procedure. By calling π the probability of exploring across the entire firms' space *once* in a given number n of attempts to change a randomly chosen neighbour and defined as $\pi = \frac{1}{n}$, a definition of τ is obtained: $\tau = \frac{1}{\pi}$; $\tau = 1, 2, 3, \dots \rightarrow \infty$. When $\tau = 1$, firms are called to change the chosen neighbour by always looking over the whole economy, when $\tau = 2$, the routine is once over the whole economy and once inside their own downcast neighbour's neighbourhood, when $\tau = 3$ once outside and two inside, when $\tau \rightarrow \infty$ firms are always constrained to draw from their neighbour's neighbourhood. τ , therefore, sets a scale of constrained behaviour: from entirely free to wholly constrained.

Simulations results show that for $\tau = 1$, the free to choose routine, average performance is relatively mediocre and no paradigm setters emerge. As τ rises, as behaviour adopts more constraining routines, performance also rises and paradigm setters emerge. This trend continues for ever more constraining routines up to τ^* , that simulations reveal to be in the 3-4 area, but then average performance begins to fall and finally plunge as $\tau \rightarrow \infty$. This pattern holds for all mean waiting times μ . These results are due to the fact that when the routine is entirely unfettered when it comes to changing to a better neighbour, firms draw from an urn in which highly performing and just barely better ones have a nearly equal chance of being chosen. As routines become more constrained and more frequently focused on an eschewed neighbour's neighbourhood they stand a higher probability of getting a good performing neighbour by searching locally: when searching chances on a well performing neighbourhood, it pays to lock in there. Yet, as chance may lower by a random event the performance of even highly performing firms, it pays to look outside for better alternatives: it is, in other words, good to look 'abroad' to renew the stock of useful capabilities to learn from. Quite generally, specialisation in a section of the search space is good up to a point; to the point where allowance is made for broader searching and learning.

5 Conclusions

It is now possible to draw some tentative conclusions. This paper is about the role played by constraints in the process of searching over a space of possibilities. The main message that is set out in foregoing pages is that this process normally implies a fine balance between relations that foster specialisation between nodes and relations that, on the contrary, promote generalised if not indiscriminate exploration of this space. This balance hangs on fine-tuning some parameters that produce the mentioned phenomena as well as a wide variety of intermediate states including power law distributions. It is argued that there exists a mounting body of evidence showing that the development and self organisation of constraints operates to determine where the balance between these two features of networks actually lies. The statistical properties of the two opposing configurations indicate that a fitness trade-off emerges. This feature clearly appears in different but complementary constraint settings. Since the relational architecture of networks determines the search-for-fitness interaction, the way the inherent topology is laid out bears influence on their performance, be they biological or social ones. Thus, constraints that generate clustering and hence specialised relations producing positive externalities between nodes are conducive to greater fitness, yet they may hamper fruitful linkages with other clusters that carry idiosyncratic properties, specific functional characteristics, specific knowledge or innovative capabilities. Each cluster explores, in this sense, a space of possibilities whose limits can be overcome by letting a few links reach out to other clusters in such a way that some nodes may function as informational hubs expanding the search space that each is able to explore. Excessive linkage generates connectivity randomness and wastage of searching energy. This trade-off is observed when network development algorithms are considered. In this framework, the actual adaptive method implemented to weigh the strengthening of a link between nodes against a threshold compounding past experience decides of the final network topology and whether specialisation rather than across-the-board linking prevails.

Fitness trade-offs conclusively appear when interdependence between nodes and, in a co-evolutionary context, between networks is the important feature considered. Here constraints making nodes reciprocally interdependent and mutually limiting while producing fitness mediocrity may also favour highly performing equilibria when pitted against the constraints that make whole networks mutually interdependent. The balance stems from the fact the former focus or specialise the search space by increasing the probability of locking-in into equilibrium of local fitness maxima while the latter induces continuous exploration of landscapes and hence fitness oscillations. Fine tuning and self organising the two is the key, on the one, hand to achieve a relatively 'optimum' fitness and, on the other, to avoid chaos. Finally, the specialisation versus generalised searching trade-off can clearly be observed when different behavioural routines, in social systems, or procedural methods, in biological ones, are considered. The models discussed in previous pages show that routines applying constraints that focus searching in specific sections of the space of possibilities can be conducive to

high fitness but only to the point that they do not prevent the renewal of positive probabilities of better fitness to be retrieved in far lying parts of the space. This feature emerges in models that consider searching in complex technological spaces as well as in those that investigate knowledge spillovers. In both, it is not necessarily the routines that allow the greater search freedom that prove to be the fittest but those that, once again, balance out focused and specialised searching within a local neighbourhood with the once-in-while option of spanning the whole system. Fine-tuning and self organising the two emerge as the winning solution.

6 References

- Albert R., Barabasi A.L. (2002): 'Statistical Mechanics of Complex Networks'. *Reviews of Modern Physics*, Vol. 74, pp.47-97.
- Andergassen R., Nardini F., Ricottilli M. (2005a): 'Firms' Network Formation through the Transmission of Heterogeneous Knowledge'. *Working Paper of the Department of Economics*, University of Bologna, n. 543.
- Andergassen R., Nardini F., Ricottilli M. (2005b): 'Technological Paradigms and Firms' Interaction'. In Léskow J., Puchet M., Punzo L. (eds.): 'New Tools of Economic Dynamics'. *Lecture Notes in Economics and Mathematical Systems*. Vol. 551. Springer Verlag, Berlin.
- Andergassen R., Nardini F., Ricottilli M. (2006): 'The Emergence of Paradigm Setters through Firms' Interaction and Network Formation'. in Namatame A., Kaizouji T., Aruka, Y.: 'The Complex Networks of Economic Interactions'. *Lecture Notes in Economics and Mathematical Systems*, Vol. 567. Springer Verlag, Berlin.
- Andergassen R., Castellani G.C., Nardini F. Remondini D., Ricottilli M. (2006): 'Node Evolution from Randomness to Specialisation'. *Mimeo*. Centro Interdipartimentale 'Luigi Galvani', Bologna.
- Auerswald P., Kauffman S., Lobo J., Shell K. (2000): 'The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning by Doing'. *Journal Of Economic Dynamics and Control*, Vol.24, pp. 369-450.
- Boccaletti S., Latora V., Loreno Y., Chavez M., Huang D.U. (2006): 'Complex Networks: Structure and Dynamics', *Physics Reports*, 424, pp.175-308.
- Castellani G.C., Intrator N., Shouval H., Cooper L.N. (1999): 'Solutions to the BCM Learning Rule in a Network of Lateral Interacting Nonlinear Neurons'. *Network*, Vol. 10, pp.111-121.
- Intrator N., Cooper L.N. (1992): 'Objective Function Formulation of the BCM Theory of Cortical Plasticity: Statistical Connections, Stability Conditions'. *Neural Networks*, Vol.5, pp 3-17.
- Kauffman S.A. (1989): 'Adaptation on Rugged landscapes', in Stein D. (ed): '*Lectures in the Science of Complexity*', Addison Wesley Longman.

- Kauffman S.A., Johnson S. (1990): 'Co-evolution to the Edge of Chaos', in Langton C.G. et Allii (eds.): '*Artificial Life II: Studies in the Science of Complexity*', Addison Wesley Longman.
- Kauffman S.A., Macready W. (1995): 'Technological Evolution and Adaptive Organisations', *Complexity*, 1, 26-43.
- Ricottilli M. (1999): 'Innovative Change, Search and Division of Labour'. *Proceedings of the EAEPE Conference*, Lisbon.
- Silverberg G., Verspagen B. (2005a): 'A Percolation Model of Innovation in Complex Technology Spaces'. *Journal of Economic Dynamics and Control*, Vol. 29, pp. 225-244.
- Silverberg G., Verspagen B. (2005b): 'Self-organization of R&D Search in Complex Technology Spaces'. *Merit-Infonomics Research Memorandum Series*, n.017.
- Watts D.J., Strogatz S.H.(1998): 'Collective Dynamics of "Small Worlds" Networks. *Nature*. Vol. 393,pp. 440-442.