



Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)
Tel. +39-050-883-343 Fax +39-050-883-344
Email: lem@sss.up.it Web Page: <http://www.lem.sssup.it/>

LEM

Working Paper Series

**Toward Formal Representations of Search
Processes and Routines in Organizational
Problem Solving. An Assessment of the State of
the Art**

Giovanni DOSI*
Marco FAILLO*^o
Luigi MARENGO*

* LEM, Sant'Anna School of Advanced Studies, Pisa, Italy
^o University of Trento, Italy

2006/10

April 2006

ISSN (online) 2284-0400

TOWARD FORMAL REPRESENTATIONS OF SEARCH PROCESSES AND ROUTINES IN ORGANIZATIONAL PROBLEM SOLVING. AN ASSESSMENT OF THE STATE-OF-THE-ART

Giovanni Dosi*, Marco Faillo** and Luigi Marengo*

•

•

* LEM -*Sant'Anna School of Advanced Studies, Pisa*

University of Trento, Italy

(this version:8 april 2006)

Abstract

This paper presents a critical overview of some recent attempts at building formal models formalizations of organizations as information-processing and problem-solving entities.

We distinguish between two classes of models according to two distinct objects of analysis. The first class includes models mainly addressing information processing and learning and analyze the relations between organizational performance, learning patterns and the structure of information flows. The second class includes models focusing upon the relationship between the division of cognitive labor and search process in some problem-solving space, addressing more directly the notion of organizations as repositories of problem-solving knowledge. Here the focus is on the problem-solving procedures which the organization embodies.

The results begin to highlight important comparative properties regarding the impact on problem-solving efficiency and learning of different forms of hierarchical governance, the dangers of lock-in associated with specific forms of adaptive learning, the relative role of “online” vs. “offline” learning, the impact of the “cognitive maps” which organizations embody, the possible trade-offs between accuracy and speed of convergence associated with different “decomposition schemes”.

We argue that these are important formal tools towards the development of a comparative institutional analysis focusing on the distinct properties of different forms of organization and accumulation of knowledge.

Keywords: Information processing, Problem-solving, Organizational structure.

1. Introduction¹

This work is meant to offer a critical overview of the achievements and challenges ahead facing explicit formalizations of organizations as information-processing and problem-solving entities.

The importance of the information-processing arrangements is well acknowledged within both *agency* and *capability*-based theories of the firm, even if only the latter focuses on the *problem-solving* features of organizations.

However, most formal representations of such activities tend to offer highly *blackboxed* accounts. In that agency models are an extreme case to the point where the whole activity of information processing is compressed in some function maximization conditional on the appropriate processing of the available information. On the contrary, here we shall survey those endeavours which try to account for organizational information processing and problem-solving in terms of explicit sequences of activities and procedures nested into specific organizational arrangements prescribing "who send which signals to whom" and "who does what and in which sequence".

The appreciative theories upon which such model draw represent a small – but not negligible and growing – minority of the economic profession who place their “primitives” of the nature of economic organizations are placed in their *problem-solving features*, in turn nested in ubiquitous forms of human “bounded rationality”, grossly imperfect processes of learning and diverse mechanisms of social distribution of “cognitive labor”. The root of this approach can be found in the works of Herbert Simon, James March, Alfred Chandler and Richard Nelson and Sidney Winter².

The problem-solving activities of the firm can be conceived as combinations of physical and cognitive acts, within a procedure, leading to the achievement of a specific

¹ The work draws upon other works of the authors, in particular: Cohen *at al.* (1996), Dosi, Nelson and Winter (2000), Marengo and Dosi (2005), which the reader is referred to for further details.

² See Chandler (1977), Cyert and March (1963), March and Simon (1993), Nelson and Winter (1982), Simon (1962) and (1981).

outcome. Its internal organization determines the distribution of the informational inputs across specific task units and, as such, the division of the cognitive labor. The general idea is that firms possess the specific problem-solving competencies associated with their own operational procedures and routines, in turn embedded into the patterns of intra-organizational division of labor and assignments of decision entitlements.

An illustrious antecedent of this view dates back, indeed, to Adam Smith's "Pin Factory" example in the *Wealth of Nation*:

"One man draws out the wire, another straightens it, a third cuts it, a fourth points it, a fifth grinds it at the top for receiving the head; to make the head requires two or three distinct operations; to put it on, is a peculiar business, to whiten the pins is another; it is even a trade by itself to put them into the paper; and the important business of making a pin is, in this manner, divided into about eighteen distinct operations, which, in some manufactories, are all performed by distinct hands, though in others the same man will sometimes perform two or three of them."(Smith, 1776).

And, relatedly, such patterns of division of labor match specific channels of information flows and "lines of command".

How does one formalize these basic intuitions?

It is fruitful to distinguish between two classes of models according to two distinct objects of analysis. The first class includes models mainly addressing information processing and learning. Here the focus is on the relation between organizational performance, learning patterns and the structure of information flows. Agents are adaptive learners who adjust their information processing capability (i.e. their knowledge of the environment) through local trial-and-error.

The second class includes models focusing upon the relationship between the division of cognitive labor and search process in some problem-solving space, addressing more directly the notion of organization as repositories of problem-solving knowledge. Here the focus is on the problem-solving procedures which the organization embodies. After all, managing an organization, designing and producing cars or software packages,

discovering a new drug, etc. can be seen as a complicated problem whose “solutions” are made of a large number of cognitive and physical acts. This kind of activities imply the coordination of large combinatorial spaces of components.

At one end, components which make up an artifact can take a number of alternative states: so, for example, in the case of the production of a car, one combines different characteristics of the engine, alternative designs, different materials etc. Conversely, innovative search may be straightforwardly represented in form of combination of multiple “cognitive acts” eventually yielding the solution of the problem at hand, e.g. the discovery of a new molecule with the required characteristics, a reasonable and coherent software package, etc. Note that in both examples the existence of strong interdependencies among the components – which often are only partially understood by all agents involved - implies that the effect on the system’s performance of a change in the state of a single component depends on the values assumed by the other ones. An implication is also that in this kind of problems it is impossible to optimize the system by optimizing each single component.

By applying this view to organizational analysis one can conceive economic organizations as bundles of routines, procedures, rules characterized by strong interrelations which often are opaque to organizational members. Notice first the partial “opaqueness” of the mappings between actions and outcomes is quite in tune with “garbage can” interpretation of organizational dynamics (Cohen et al. 1972). Second it is well corroborated by plenty of evidence regarding widespread difficulty in *replication* and *transfers* of incumbent organizational arrangements (Winter and Szulanski, 1998, 2002; Zander and Kogut (1995)). Third, an obvious implication of such partly opaque interrelatedness is also that the introduction of a new routine which has proven superior in another situation might have negative effects on the performance of the organization if other interrelated components are not appropriately co-adapted (Marengo and Dosi, 2003: 8-9; Marengo et al 2000).

2. Information processing and structural learning

Marengo (1992) and Marengo (1996) present a model which focuses upon the modification of agents' information processing capabilities, i.e. a process of "structural" learning. Individual agents are imperfect adaptive learners, as they adjust their information processing capabilities through local trial-and-error. This adaptive learning is (at least partly) driven by the information coming from the environment and/or from other members of the organization. The model shows that the architecture of such information flows plays a crucial role in determining the learning patterns and the performance characteristics of the organization.

One begins by considering a standard problem of individual decision making, which will be then extended to a collective one. Let

$$S = \{s_1, s_2, \dots, s_N\}$$

be the set of the N possible states of nature and

$$A = \{a_1, a_2, \dots, a_N\}$$

the set of the k possible actions the decision-maker can undertake. The payoff to the agent is given by a function:

$$\Pi: A \times S \rightarrow \mathbb{R}$$

where the agent's payoff to action a_i when the state of the world s_i occurs will be indicated by π_{ih} .

The action the agent chooses depends obviously on the level of his or her knowledge about the state of the world. The agent's state of knowledge (or information processing capabilities) can be represented by a collection of subsets $P(s_i) \subseteq S$ where $P(s_i)$ is the set of states of the world which the agent considers as possible (or cannot tell apart) when the real state is s_i .

The basic component of this learning system is a condition-action rule, where the execution of a certain action is conditional upon the agent's perception that the present

state of the world falls in one of the categories he or she has defined in his \mental model. The condition part is a category, that is a subset of the states of the world and is activated when the last detected state of the world falls in such a subset. Practically, the condition is a string of n symbols (as many as the states of the world) over the alphabet $\{0,1\}$ and it is satisfied whenever the last state of the world corresponds to a position where a “1” appears. All in all, the condition:

$$c_1c_2\dots c_N \text{ with } c \in \{0,1\}$$

is satisfied when, if s_i is the last observed state of the world, we have $c_i = 1$. Thus, a set of conditions defines a subset of the power set of S. It is important to notice that each condition defines one subjective state (or category) of the world, as perceived by the agent and defines its relationship with the objective true states of the world. This relationship remains anyway unknown to the decision maker, who knows only the subjective states.

The action part is instead a string of length k (the number of the agent's possible actions) over the same alphabet and with the following straightforward interpretation:

$$a_1a_2\dots a_k \text{ with } a_i \in \{0,1\}$$

which has one and only one position which equals “1” and “0's” everywhere else.

The decision maker can be therefore represented by a set of such condition-action rules:

$$R_1 = \{R_1, R_2, \dots, R_q\}$$

where:

$$R_1 : c_1, c_2 \dots c_N \Rightarrow a_1 a_2 \dots a_k \text{ with } c_i, a_h \in \{0,1\}$$

In addition, each rule is assigned a “strength” and a “specificity” measure.

Strength basically measures the past usefulness of the rule, that is the rule's cumulated payoff. Specificity measures the strictness of the condition: the highest specificity (or lowest generality) value is given to a rule whose condition has only one symbol “1” and therefore is satisfied when and only when that particular state of the world occurs,

whereas the lowest specificity (or the highest generality) is given to a rule whose condition is entirely formed by “1's” and is therefore always satisfied by the occurrence of any state of the world.

In this *genre* of models, at the beginning of each simulation the decision maker is supposed to be completely ignorant about the characteristics of the environment he or she is going to face: all the rules initially generated have the highest generality, meaning that all their conditions are formed entirely by 1's. The action parts are instead randomly generated.

The decision maker is also assumed to have limited computational capabilities, therefore the number of rules stored in the system at each moment is kept constant and relatively small in comparison to the complexity of the problem which is being tackled.

This set of rules is processed in the following steps throughout the simulation process:

1. *Condition matching*: a message is received from the environment which informs the system about the last state of the world. Such a message is compared to the condition of all the rules and the rules which are matched, i.e. those which apply to such a state of the world, enter the following step.
2. *Competition among matched rules*: all the rules whose condition is satisfied compete in order to designate the one which is allowed to execute its action. To enter this competition each rule makes a bid based on its strength and on its specificity. In other words, the bid of each matched rule is proportional to its past usefulness (strength) and its relevance to the present situation (specificity):

$$Bid(R_i, t) = (k_1 + k_2 Specificity(R_i)) Strength(R_i, t)$$

where k_1 and k_2 are constant coefficients. The winning rule is chosen randomly, with probabilities proportional to such bids.

3. *Action and strength updating*: the winning rule executes the action indicated by its action part and has its own strength reduced by the amount of the bid and increased by the payoff that the action receives, given the occurrence of the “real” state of the world. If the j^{th} rule is the winner of the competition, we have:

$$Strength(R_j, t+1) = Strength(R_j, t) + Payoff(t) - Bid(R_j, t)$$

4. *Generation of new rules*: the system must be able not only to select the most successful rules, but also to discover new ones. This is ensured by applying genetic operators which, by recombining and mutating elements of the already existing and most successful rules, introduce new ones which could improve the performance of the system. Thus new rules are constantly injected into the system and scope for new opportunities is always made available and such new rules are obtained by recombining and/or locally modifying existing knowledge.

Genetic operators generate new rules which explore other possibilities in the vicinity of the currently most successful ones, in order to discover the elements which determine their success and exploit them. Search is not completely random but influenced by the system's past history. New rules take the place of the currently weakest ones, so that the total number of rules is kept constant.

In Marengo (1992) and (1996) two genetic operators have been used for the condition and one for the action part. The latter is a simple type of local search and is simply a mutation in the “vicinity”: the action prescribed by the newly generated rule is chosen (randomly) in the close proximity of the one prescribed by the parent rule. More concretely, a mutation in the action part will probabilistically mutate the product type prescribed by the rule into one of the neighbouring product types.

The two operators used for the condition part deserve more attention because of their role in modelling the evolution of the state of knowledge embedded into the system. They operate in opposite directions:

- *Specification*: a new condition is created which increases the specificity of the parent one. Wherever the parent condition presents a 1, this is mutated into a 0 with a given small probability;
- *Generalization*: the new condition decreases the specificity of the parent one. Wherever the latter presents a 0, this is mutated into a 1 with a given small probability.

Note that *specification* and *generalization* stand for two possible "cognitive" strategies which tend to drive the learning system towards, respectively, specific rules which apply to more specific states of the world and more general rules which instead cover a wider set of states of the world. Different degrees of specification and generalizations can be simulated both by means of different combinations of these two genetic operators and by varying the coefficient k_2 with which specificity enters the bid equation: the higher this coefficient, the more highly specific rules will be likely to prevail over general ones. The simulations discussed below use a *specificity coefficient* to summarize the overall inclination of the system toward the search for specific rules, such coefficient will represent both the value k_2 in the bid equation and the probability of application of the genetic operator specification every time the genetic operators routine is called.

The model outlined so far can be used to study a variety of *coordination problems conditional on changing environmental states*. Basically, an organization has to respond to an exogenous *and changing* environment by implementing some collective action.

Suppose for instance that a firm can produce a certain number of product types, which are demanded by an exogenous market, and that the production process is divided into several parts, each of them being carried out by a different shop. The problem is therefore to detect correctly which product type is being demanded (state of the world) and to coordinate the actions of the shops so that the correct production process is implemented.

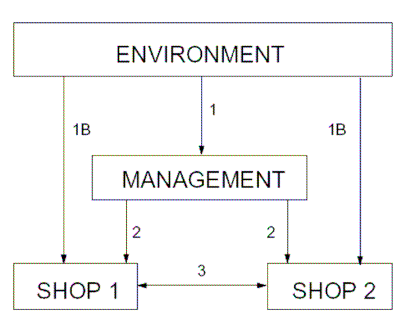
More specifically, suppose that there exist eight possible product types, called respectively "1", "2", . . . , "8". The firm's production possibilities set is represented by sequences of operations which can be of two types (A and B). Such sequences have all the same length and map into a product type, which is conventionally designated by the number of operations of type "1" which are utilized in its production. For example the product of type "8" is produced by all and only the production processes which contain eight operations of type "1". Each production process is divided into two parts (of the same length) which are carried out separately by each of two shops. The problem of the

firm is therefore to forecast the product type which will be demanded by the market and to implement the correct production process by coordinating the operations of the two shops. The payoff is the following: if the firm produces the correct product type it receives a payoff of 5 units; if it does not produce the correct output it receives a negative payoff, given by the distance of the actual product type from the required one (for example, if the market demands type "7" but the firm produces type "5", it will receive the payoff -2).

Suppose now that the all the decision-making units which the organization is made of are represented by agents whose knowledge of the state of the world evolves exactly in the way presented above.

A first bunch of simulations test the behaviour of a simple but quite general organizational structure (visualized in Figure 1), composed by a "management" and two shops. The management observes the environmental message (the last state of the world), interprets it according to its, evolving, "model of the world", and sends a message to the two shops.

Figure 1. Organizational informational flows (Marengo, 1992, 1996)



Each of the two shops can, in general, observe three kinds of signals and develop an interpretative model for each them. These signals are, respectively, the environmental signal (last observed state of the world), the message sent by the management (and based on its own interpretation of the environment), and the signal sent by the other shop (i.e. its last action). The latter two messages are coordinating devices, respectively a centralized and a decentralized one, which allow the shops to coordinate their action, whereas the former allows the two shops to form their own independent (from the management's) model of the world.

The weights with which these three types of messages enter the shops' decision processes define the organizational balance between differentiation and commonality of knowledge. Such weights are represented by the specificity coefficients which express the agent's search for a precise model which interprets the corresponding type of message. A high specificity coefficient for the shops' condition parts which classify messages coming from the environment (messages of type 1B in Figure 2) implies that shops are aiming at building a detailed individual model of the world. A low coefficient implies instead that shops do not pay much attention to the environment. When the coefficient is equal to zero we have an organization in which shops do not form any autonomous model of the world but rely entirely on the world's interpretation given by the management (messages of type 1 and 2).

A high specificity coefficient for the condition part which classifies messages coming from the management (messages of type 2 in Figure 2) implies that shops attribute great importance to the correct interpretation of the coordinating messages which are sent by

the management. A low coefficient implies instead that shops are not seeking careful coordination on the organizational collective knowledge. When the coefficient is equal to zero we have an organization without any form of centralized coordination, i.e. the management has no role.

Finally, a high specificity coefficient for the condition part which classifies messages coming from the other shop (messages of type 3 in Figure 2) implies that shops are attaching high importance to mutual, decentralized coordination. When the coefficient is equal to zero we have an organization without any form of decentralized coordination, i.e. no inter-shop communication.

Marengo (1992) and Marengo (1996) present a set of simulations, whose main results can be summarized as follows.

In stationary environments, i.e. when the state of the world does not change, agents can achieve coordination without building any model of the environment and resorting only to trial-and-error with selection. If instead they try to learn, i.e. to build such a model and constantly improving it, they need also to learn a model for the interpretation of coordinating messages: messages 1 and/or 1B are not sufficient, and messages 2 or 3 are also needed.

If the environment undergoes cyclical and predictable changes, high specificity coefficients on the shops' conditions which classify environmental messages (message 1B) are needed in order to exploit the environmental regularity.

Shops need to have a direct access to environmental information in order to develop the necessary decentralized learning.

Finally, if the environment undergoes frequent and unpredictable changes, the organization has to develop stable routines which give a "satisficing" average result in most conditions. In this case decentralized learning is detrimental, because the stability of such routine is continuously jeopardized by individual efforts to grasp the unpredictable environment. Shops should rely on the management's message.

All in all, in order to exploit a regularly changing environment a high amount of knowledge about the environment itself is required: the model must distinguish between the states of the world and connect them diachronically.

It is not surprising therefore that the most appropriate organization in such circumstances is the one which, by partly decentralizing the acquisition of knowledge about the environment, can achieve higher levels of sophistication in its model of the world, provided the coordination mechanisms - which are here centralized - are powerful enough to enable the organization to solve conflicts of representations. On the other hand, this very decentralization of the acquisition of knowledge can be a source of loss when it is more profitable for the organization to cling to a robust and stable set of routines. This situation requires strong coordination in order to make the entire organization implement coherently such a set of robust routines. Autonomous and decentralized experimentation can only disrupt such a coherence.

In our view, one ought to consider the foregoing models as a template for a largely unexplored family of exercises which takes seriously on board (i) informational imperfections; (ii) "boundedly rational" information processing; (iii) adaptive learning; and (iv) inter-organizational differences in information channels and decision rules. Indeed in these types of exercises, "balckboxing" is reduced to a minimum in so far as flows of information and decision acts are explicitly modelled. The downside sets precisely in the associated difficulty in identifying robust traits of whatever organizational arrangements which yield revealed "better" or "worse" performances.

3. Models of evolution in the space of "traits" and problem solving

In the last few years a new family of evolutionary models of organizations has developed inspired by biologist Stuart Kauffman's so-called "NK model" (Kauffman 1993). His model of selection and adaptation in complex environments represents evolving entities characterized by non-linear interactions among their elements. Kauffman developed the so called "NK-model" primarily to deal with the evolution of populations of biological entities described by a string of "genes", but its formal structure allows for various applications in other domains. The model, indeed, has lent itself to a growing number of applications, extensions and modifications within the realm of organization studies. In this section we will review some of them, well short of

a comprehensive survey, with the primary purpose to flag some of the main results and incumbent challenges.

3.1.1 *Organizational dynamics on complex selection landscapes*

With such a purpose let us build on one of the earliest attempt to apply the "NK" approach to organizational analysis, presented by Levinthal (1997), who assumes that an organization can be represented as a string of (binary) traits (e.g. policies, rules, routines, standard operating procedures, etc.) linked together by a thread of interdependencies which map into an equally stylized environment delivering performance feedbacks which select in favor/against such configuration of traits.. More formally, an organization is described by a string of N loci which refer to the set of elements ($i=1...N$) that make up the system. For each element i , there exist A_i possible states³. The set of all possible configurations (strings) of system's elements $A_1 \times A_2 \times \dots \times A_N$ is called the possibility space of a system.

Next, define a fitness function $F: A_1 \times A_2 \times \dots \times A_N \rightarrow [0,1]$ which assigns a (normalized) real number to each possible string as a measure of its relative performance.

The distribution of fitness values to all possible configuration defines the *fitness landscape* of the system. This landscape can be explored in search for the configuration with the maximum fitness value, moving from one configuration (a point in the fitness landscape) to another, by changing the value of one element. This "adaptive walk" ends when a configuration is reached which has not immediate neighbours with better fitness. Of course, if the "fitness contribution" of each trait were perfectly decomposable – as it is most obviously in e.g. standard (utterly "blackboxed") production function accounts – the usual "accounting" assumptions would be likely to apply: "more of x " contributes $f(x)$ to the fitness of the entity, etc. However if complementarities applies the map dramatically changes. Here the fundamental parameter, the K-value, refers to the number of "epistatic" relations among elements (the structure of the system). The existence of these relations imply that the contribution of one element to the overall

³ In most applications and in all those we consider in this paper, the number of states is reduced – for the sake of simplicity – to two: $A_i = \{0,1\}$.

fitness of the system is dependent both upon its own state and upon the state of K other elements. In the case, for example, of a system characterized by $K=3$, then the contribution of each element to the system's performance depends on the value assumed by other three elements to which it is interrelated. Two limit cases of complexity can be distinguished: *minimum complexity* when $K=0$, and *maximum complexity* when $K=N-1$. Consider for a example a system characterized by $N=3$, $A_i = [0,1]$ and $K=0$. Following Kauffman we draw the fitness values of elements for the two possible states randomly from a uniform distribution between 0 and 1. The fitness of the string as a whole is then defined as the mean value of the fitness values of elements:

$$F = \frac{\sum_{i=1}^N f_i}{N}$$

In Levinthal's simulations, populations of randomly generated structures (organizations) evolve on a fitness landscape, whereby the evolution is driven by variation selection and retention processes.

Variation, i.e. the generation of variety, is provided by two mechanisms:

- *local search*: one-feature mutation with retention of strings with higher fitness value.
- *Radical changes ("long jumps")*: mutation of many (possibly all) features with retention of string with higher fitness value.

Selection is obtained by simple birth and death process: organizations die with a probability inversely proportional to their relative fitness and are replaced by newly born ones. Some of these organization are randomly generated, owing possibly no resemblance to the existing ones, while others are replica of existing successful organization.

Information passes among generations by mean of two mechanisms:

- *retention*: successful existing organizations have a high probability of surviving. Their features tend therefore to survive.

- *replication*: some of the newly born organizations which replace bad performing ones, which are selected out, are copies of the most successful existing organizations. The features of the latter tend therefore to spread in the population.

Suppose that a large population of randomly generated organizations evolves according to the mechanisms of selection and information passing just mentioned and suppose that instead variation can be only local, i.e. that only one bit at a time can be mutated for every organization. Local adaptation and selection will reduce the heterogeneity of the population: bad performers will be selected out and replaced by copies of good performers. In the meantime good performers will climb with local mutations the fitness peaks they are located on.

However the final outcome of the evolution will crucially depend on the value of K , i.e. the complexity of the fitness landscape. With $K=0$ local adaptation will quickly take all the organizations to the only global optimum: thus selection and adaptation will completely wipe out the initial heterogeneity of the population and cause a fast convergence to unique optimal organizational form. For higher values of K the landscape will display an increasing number of local optima on which subset of organizations will converge according to their initial configuration. Selection and adaptation will reduce the heterogeneity but will never make it disappear.

This result, rather obvious in this framework, must not be overlooked, as it provides a simple and intuitive explanation of the persistence of heterogeneity among firms, a piece of evidence widely reported by the literature but at odds with neoclassical theory, according to which deviations from the only best practice should be only a transient property inevitably due to fade out as market selection forces operate. Note also that as K increases not only does the number of local optima increase, but also the size of the basin of attraction of each of them will shrink. It is possible therefore that none of the organizations is located in the basin of attraction of the global optimum and therefore no organization will ever find the globally optimum configuration.

In complex environment diversity of form can also emerge out of homogeneity. Levinthal (1997) that even if we start from a population of homogeneous organizations,

because of random local search they start mutating in different directions in the landscape. If $K > 0$ such initial random mutations will take organizations in the basins of attraction of different local optima. Selection and adaptation will only partially reduce such diversity.

If organizations can perform more radical changes (“long jumps”), i.e. mutate many (possibly all) features, also with large K heterogeneity tends - though very slowly - to disappear, as organization located on sub-optimal peaks can always perform -though with low probability - a radical mutation which allows them to jump on a higher fitness peak, until they reach the highest one of the global optimum. However if N is large enough such a possibility may have a very low probability and not make any real impact on the medium term evolution of the population.

Consider now the case of environmental changes, which can be modelled by re-drawing the fitness contributions of some features after the population has evolved and stabilized over the local optima. Suppose that such a change concerns only one feature and $K=0$, then if the fitness contribution of only one attribute is modified, the global optimum will either remain where it was or move to a point which is at most one mutation away. Thus, if the population has already evolved and located on the global optimum, it can easily and quickly adapt and move to the new global optimum. Simulations show that all incumbent organizations survive to such an environmental change.

If instead the complexity of the landscape is high ($K \gg 0$), even the modification of the fitness contribution of just one attribute can cause a large alteration of its shape. In high dimensional landscapes with large N local optima can move far away. This implies that a population which has settled on the local optima of the initial landscape will find it much more difficult to adapt to the change. Mortality of incumbents will rapidly rise as K increases.

If the environment changes more radically, i.e. the fitness contributions of many (possibly all) the attributes are re-drawn, we get a different picture. As we have already argued, in a “simple” landscape with $K = 0$ all organizations quickly converge to the same configuration, which correspond to the unique global optimum and diversity dies out. If a dramatic environmental shock happens for which the global optimum moves

far away from its initial position, the entire population will find itself in a low fitness area of the landscape and incumbent organizations will be outperformed by newly created ones with random configuration.

If, on the contrary, K is high the population remains distributed over a large number of local optima and there is a high likelihood that a subset of the population will find itself in or close to a high fitness portion of the landscape after the environmental shock has occurred. Preserving diversity helps the population adapt to dramatic environmental changes.

Levinthal's analysis has been expanded and broadened by a few papers which have further studied the relationship between organizational design and environmental complexity and turbulence. Among them, interesting results have been obtained by Rivkin and Siggelkow (cf. Rivkin and Siggelkow (2002), (2003), Siggelkow and Rivkin (2005)). Differently from Levinthal (1997) they introduce a representation of an organizational structure in a NK-type model. Decisions over the N policies (bits of the string) are allocated among different departments and a superordinate CEO takes the function of coordinating departmental decisions.

More in detail, each department controls a given number of policies and is engaged in increasing the fitness contribution of such policies (climbing the departmental "subspace", i.e. the landscape generated by only those policies). As – in general – any policy change in one department changes also the other departments' fitness values, each department also attaches some weight to fitness changes of other departments. This weight, ranging from 0 to 1, is a model parameter which stands for the degree of inter-department coordination.

Finally, the organization has a CEO in charge endowed with the power of taking the final decisions by selecting departments' proposals. For this purpose, the CEO asks each department i to suggest its most preferred alternatives and selects those combination of departments' proposals which deliver the highest organizational fitness. The parameters d_i measure the degree of CEO discretion: at one extreme, if d_i is equal to one for all departments, then the CEO can simply automatically approve each department's most preferred alternative, without any *de facto* selection power. At the

other extreme, if d_i is equal to the number of all envisageable alternatives for all departments, then the CEO has a *de facto* full discretionary control over all policies.

This interplay between departments and CEO creates what the authors call a set of “sticking points”, i.e. organizational configurations to which no alternative exists which can go through the approval of all subjects involved. Sticking points do not necessarily correspond to organizational local optima, as on one side cross-vetoes of departments and CEO can prevent an improvement which would increase the fitness of the organization and, on the other side, a department can, in some circumstances, implement a change which is beneficial for itself but not for the entire organization and therefore unlock the organization from local optima.

Divergence between the set of local optima and the set of sticking points is larger when the following conditions are met:

1. decisions are allocated among a larger number of departments;
2. interdependencies among policies allocated to different departments are stronger;
3. the weight which each department attributes to other departments’ fitness is lower;
4. the larger the number of proposal the CEO receives from departments if the latter give high weight to others’ fitness.

3.1.2 *Cognitive and experiential search*

Gavetti and Levinthal (2000) deepen the analysis of search processes by looking at the relations between forward-looking and backward-looking search and their effects on the performance of the system. The two search processes refer to two logics of action derived by Herbert Simon’s (1955) definition of bounded rationality. On the one side there is the cognitive and forward-looking choice based on off-line evaluation of a broad set of alternatives, even very distant from current behavior; on the other side there is experiential choice based on on-line evaluation of a limited set of alternatives which are close to current behavior.

In Gavetti and Levithal's model, the organization chooses a policy on the basis of a simplified cognition of its environment. This choice results in the identification of a set of possible actions (a template) which cannot be directly translated into actions. In this context, existing practices function as defaults for elements not specified by the cognitive representation and they allow for the identification of a specific course of action. Thus, it may happen that actors with the same cognition may engage in different behaviors.

This concept are translated into a NK-based model in which organization's limited cognition corresponds to a simplified representation of the fitness landscape which is assumed to be of lower dimensionality than the actual landscape ($N_1 < N$), but nonetheless grounded in it. This idea is captured by assuming that for each point of the cognitive representation (of the perceived landscape) there are 2^{N-N_1} points in the actual fitness landscape. The fitness value assigned to each point of the cognitive corresponds to the average fitness values of these 2^{N-N_1} point. Thus for each point in the perceived landscape there are 2^{N-N_1} arrays in the actual landscape.

Organization which choose according to its cognitive representation explores regions, and not single points, of the landscape. And the width of these regions depends on the crudeness of the representation. When both cognitive and experiential search are at work, organization identifies a pick in its perceived N_1 -dimensional landscape (by cognitive or off-line search) and then explores the remaining $N-N_1$ alternatives through a local (or on-line) search based on one bit-mutations. The role of experiential search becomes more and more important as the crudeness of the cognitive representation increases. What is important to notice is the role of the initial cognitive search in identifying, on average, superior basins of attractions. Indeed, the global pick of the representation generally corresponds to an attractive region of the actual landscape. Initial off-line search then helps in finding a good position from which the local search can start.

Gavetti and Levinthal show that in a context of competitive ecology in which low performance organizations are selected out and are replaced by new born ones, organizations which adopt a joint cognitive and experiential search dominated the

population. This becomes particularly clear with rugged actual landscapes, in which organizations which use purely experiential search are trapped into local optima.

In this framework what are the effects of adaptation through changes in the cognitive representation? Gavetti and Levinthal (2002) consider these effects both in the case of purely cognitive search and in that of joint cognitive and experiential search, also with changes in the actual fitness landscape. In the case of pure cognitive representation the organization choose an alternative on the basis of its understanding of the payoffs as characterized by a set of N_1 attributes. In this case the effects of changes in the representation depends on the complexity of the landscape (the value of K). If K is high these changes may produce good performances, as they can compensate for a poor representation of the landscape. Of course these effects depend on the nature of the change: in the case of purely off-line search, organizations which perform better are those who adopt a semi-intelligent change process, changing their representation with a probability which is inversely-proportional to their fitness and imitating leading organizations in the population. But if one considers organizations which use joint off-line and on-line search, the shift to a new representation could destroy the accumulated experience. In this context the best performances is obtained in the case of no changes at all, while no differences exist between organizations which adopt semi-intelligent changes and organizations which adopt a random change procedure.

Changes in the representation can enhance organization's performance when the landscape itself changes as the new representation may more effectively identify new superior basin of attraction, and this can compensate for the loss of experiential wisdom. Gavetti and Levinthal (2002) shed light on the role of cognitive search in conditioning experiential leaning by constraining the local search to the most promising regions of the landscapes. The analysis of the interplay between the two logics of action indifferent contexts represents a significant progress with respect to Levinthal's (1997) model in which organizational search process is reduced either to "one-bit mutation" search or to totally random "big jumps" .

A further step in the direction of opening up the “organizational problem solving black box” entails an explicit representation of organizational problem solving procedures, their emergence and their dynamics.

3.1.3 Problem solving organization and the division of labor

Following Simon (1981), Marengo and Dosi (2005)⁴ focus on strategies for the reduction of problem complexity through a division of problem solving labor, that results in the decomposition of large and complex problems into smaller sub-problems which can be solved independently. They argue that the deriving process of division of labour is a major and long neglected driving force in explaining economic organization. In particular, traditional organizational economics has concentrated upon the governance of transaction and contractual relations between given “technologically separable” units, but does not tackle the analysis of where such technologically separable units come from nor, more importantly, of whether organizational structures have some

This issue is relevant both because it is clear that most processes of division of labour take place within organizations and, relatedly, because empirical evidence shows that most of the times technologies are born in a highly integrated fashion, then they possibly undergo vertical disintegration (and sometimes a subsequent re-integration) along the lines defined by the within-firm division of labour. In other words, we could say that “in the origin there were organizations” and then markets develop along the lines defined by the division of labour within firms, rather than the other way round as postulated by transaction costs economics.

Marengo and Dosi (2005) put forward a problem-solving approach to economic organization where different organizational structures (with varying degrees of vertical integration) are compared in terms of their dynamic problem-solving properties determined by division of labour and task decomposition. The basic assumption is that

⁴ See also Marengo et al. 2000.

solving a given problem requires the coordination of N atomic “elements” or “actions” or “pieces of knowledge”, which we can generically call components, each of which can assume some number of alternative states. The one-bit mutation algorithm at the basis of the NK model, can be conceived as a particular case in which the problem is fully decomposed and the search process is fully decentralized: each sub-problem consist of a single component (bit). As showed by Kaufmann (1993), this algorithm is very quick, but it allows to reach only the local optimum whose basin of attraction contain the initial configuration. On the opposite there is the case of no decomposition at all, or total centralization, corresponding to a strategy in which all the components (bits) are simultaneously mutated. In this case the global optimum can be reached by exploring all the possible configurations. In between there are all the other possible divisions of labor strategies.

Note that the effectiveness of the decomposition, in terms of system optimization, is strongly affected by the existence of interdependences among the components of the problem: separating interdependent components and then solving each sub-problem independently will prevent overall optimization. Note also that, as pointed out by Simon, because of the opaqueness of the interrelations between component, optimal decomposition – a division of labor that separate into sub-problems only the components that are independent from each other - cannot be achieved by bounded rational agents that normally are bound to adopt near-decompositions, trying to put together within the same sub-problem only those components whose interdependences are important for the performance of the system.

A last, but central aspect that must be considered is the fact the that the search space in not given exogenously, but is constructed by individuals that possess subjective representations of the structure of the problem. The point is that the distance between the real structure of the problem (its real decomposition) and the subjective representation that individuals have of it has a dramatic effect on the problem solving outcome.

More formally, one can characterize a problem by the following elements:

The set of components: $C = \{c_1, c_2, \dots, c_N\}$, where each component can take one out of a final number of states. Normally, without loss of generality, a binary set of components is assumed for simplicity: $c_i \in \{0, 1\} \forall i$.

A configuration, that is a possible solution to the problem, is a string $x^i = c_1^i c_2^i \dots c_N^i$

The set of configurations: $X = \{x^1, x^2, \dots, x^{2^N}\}$

An ordering over the set of possible configurations: we write $x^i \geq x^j$ (or $x^i > x^j$) whenever x^i is weakly (or strictly) preferred to x^j .

A problem is defined by the pair (X, \geq) .

As the size of the set of configurations is exponential in the number of components, whenever the latter is large, the state space of the problem becomes much too vast to be extensively searched by agents with bounded computational capabilities. One way of reducing its size is to decompose⁵ it into sub-spaces.

Let $I = \{1, 2, \dots, N\}$ be the set of indexes and let a block⁶ $d_i \subseteq I$ be a non-empty subset of it, we call the size of block d_i its cardinality $|d_i|$. Let us define a decomposition of the problem (X, \geq) as a set of blocks:

$$D = \{d_1, d_2, \dots, d_k\} \text{ such that } \bigcup_{i=1}^k d_i = I$$

Note that a decomposition does not necessarily have to be a partition.

Given a configuration x^i and a block d_j , the block-configuration $x^i(d_j)$ is the substring of length $|d_j|$ containing the components of configuration x^i belonging to block d_j :

$$x^i(d_j) = x_{j_1}^i x_{j_2}^i \dots x_{j_{|d_j|}}^i \quad \forall j_h \in d_j$$

We also use the notation $x^i(d_j)$ to indicate the substring of length $N - |d_j|$ containing the components of configuration x^i not belonging to block d_j .

⁵ A decomposition can be considered as a particular case of search heuristics: search heuristics are, in fact, ways of reducing the number of configurations to be considered in a search process.

⁶ Blocks in our model can be considered as a formalization of the notion of modules used by the flourishing literature on modularity in technologies and organizations (Baldwin and Clark, 2000) and decomposition schemes are a formalization of the notion of system architecture which defines the set of modules in which a technological system or an organization are decomposed.

Two block-configurations can be united into a larger block-configuration by means of the \vee operator so defined: $x(d_j) \vee y(d_h) = z(d_j \cup d_h)$ where $z_k = x_k$ if $k \in d_j$ and $z_k = y_k$ if $k \in d_h$

Let us then define the size of a decomposition as the size of its largest defining block:

$$|D| = \max\{|d_1|, |d_2|, \dots, |d_k|\}$$

Coordination among blocks in a decomposition may either take place through market-like mechanisms or via other organizational arrangements (e.g. hierarchies). Dynamically, when a new configuration appears, it is tested against the existing one according to its relative performance. The two configurations are compared in terms of their ranks and the superior one is selected, while the other one is discarded.

More precisely, let us assume that the current configuration is x^i and take block d_h with its current block-configuration $x^i(d_h)$. Let us now consider a new configuration $x^j(d_h)$ for the same block, if:

$$x^j(d_h) \vee x^i(d_h) \geq x^i(d_h) \vee x^i(d_h)$$

then $x^j(d_h)$ is selected and the new configuration $x^j(d_h) \vee x^i(d_h)$ is kept in place of x^i , otherwise $x^j(d_h)$ is discarded and x^i is kept.

It might help to think in terms of a given division of labor structure (the decomposition scheme) within firms, whereby individual workers and organizational sub-units specialize in various segments of the production process (a single block). Decompositions, however, sometimes determine also the boundaries across independent organizations specialized in different segments of the whole production sequence.

Note that, dynamically, different inter-organizational decompositions entail different degrees of decentralization of the search process. The finer the inter-organizational decompositions, the smaller the portion of the search space which is being explored by local variational mechanisms and tested by market selection. Thus there is inevitably a trade-off: finer decompositions and more decentralization make search and adaptation faster (if the decomposition is the finest, search time is linear in N), but on the other hand, they explore smaller and smaller portions of the search space, thus decreasing the likelihood that optimal (or even good) solutions are ever generated and tested.

A decomposition is a sort of template which determines how new configurations are generated and can be tested afterward by the selection mechanism. In large search spaces in which only a very small subset of all possible configurations can be generated and undergo testing, the procedure employed to generate such new configurations plays a key role in defining the set of attainable final configurations.

We will assume that boundedly rational agents can only search locally in directions which are given by the decomposition: new configurations are generated and tested in the neighborhood of the given one, where neighbors are new configurations obtained by changing some (possibly all) components within a given block.

Given a decomposition $D = \{d_1, d_2, \dots, d_k\}$, we say that a configuration x^i is a preferred neighbor or simply a neighbor of configuration x^j with respect to a block $d_h \in D$ if the following three conditions hold:

1. $x^i \geq x^j$
2. $x_k^i = x_k^j \quad \forall k \notin d_h$
3. $x^i \neq x^j$

Conditions 2 and 3 require that the two configurations differ only by components which belong to block d_h . According to the definition, a neighbor can be reached from a given configuration through the operation of a single decentralized coordination mechanism.

We call $H_i(x, d_i)$ the set of neighbors of a configuration x for block d_i .

The set of best neighbors $B_i(x, d_i) \subseteq H_i(x, d_i)$ of a configuration x for block d_i is the set of the most preferred configurations in the set of neighbors:

$$B_i(x, d_i) = \{y \in H_i(x, d_i) \text{ such that } y \geq z \quad \forall z \in H_i(x, d_i)\}$$

By extension from single blocks to entire decompositions, we can give the following definition of the set of neighbors for a decomposition as:

$$H(x, D) = \bigcup_{i=1}^k H_i(x, d_i)$$

A configuration is a local optimum for the decomposition D if there does not exist a configuration y such that $y \in H(x, D)$ and $y > x$.

A search path or, for short, a path $P(x^i, D)$ from a configuration x^i and for a decomposition D is a sequence, starting from x^i , of neighbors:

$$P(x^i, D) = \{x^i, x^{i+1}, x^{i+2}, \dots\} \text{ with } x^{i+m+1} \in H(x^{i+m}, D)$$

A configuration x^j is reachable from another configuration x^i and for decomposition D if there exists a path $P(x^i, D)$ such that $x^j \in P(x^i, D)$.

Suppose configuration x^j is a local optimum for decomposition D ; we call the basin of attraction of x^j for decomposition D the set of all configurations from which x^j is reachable:

$$\Psi(x^j, D) = \{y, \text{ such that } \exists P(y, D) \text{ with } x^j \in P(y, D)\}$$

Now let x^0 be the global optimum and let $Z \subseteq X$ with $x^0 \in Z$. We say that the problem (X, \geq) is locally decomposable in Z by decomposition D if $Z \subseteq \Psi(x^0, D)$. If $Z = X$, we say that the problem is globally decomposable by decomposition D .

We can soften the perfect decomposability requirement into one of near-decomposability: we no longer require the problem to be decomposed into completely separated sub-problems, i.e. sub-problems which fully contain all interdependencies, but we might be happy to find sub-problems which contain the most relevant interdependencies, while less relevant ones can persist across sub-problems. In this way, optimizing each sub-problem independently will not necessarily lead to the global optimum, but to a “good” solution. In other words, we construct near-decompositions which give a precise measure of the trade-off between decentralization and optimality: higher degrees of decentralization, while generally displaying a higher adaptation speed, are likely to be obtained at the expense of the asymptotic optimality of the solutions which can be reached.

Let us arrange all the configurations in X by descending rank $X = \{x^0, x^1, x^2, \dots\}$ where $x^i \geq x^{i+1}$, and let $X_\mu = \{x^0, x^1, \dots, x^{\mu-1}\}$ be the ordered set of the best μ configurations.

We say that X_μ is reachable from a configuration $y \notin X_\mu$ and for decomposition D if there exists a configuration $x^i \in X_\mu$ such that $x^i \in P(y, D)$.

We call the basin of attraction $\Psi(X_\mu, D)$ of X_μ for decomposition D the set of all configurations from which X_μ is reachable. If $\Psi(X_\mu, D) = X$ we say that D is a μ -decomposition for the problem. μ -decompositions of minimum size can be found

algorithmically with a straightforward generalization of the above algorithm which computes minimum size optimal decompositions.

It is straightforward to show (Marengo and Dosi 2005) that as μ increases we can generally find finer near-decompositions. This shows that the organizational structure sets a balance in the trade-off between search and adaptation speed and optimality. It is easy to argue that in complex problem environments, characterized by strong and diffused interdependencies, such a trade-off will tend to produce organizational structures which are more decomposed and decentralized than what would be optimal given the interdependencies of the problem space.

Different organizational forms implement different decomposition heuristics and might be characterized by different representations of the problem and therefore present different properties in terms of the effectiveness and efficiency of the derived search processes (see Marengo, Pasquali and Valente (2005) for a theoretical discussion of the topic). In particular a trade-off exists between complexity and optimality: a finer decomposition makes search faster, but the exploration of smaller portion of the search space reduces the likelihood to generate and then select an optimal solution. The application of these ideas to organizational design leads to the comparison, in terms of relative performance, between not decomposed tasks (organization-embodied) and decomposed tasks (coordinated via market-like mechanism or via simple organizations structured as sets of perfectly independent tasks). One of the main conclusions is that is that the advantages of decentralization (faster adaptation) usually imply a cost in terms of sub-optimality (impossibility to reach global optima). This casts strong doubts on the efficacy of market selection processes as substitutes for individual optimization: selection is not able to select out sub-optimal features nor to select for optimal ones if both are somehow complementary to each-other in actual organizations and technologies.

Modeling the coupling mechanisms between capabilities and governance

Marengo and Dosi (2005), as well as most of contributions of this genre, while concentrating on the problem-solving features of organizational dynamics, censor any

incentive compatibility issue. An attitude that, as noted above, is quite typical within the capability-based framework.

There is nothing, however, preventing this type of analysis to go beyond the exclusive focus on firms as *loci* of coordination and as *loci* of creation, implementation, storage and diffusion of productive knowledge⁷ and explicitly take on board the issues of incentive governance and control discussed qualitatively in Coriat and Dosi (1998). Attempts in this direction are formal analyses by Dosi, Levinthal and Marengo (2002; 2003) which incorporate issues of conflict of interests, power and control over agents' decisions within the analytical framework of Marengo and Dosi (2005) and Marengo et al (2000) and to discuss the interaction between problem representation and incentive mechanisms. In particular, the double role of problem representation is stressed: on the one hand it defines the "cognitive" structure of the problem and the consequent decomposition which is adopted (definition of teams as subsets or blocks of components); on the other hand, it has important consequences for a reward mechanism based on the distinction between organization's (system) and team's (block) performance as it defines what organization conceives as a team.

The analysis starts by considering the conflicts of interest among problem solving teams generated by the adoption team-level incentive mechanisms. While under a global reward an alternative (a particular configuration of sub-problem's components) is selected if it improves the overall organization's performance, with a team-level reward mechanism a would-be alternative is accepted if it enhances the performance of the unit even if it degrades the overall organization's performance. It can be shown that if the organization's representation of the problem is not correct (it does not correspond to the right structure of the problem in terms of interrelations among components) the adoption of a global reward allows the organization to reach a global optimum. But what is more interesting is that, even if the representation of the problem is not correct,

⁷ A more complete "co-evolutionary" picture is discussed by Dosi (1995). Organizations are assumed to be characterized by six correlated dimensions: the distribution of formal authority; the distribution of power; the incentive structure; the structure of information flows; the distribution of knowledge and competence. In this context organization dynamics can be conceived as a process of adaptation and selection according to multiple, and possibly conflicting, objectives.

the adoption of a team-level reward structure tends, in the long run, to produce performances that are similar to the global- reward one. Thus, goal conflicts prevent organization to remain absorbed in local optima and act as substitute for a correct representation of the problem (Dosi, Levinthal and Marengo: 2002).

Power is introduced by allowing one team (a block in the decomposition) to stop the mutation of any other blocks that decreases its own performance (veto power). The evidence suggests that, under specific conditions, the adoption of such a mechanism lead to good solutions. In particular the a team reward scheme with veto power is superior to the global reward structure when the organizational representation of the problem is based on a finer decomposition than the real one and the latter is not too complex. This is due to the fact that veto power interrupts the cycling among possible solutions generated by a team-based reward structure preserving the advantages in terms greater search effort which are typical of this reward mechanism.

A principal-agent-like model of interaction is reproduced considering the case of control over the decisions of other organizational members by a principal, the residual claimant of the total payoff, who can “order” others to keep performing a given action or to switch to a different one. This activity is considered to have a cost which depends on the span of control, i.e. the dimension of each sub-unit, and it is higher when the principal wants to induce a change in agent’s action than when he wants to elicit the same behaviour (the principal’s profit is defined as the total output of the organization minus the “elicitation cost”). When actions are interdependent, the control function, as any other problem-solving activity, cannot be entirely decomposed. Thus, the interaction between a cognitive dimension and a control dimension has to be considered. The effects on total performance and the principal’s profit are analyzed considering four different cases: right, almost right, wrong and minimal (one-component units) perceived decomposition by agents, with reference to different decompositions of the underlying problem and the “correctness” of the decomposition itself.

Obviously if the organizational decomposition is the “true” one, perfectly knowledgeable agents not facing any incentive compatibility problem would make costly control redundant. However, interestingly, when the organization has a wrong

representation of the problem space (and in particular underestimates the span of interdependencies), agents subject to costly control may generate a better performance than the one produced by perfectly ‘cooperative’ agents.

Finally Dosi, Levinthal and Marengo (2002) analyze more explicitly the double role of problem representation. The work examines, in particular, by means of a simulation model, the relations between cognitive decompositions and operational decompositions. The former establish search heuristics and targets, whereas the latter implement search processes driven by those targets. The exercise shows that if cognitive decompositions are correct then it is efficient to have maximum division of labor at the operational level, as this increases speed and accuracy of adaptation to targets. On the contrary, if cognitive decompositions do not correspond to the “true” ones, coarser division of labor at the operational level ensures less accurate but prompt adaptations to the imperfectly set target.

4. Conclusions

Parallel to the qualitative analyses of organizations as structured bundles of problem-solving capabilities (for a critical review of the literature cf. Dosi, Faillo and Marengo, 2006), a growing number of contributions have begun to offer formal accounts of such organizational properties and their dynamics. The formal instruments are diverse, and include NK models representing organizations as ensembles of interrelated “traits” mapping into some overall environmental fitness of the firm; classifiers system representations of the problem-solving procedures triggered by diverse internal or environmental states; decomposition schemes of Simonian ascendancy allowing the analysis of the performance properties of different “representations” in the problem-solving space and different patterns of division of cognitive and operational labour.

The results begin to highlight important comparative properties regarding, among other, the impact on problem-solving efficiency and learning of different forms of hierarchical governance, the dangers of lock-in associated with specific forms of adaptive learning,

the relative role of “online” vs. “offline” learning, the impact of the “cognitive maps” which organizations embody, the possible trade-offs between accuracy and speed of convergence associated with different “decomposition schemes”.

In a nutshell, one has finally begun to develop formal instruments allowing exercises of comparative institutional analysis (cf. Aoki, 2001), focusing on the distinct properties of different forms of organization and accumulation of knowledge. It is a work which is only at its exciting start.

References

- Aoki, M. (2001), *Toward a comparative institutional analysis*, Cambridge MA, MIT Press.
- Augier, M. and March J.G. (eds.) (2002) *The Economics of Choice, Change and Organization: Essays in Memory of Richard M. Cyert*. Cheltenham UK and Northampton, MA, USA, Edward Elgar.
- Becker, M. (2004), “Organizational Routines: A Review of the Literature”, *Industrial and Corporate Change*, 13(4), pp. 643-678.
- Chandler, A., Hagström, P. and Sölvell, Ö. (eds.) (1998). *The Dynamic Firm*, Oxford/New York, Oxford University Press.
- Chandler, A. (1977). *The Visible Hand: the Managerial Revolution in American Business*, Cambridge, Mass., Harvard University Press.
- Cohen, M.D., March J.G. and Olsen, J.P. (1972), “A Garbage Can Model of Organizational Choice”, *Administrative Science Quarterly*, 17(1), 1-25.
- Cohen, M. D., Burkhart, R., Dosi, G., Egidì, M., Marengo, L., Warglien, M. and Sydney Winter (1996) “Routines and Other Recurring Action Patterns of Organizations: Contemporary Research Issues”.

- Coriat, B. and Dosi, G. (1998). 'Learning How to Govern and Learning How to Solve Problems: On the Co-Evolution of Competences, Conflicts and Organizational Routines', in Chandler et al. (1998).
- Cyert, R.M. and March J.G. (1963). *A Behavioral Theory of the Firm*, Cambridge: Blackwell.
- Dosi, G. (1995) "Hierarchies, Markets and Power: Some Foundational Issues on the Nature of Contemporary Economic Organization", *Industrial and Corporate Change*, 4(1).
- Dosi G., Faillo M. and Marengo L. (2006) "Problem-Solving and Governance in the Capability-based view of the Firm: The Roles and Theoretical Representations of Organizational Routines", forthcoming in M. Becker, *Handbook of Organizational Routines*, Cheltenham UK and Northampton, MA, USA, Edward Elgar..
- Dosi G., Levinthal D. and Marengo L. (2002) "The Uneasy Organizational Matching Between Distribution of Knowledge, Division of Labor and Incentive Governance", *LEM Working Papers*, Pisa, *Sant'Anna School of Advanced Studies*.
- Dosi G., Levinthal D. and Marengo L. (2003) "Bridging Contested Terrain. Linking Incentive-Based and Learning Perspectives on Organizational Evolution", *Industrial and Corporate Change*, 12 (1), pp. 413-436
- Dosi G. and Marengo L. (1995) 'Toward a Theory of Organizational Competencies', in R.W. England (ed.), *Evolutionary Concepts in Contemporary Economics*, Ann Arbor, Michigan University Press, pp. 157-178.

- Dosi G. and Marengo L. (2000) “On the Tangled Discourse between Transaction Cost Economics and Competence-Based Views of the Firm” in Foss, N. and Mahnke V. (eds.), *Competence, Governance and Entrepreneurship*, Oxford, Oxford University Press.
- Dosi, G., Nelson R. and Winter S.G.(eds.) (2000). *The Nature and Dynamic of Organizational Capabilities*. Oxford/New York, Oxford University Press.
- Dosi, G. and Winter, S.G. “Interpreting Economic Change: Evolution, Structures and Games” in Augier and March (2003).
- Ethiraj, S and D. Levinthal (2004). Modularity and Innovation in Complex Systems. *Management Science*, 50, 159-173.
- Gavetti, G. and Levinthal, D. (2000) “Looking Forward and Looking Backward: Cognitive and Experimental Search”, *Administrative Science Quarterly*, 45 pp. 113–137.
- Jones, Oswald and Martin Craven (2001), “Beyond the routine: innovation management and the Teaching Company Scheme”, *Techovation*, 21, pp. 267-279.
- Kauffman, S.A. (1993) *The Origins of Order*, Oxford: Oxford University Press.
- Kreps, D.M. (1990). ‘Corporate culture and economic theory’ in Alt, J.E., Sheplisle K.A.(eds.) (1990) *Perspectives on Positive Political Economy*, Cambridge University Press.
- Lazaric, N. and Denis, B. (2004), “Routinization and Memorization of Tasks inside a Worskshop: An Illustration through a Case Study”, forthcoming in *Industrial and Corporate Change*.
- Levinthal, D. (1997), “Adaptation on rugged landscapes”, *Management Science*, 43, 934-950.

- Levinthal, D. (1998), "The Slow Pace of Rapid Technological Change: Gradualism and Punctuation in Technological Change", *Industrial and Corporate Change*, 7 (2), pp. 217-247
- March, J.G. and Simon H.A. (1993), *Organizations*, New York, Wiley, 2nd Edition.
- Marengo, L. (1992), "Coordination and Organizational Learning in the Firm", *Journal of Evolutionary Economics*, 2, 313-326.
- Marengo, L. (1996), "Structure, Competence and Learning in an Adaptive Model of the Firm" in G. Dosi and F. Malerba (eds.) *Organization and Strategy in the Evolution of the Enterprise*, London, MacMillan, 124-154.
- Marengo, L., Dosi, G., Legrenzi, P. and Pasquali, C. (2000). 'The Structure of Problem-solving Knowledge and the Structure of Organizations', *Industrial and Corporate Change*, 9:757-788.
- Marengo, L., and G. Dosi (2005): "Decentralization and market mechanisms in collective problem-solving," *Journal of Economic Behavior and Organization*, forthcoming.
- Marengo, L. Pasquali C. and Valente M. (2005) 'Decomposability and Modularity of Economic Interactions' in Callebaut, W. and Rasskin-Gutman, D. (eds) *Modularity: Understanding the development and evolution of complex natural systems*, Cambridge MA, MIT Press, pp- 835-97.
- Narduzzo, A., E. Rocco and M. Warglien (2000)(A). 'Talking about Routines in the Field', in Dosi, G., Nelson R. and Winter S.G.(eds.) (2000), *The Nature and Dynamics of Organizational Capabilities*, Oxford: Oxford University Press, pp. 27-50.

- Nelson R. (1991), "Why do firms differ and how does it matter?", *Strategic Management Journal*, 12, 61-74
- Nelson, R. (1994), "The Role of Firm Differences in an Evolutionary Theory of Technical Advance" in Magnusson (ed.), *Evolutionary and Neo-Schumpeterian Approaches to Economics*, Boston: Kluwer.
- Nelson, R. and Winter, S. (1982) *An Evolutionary Theory of Economic Change*, Belknap Harvard Press.
- Page, S. E. (1996) "Two measures of difficulty", *Economic Theory*, 8 pp.321–346.
- Pavitt K. (2002), "Innovating Routines in the Business Firm: What Corporate Task Should They Be Accomplish?", *Industrial and Corporate Change*, 11(1), pp.117-133.
- Rivkin, J. W: and Nicolaj Siggelkow. 2002. "Organizational Sticking Points on NK Landscapes." *Complexity*, 7 (5), pp. 31-43.
- Rivkin, J. W. and N. Siggelkow.(2003). "Balancing Search and Stability: Interdependencies Among Elements of Organizational Design." *Management Science*, 49, pp. 290-311.
- Siggelkow, N. and J. W. Rivkin,.(2005). "Speed and search: Designing organizations for turbulence and complexity" *Organization Science*, 16, pp. 101-122.
- Simon, H. A (1962), "The Architecture of Complexity", *Proceedings of the American Philosophical Society*, 106, 467-482.
- Simon, H.A. (1981), *The Science of Artificial*, 2nd Ed., Cambridge, MA: MIT Press.

- Schumpeter, J.A. (1950), *Capitalism, Socialism and Democracy*, New York: Harper.
- Teece, D.J., Pisano G. and Shuen, A. (1997). "Dynamic Capabilities and Strategic Management", *Strategic Management Journal*, 18: 509-33; a revised version is published in Dosi, Nelson and Winter (2000).
- Winter S. 1987, Knowledge and Competence as Strategic Assets, In D.J. Teece (ed.), *The competitive Challenge: Strategies for Industrial Innovation and Renewal*: Cambridge, M.A.: Ballinger. Pp.159-184.
- Winter, S.G. and Szulanski, G. (1998), "Replication as Strategy", *Organizational Science*, 12(6): 730-743.
- Winter, S.G. and Szulanski, G. (2002), "Replication of Organizational Routines: Conceptualizing the Exploitation of Knowledge Assets", in N. Bontis and C.W. Choo (eds.) *The Strategic Management of Intellectual Capital and Organizational Knowledge: A Collection of Readings*, 207-221.
- Winter, S.G. (2005) 'Toward and Evolutionary Theory of Production' in K. Dopfer (ed.) *The Evolutionary Foundations of Economics*, Cambridge University Press.
- Zander, U. & Kogut, B. 1995. "Knowledge and the speed of transfer and imitation of organizational capabilities: an empirical test". *Organization Science*, 6(1): 76-92.
- Zollo, M. and Winter, S. G. (2002), "Deliberate Learning and the Evolution of Dynamic Capability", *Organization Science*, 13(3), pp. 339-351.