LEM
Working Paper Series

Divison of Labor, Organizational Coordination and Market Mechanism in Collective Problem-Solving

Luigi Marengo [†]

Giovanni Dosi [‡]

[†]*Univeristà di Teramo, Teramo, Italy*
[‡]*Scuola Superiore Sant'Anna, Pisa, Italy*

Division of Labor, Organizational Coordination and Market Mechanisms in Collective Problem-Solving[*]

Luigi MARENGO[†]

University of Teramo

Teramo, Italy

Giovanni DOSI [‡]

S. Anna School of Advanced Studies

Pisa, Italy

Abstract

This paper builds upon a view of economic organizations as problem-solving arrangements and presents a simple model of adaptive problem-solving driven by trial-and-error learning and collective selection. Institutional structures, and in particular their degree of decentralization, determines which solutions are tried out and undergo selection. It is shown that if the design problem at hand is "complex" (in term of interdependencies between the elements of the system), then a decentralized institutional structure is very unlikely to ever generate optimal solutions and therefore no selection process can ever select them. We also show that nearly-decomposable structures have in general a selective advantage in terms of speed in reaching good locally optimal solutions.

[†]**Corresponding author**: DSGSS, Università di Teramo, Loc. Colleparco, 64100 Teramo, Italy, e-mail: marengo@unite.it

[‡]Scuola Superiore S. Anna, P.za Martiri della Libertà 33, 56127 Pisa, Italy, e-mail: gdosi@sssup.it

1 Introduction

One way to describe any economy or, for that matter, any economic organization is as a huge ensemble of partly interrelated tasks and processes which, combined in certain ways, produce "well constructed" goods and services. It is a perspective which dates back at least to Adam Smith who identified a major driver of productivity growth in the progressive division of tasks themselves and the associated specialization among workers. More recently, several seminal works of Herbert Simon have explored the general structure of problem-solving activities of which the activities of technological search and economic production are just subsets (Simon 1969). From different angles, several investigations from the perspective of team theory have addressed the symmetric problem concerning coordination amongst multiple interrelated tasks (cf. Marschak and Radner (1972), Radner (2000), Becker and Murphy (1992) among others). And, finally, a growing literature has focussed on the "cognitive" characteristics of organizations (cf., among others, Richardson (1972), Langlois and Robertson (1995), Loasby (1998), Teece et al. (1994), Dosi et al. (2000)).

The contribution which follows has its roots in the foregoing perspectives and focuses on the comparative properties of different decomposition schemes i.e., intuitively, different patterns of division of labor within and across organizations.

Since a good deal of current interpretations at least of the vertical boundaries of economic organizations is grounded on transaction cost considerations, this is also a good place to start. Indeed, as we shall argue in section 2, the latter do tell part of the story but fail to account for those powerful drivers of intra- and inter-organizational division of labor which have to do with the nature of problem-solving knowledge, addressed by more "cognitive" approaches to organizational analysis (section 3). Next, building on the discussion of some fundamental features of problem-solving (section 4), a rather novel formalization of the decompositions of problems and tasks is presented in section 5 (perfect decompositions) and section 6 (near decompositions). Section 7 discusses some analytical and simulation-based properties of the model regarding the relative efficiency and speed of adaptation of diverse set-ups characterized by different boundaries between organizations and markets. Finally, in section 8 we draw some conclusions.

2 Problem solving tasks vs. transactions

Think of an industry or the whole economy as a sequence of tasks leading from, say, raw materials to final products. How does one "cut" such sequences within single organizations and across them?

As known, transaction costs economics (TCE henceforth) albeit rather silent on intra-organizational division of tasks, focuses upon the latter issue

– i.e. the *vertical* boundaries of organizations.

In a nutshell, TCE (cf. Williamson (1975), Williamson (1985) and Riordan and Williamson (1985), and the seminal argument first developed by Coase (1937)) starts from a hypothetical "state of nature" (a logical, if not a historical one) in which all coordination of transactions across technological separable units takes place within markets[1], and predicts that whenever the working of the market price mechanism incurs costs which are higher than the corresponding costs of bureaucratic governance, then the latter can prevail on the grounds of higher allocative efficiency.

This explanation however, although certainly capturing some determinants of the governance structure, does not tell in our view the whole story. Let us just mention three major difficulties of the theory which are crucial for the argument which follows (for a broader critical appraisal of transaction costs theory cf., for instance, Granovetter (1985) and Dow (1987)).

First, the logical process traced by Coase and Williamson often conflicts with actual historical records: with some remarkable exceptions, most technologies and industries are born with a highly vertically integrated structure, undergo a process of disintegration as the industry grows in the expansion phase and then re-integrate in the maturity phase, but often along profiles of integration which differ significantly from those of the original infant industry. Thus the degree of vertical integration of an industry undergoes major changes along its life cycle (Klepper 1997) and historical evidence seems to turn the transaction costs argument the other way round. One could say that "at the beginning" there were hierarchies and then they partly disintegrated giving rise to markets. Actually, it is the very process of division of labor, usually taking place within hierarchical organizations, which creates the opportunity for markets to exist: thus Williamson's story on markets as original state of nature presents clear limitations even as a logical instrument. Transaction cost views of vertical integration (and, for that matter, all standard vertical integration models based on information and agency problems, cf. Perry (1989)) appear relatively more appropriate to describe the processes of growing vertical integration which take place in mature industries where the division of labor is relatively stable and allocative efficiency requirements tend to prevail. On the contrary these models seem to have limited explanatory power when the early stages of the industry life-cycle are considered and whenever the main activity in which firms are engaged is the design of effective solutions to new technological and organizational problems.

Second, and relatedly, the transaction costs perspective deals with the efficiency of different governance structures in managing transactions across given technologically separable interfaces: technology and the division of labor are taken as given and organizational structures are derived. But the story could

---

[1] "...in the beginning there were markets" (Williamson 1975, p. 20)

be very different if one assumed that technology and the division of labor were themselves at least partly determined by the organizational structure: for example, one would easily obtain multiple organizational/technological equilibria (Pagano 1992, Aoki 2001) and a strong institutional path-dependency (David 1994). This point has been repeatedly emphasized also by scholars of the so-called "radical school" (cf., for instance, Bowles (1985) and Marglin (1974)) who had in mind an opposite view of the world in which it is primarily the governance structure which determines the technology and not the other way round. Even without taking a position in this old debate, it seems hardly questionable that most of the processes of division of labor take place within organizations. Therefore the latter cannot be taken as irrelevant with respect to where the technologically separable interfaces are placed and what their economic characteristics are. Moreover, a technologically separable interface requires well defined sets of codified standards for compatibility, especially if it has to be managed by transactions in a competitive market. As it is well known from the literature on technological standards (David and Greenstein 1990), they emerge either as unplanned conventions or as outcomes of deliberative processes (or combinations of the two) and in turn have a relevant influence on the directions of further technological change and division of labor. Again, markets cannot be original and spontaneous "states of nature", but require all sorts of institutional and technological conditions, some of which are put in place by explicit organizational planning. Moreover, once established, standards shape specific technological trajectories, limiting the directions of innovation.

Finally, a third weakness of the transaction costs approach resides in its account (or, better, the lack of it) of the processes through which superior governance structures do emerge. Proving that a given governance structure is more efficient than another one is not an explanation of its emergence through "spontaneous" processes driven by market selection: a selection mechanism can indeed, under some condition, select for fittest structures, but only if the latter exist in the first place[2]. Selection can account for convergence of a population to some given form, not for the emergence of such a form. The variational mechanisms through which new structures are generated and thus tested by the selection process are essential in determining the outcome of selection itself. If the set of possible structures is "large", only a small subset of it can ever be generated by any computationally feasible mechanism: thus we have to specify what it is the likelihood that such a subset include also the optimal structure.

---

[2] "...in a relative sense, the *fitter* survive, but there is no reason to suppose that they are *fittest* in any absolute sense" (Simon 1983, emphasis in original). On this point see also Winter (1975).

3 "Cognitive" perspectives on organizations: some roots in the literature

As already mentioned in the introduction, a respectable tradition, dating back at least to Adam Smith, attempts to identify the efficiency properties of different organizational forms by looking at the patterns of division of labor and at the learning opportunities which they entail, *quite independently from any issue of incentive compatibility and transaction governance.* Smith's famous example of the pin factory vividly illustrates the relationship between division of tasks, improvements of operational skills and opportunities for mechanization of production (see also the discussion in Leijonhufvud (1986) and Langlois and Robertson (1995)).

It is true, however, that what we could call a "procedural", knowledge-centered, approach to production and coordination patterns has been dormant for a long time. Rather, in mainstream economics the prevailing style of analysis has rested upon a thorough "blackboxing" summarized into production functions of various sorts.

In such a view the procedural aspects of production processes and, dynamically, of learning processes are explicitly censured. With that goes away also any investigation of the sequences of operations which are "legal", in the sense of being able to ultimately yield the desired output, and of their relative efficiencies. Of course one may always claim that these are issues for engineers and not for economists, but then the economists' analysis of the patterns of production and coordination also looses any reference to the underlying patterns of knowledge distribution and learning.

Certainly, Smith's seminal insights have been followed by some other major contributions to "procedural" analysis of the links between division of labor, production patterns and organizational forms. In the 19th century, Karl Marx's investigation of the capitalist factory system is an outstanding one, and Babbage's is another; in the 20th century the work of Georgescu-Roegen comes to mind; while across the two centuries several authors of the Austrian school have contributed to keep alive the interest in the importance of the links between forms of economic organization and the patterns of knowledge distribution within society (for a thorough discussion of several of these contributions see Langlois (1986) and Morroni (1992)).

All this notwithstanding, it is fair to say that a new impetus to procedural, knowledge centered, analysis of production and economic organization has mostly occurred over the last four decades. This has come together with the development and partial convergence of four interpretative perspectives, namely (i) the investigations by Herbert Simon and colleagues of the properties of problem-solving procedures in their relation to some measure of complexity

of the problems themselves[3]; (ii) behavioral theories of organizations in general and firms in particular[4]; (iii) evolutionary theories of economic change, with their emphasis on the process features of organizational knowledge and its partial embeddedness into organizational routines[5]; (iv) capabilities and competencies based views of firms[6].

As the intersection between these perspectives, the work which follows will try to offer a "constructive" – that is explicitly process-based – formal account of, first the links between problem-solving knowledge and division of labor within and across organizations and, second, the characteristics of diverse processes of selection amongst diverse organizational arrangements entailing distinctly different problem-solving repertoires.

The initial angle of investigation is clearly "Simonian". We put forward a notion of problem complexity which builds upon and refines Simon's ideas of decomposability and near-decomposability of complex problems (Simon 1969). An advantage of our notion is that it also allows straightforward mappings into selection dynamics wherein problem-solving entities are nested[7].

As we shall see, our notions of decomposability and the related one of problem complexity, bear upon the presence or absence of interrelations among the elementary activities which make up the overall problem-solving process.

It seems quite natural indeed to assume that business firms and other economic organizations fully belong to this category of complex entities made up of many non-linearly interacting elements.

One of the conjectures we shall investigate concerns in fact the evolution of vertical integration in terms of the characteristics of problem-solving tasks. The main argument can be stated as follows: the division of problem-solving labor into decentralized decision units coordinated by markets determines which solutions (i.e. technological and organizational designs) can be generated and then tested by the selection process. On the one hand this division is necessary for boundedly rational organizations in order to reduce the dimension of the search space, but, on the other hand, it might well happen that the division of physical and cognitive labor is such that the best designs will never be gener-

---

[3]Cf. for instance Simon (1969) and Simon (1983)

[4]To mention just the seminal works, cf. March and Simon (1958) and Cyert and March (1963)

[5]Cf. Nelson and Winter (1982) and also Nelson (1981) and Winter (forthcoming) more specifically on production theory, and Cohen et al. (1996) on routines.

[6]Cf., among others, Teece et al. (1997), Dosi et al. (2000). Important inspiring antecedents of this view are in the works of Penrose (1959) and Richardson (1972).

[7]Our model is also strictly related to the growing literature on modularity in technologies and organizations (Langlois and Robertson 1995, Baldwin and Clark 2000) and represents a formalization of a problem-solving approach to modularity. Problem decompositions define the modules on which selection applies. As we shall see, this problem-solving approach may bring a different perspective and different conclusions from the one based upon option value proposed by Baldwin and Clark (2000).

ated and therefore never selected by any selection mechanism whatsoever. In particular, we show that, everything else being equal, the higher the degree of decentralization the smaller the portion of the search space which is explored and the lower therefore the probability that the optimal solutions are included in such a portion of space. Finally, one can easily prove that computing the optimal division of problem-solving activities is more difficult than solving the problem itself, thus we cannot assume that boundedly rational agents in search for solutions to a given problem do possess the right decomposition of the problem itself[8].

In particular, if the entities under selection are made of many components which are interacting in complex way, the resulting selection landscape will present many local optima (Kauffman 1993) and selection forces will be unlikely to drive such entities to the global optima: sub-optimality and diversity of organizational structures can persistently survive in spite of strong selection forces (Levinthal 1997). Sub-optimality is due to the persistence of inferior features which cannot be selected out because of their tight connections with other favorable features: this indeed is the rule in strongly interconnected systems. In other words, whenever the entities under selection have some complex internal structure, the power of selective pressure is limited by the laws governing internal structures. In fact, one of the purposes of the present work is to provide a measure of these trade-offs and establish under which conditions either forces prevails.

4 Problem-solving: some special features

Problem-solving activities (which include, to repeat, most activities of production and innovation) present some distinctive features which make them difficult to analyze with standard economic tools. First of all, they involve, or are the outcome of, search in large combinatorial spaces of components which must be closely coordinated. Interdependencies among such components are only partly understood and can be only locally explored through e.g. trial-and-error processes, rules of thumb or the application of expert tacit knowledge.

Consider the following cases:

- The *design of complex artifacts* (e.g. an aircraft). It requires the coordination of many different design elements (engine type and power, wing

---

[8]One important caveat must be considered here: this paper assumes that there is a set of atomic components which cannot be further decomposed. This necessary analytical assumption does not allow to capture another important advantage of division of labor which is the possibility of further divisions: once a task has been specified, a new process of subdivision can be autonomously carried out on it. Contrary to what is assumed for simplicity in the model which follows, there is no given lower bound to the process of in-depth hierarchical decomposition.

size and shape, materials used, etc., each of them in turn composed of many elements) whose interactions can only partly be expressed by general models, but have to be tested through simulation, prototype building and trial-and-error exercises, where tacit knowledge plays a key role.

- The *solution to a difficult game* (e.g. solving a Rubik cube or playing chess). An effective solution is a long sequence of moves, each of which is chosen out of a set of possibilities which is large enough to make the exploration of the entire tree of the game computationally impossible for boundedly rational agents. The relations among such moves in a sequence - e.g. what changes we get in the overall performance of the solution if we change, say, the *i-th* move in the sequence we play - are only partly understood. Actually understanding it fully would imply the knowledge of the entire game tree [9].

- *Managing organizations* such as business firms. The latter are complex multi-dimensional bundles of routines, decision rules, procedures, incentive schemes, etc., whose interplay is largely unknown also to those who manage the organization itself – witness also all the problems and unforeseen consequences whenever managers try to promote changes in the organization.

Moreover, since components within a problem most often present strong interdependencies, the search space of a problem typically presents many local optima. Marginal contributions of components can rapidly switch from negative to positive values, depending on which value is assumed by the other components [10]. For instance, adding a more powerful engine could amount to decrease the performance and the reliability of an aircraft (Vincenti 1990) if other components are not simultaneously adapted. In a chess game, a notionally optimal strategy could involve - say- castling at a given moment in the development of the game, but the same castling as a part of some sub-optimal (but effective) strategy could turn out to be a losing move. Finally, introducing some routines, practices or incentive schemes which have proven superior in a given organizational context, could prove harmful in a context where other elements are not appropriately co-adapted.

As a consequence, in presence of strong interdependencies one cannot optimize a system by optimizing separately each element it is made of. Consider a

---

[9]In fact, one of the fundamental problems faced by human and artificial players is to build effective heuristics to evaluate the goodness of positions during the game, without the knowledge of the entire tree.

[10]Similar aspects are present even in the simplest production technologies. Consider for instance team production as exemplified by Alchian and Demsetz (1972): two workers lifting a heavy load. Additional individual effort generally rises team production, but when the levels of effort applied by the two are disproportionate, this might result in the load being turned over and falling, thus sharply decreasing team production.

problem which is made of N elements and whose optimal solution is $x_1^* x_2^* \ldots x_N^*$ while the current state is $x_1 x_2 \ldots x_N$. In the presence of strong interdependencies it might well be the case that some or even all solutions of the kind $x_1 x_2 \ldots x_i^* \ldots x_N$ show a worse performance than the current one [11].

However, as pointed out by Simon (1969), problem-solving by boundedly rational agents must necessarily proceed by decomposing any large, complex and intractable problem into smaller sub-problems which can be solved independently, i.e. by promoting what could be called the division of problem-solving labor. At the same time, note that the extent of the division of problem-solving labor is limited by the existence of interdependencies. If sub-problem decomposition separates interdependent elements, then solving each sub-problem independently does not allow overall optimization.

It is important to remark that the introduction of any decentralized interaction mechanism, like a competitive market, for each components does not solve the problem: for instance, if we assume that in our previous example each component x_i is traded in a competitive market, superior components x_i^* will never be selected for. Thus interdependencies undermine the effectiveness of the selection process as a device for adaptive optimization and introduce forms of path-dependency with lock-in into sub-optimal states which does not originates from the frictions and costs connected to the working of the selection mechanism, but from internal structure of the entities undergoing selection.

As Simon pointed out, since an optimal decomposition - i.e. a decomposition which divides into separate sub-problems all and only the elements which are independent from each other - can only be designed by someone who has a perfect knowledge of the problem (including its optimal solution), boundedly rational agents will be normally bound to design near-decompositions, i.e. decompositions which try to put together within the same sub-problem only those components whose interdependencies are (or, we shall add, agents believe to be) more important for the overall performance of the system. However near-decompositions involve a fundamental trade-off: on the one hand finer decompositions exploit the advantages of decentralized local adaptation, that is the use of selection mechanism for achieving coordination "for free" together with parallelism and speed of adaptation. However, on the other hand, finer decompositions imply a higher probability that interdependent components are separated into different sub-problems and therefore cannot, in general, be optimally adjusted together. One of the purposes of this paper is to provide a precise measure of this trade-off and show that, in the presence of widespread interdependencies, finer than optimal decompositions have an evolutionary advantage (in terms of speed of adaptation) although they inevitably involve lock-in into sub-optimal solutions.

---

[11]Note that this notion of interdependency differs from the notion of complementarity as sub-modularity as in Milgrom and Roberts (1990): here in fact we allow for the possibility that positive variations in one component can decrease overall performance value.

One way of expressing the limits that interdependencies pose to the division of problem-solving labor is that global performance signals are not able to drive effectively decentralized search in the problem space. Local moves in the "right direction" might well decrease the overall performance if some other elements are not properly tuned. As Simon puts it, since an entity (e.g. an organism in biology or an organization in economics) only receives feedbacks from the environment concerning the fitness of the whole entity, only under conditions of near independence the usual selection processes can work successfully for complex systems (Simon 2002, p. 593).

A further aspect concerns the property that, in general, the search space of a problem is not given exogenously, but is constructed by individuals and organizations as a subjective representation of the problem itself. If the division of problem-solving labor is limited by interdependencies, the structure of interdependencies itself depends on how the problem is framed by problem-solvers. Sometimes problem-solvers make major leaps forward by re-framing the same problem in a novel way. As shown by many case studies, often major innovations appear when various elements which were already well-known for a long time are recombined and put together under a different perspective (cf. for instance the detailed account of the development of wireless communication technologies given by Levinthal (1998)). Indeed, one can go as far as saying that it is the representation of a problem which determines its purported difficulty and that one of the fundamental functions of organizations is exactly to implement collective representations of the problems they face (Loasby 2000). In the simple model of problem-solving presented in this paper finding the "correct" representation of interdependencies is more complex than solving the problem itself. However, by changing the representation, lock-ins into sub-optimal solutions can be avoided and better solutions discovered. Division of problem-solving labor is therefore very much a question of how the problem is represented[12]. Needless to say, boundedly rational individuals cannot be innocently assumed to hold optimal representations.

Given the foregoing qualitative intuitions, let us next develop a formal model which provides a precise measure of the above mentioned trade-offs.

5 Decomposition and coordination

5.1 Definitions

We assume that solving a given problem requires the coordination of N atomic "elements" or "actions" or "pieces of knowledge", which we call generically **components**, each of which can assume some number of alternative states.

---

[12]A formal treatment of the properties of different representations in a particular class of problem can be found in Marengo (2003).

For simplicity, we assume that each component can assume only two alternative states, labelled 0 and 1. Note that all the properties presented below for the two-states case can be very easily extended to the case of any finite number of states.

Introducing some notation, we characterize a problem by the following elements:

The set of **components**: $C = \{c_1, c_2, \ldots c_N\}$ with $c_i \in \{0, 1\}$

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

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

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

In order to avoid some technical complications, we assume for the time being that there exists only one configuration which is strictly preferred to all the other configurations (i.e. a unique global optimum). This simplifying assumption will be dropped in section 6 below.

A **problem** is defined by the couple $(X, \succeq)$.

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 decompose[13] it into sub-spaces.

Let $I = \{1, 2, \ldots, N\}$ be the set of indexes and let a **block**[14] $d_i \subseteq I$ be a non-empty subset of it, we call the **size of block** d_i, its cardinality $|d_i|$. We define a **decomposition scheme** (or simply **decomposition**) of the problem $(X, \succeq)$ a set of blocks:

$$D = \{d_1, d_2, \ldots, d_k\}$$

such that $\bigcup_{i=1}^{k} d_i = I$

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

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

$$x^i(d_j) = x_{j_1}^i x_{j_2}^i \ldots x_{j_{|d_j|}}^i$$

---

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

[14]Blocks in our model can be considered as a formalization of the notion of modules used by the growing 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

for all $j_h \in d_j$.

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.

Two block-configurations can be united into a larger block-configuration by means of the $\vee$ operator so defined:

$$x(d_j) \vee y(d_h) = z(d_j \cup d_h) \text{ where } z_\nu = \begin{cases} x_\nu \text{ if } \nu \in d_j \\ y_\nu \text{ otherwise} \end{cases}$$

We can therefore write $x^i = x^i(d_j) \vee x^i(d_{-j})$ for any d_j.

Moreover, we define the **size of a decomposition scheme** as the size of its largest defining block:

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

Coordination among blocks in a decomposition scheme 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[15].

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

$$x^j(d_h) \vee x^i(d_{-h}) \succeq x^i(d_h) \vee x^i(d_{-h})$$

then $x^j(d_h)$ is selected and the new configuration $x^j(d_h) \vee x^i(d_{-h})$ is kept in the 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 structure of division of labor (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 do 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 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

---

[15]As a first approximation, we assume that this sorting and selection mechanism is errorless and operates at no cost and without any friction.

smaller and smaller portions of the search space, thus decreasing the likelihood that optimal (or even good) solutions are ever generated and tested. In the following we try to provide a precise characterization of this trade-off and its consequences.

5.2 Selection and search paths

A decomposition scheme 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 scheme: 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 scheme $D = \{d_1, d_2, \ldots, d_k\}$, 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. $x^i \succeq x^j$

2. $x^i_\nu = x^j_\nu \ \forall \nu \notin d_h$

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:

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

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

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

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

A search path or, shortly, 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:

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

A configuration x^j is **reachable** from another configuration x^i and for decomposition D if there exist 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 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[16] and let $Z \subseteq X$ with $x^0 \in Z$, we say that the problem $(X, \succeq)$ is locally decomposable in Z by the scheme D if $Z \subseteq \Psi(x^0, D)$. If $Z = X$ we say that the problem is globally decomposable by the scheme D[17].

Among all the decomposition schemes of a given problem, benchmark cases are those for which the global optimum becomes reachable from any starting configuration. One such decomposition always exists, and is the degenerate decomposition $D = \{\{1, 2, 3, \ldots, N\}\}$ for which of course there exists only one local optimum and it coincides with the global one. But obviously we are interested in smaller decompositions – if they exist – and in particular in those of minimum size. The latter decompositions represent the maximum extent to which problem-solving can be subdivided into independent sub-problems coordinated by decentralized selection, with the property that such selection processes can eventually lead to optimality irrespective of the starting condition. On the contrary, finer decompositions will not in general allow decentralized selection processes to optimize (unless the starting configuration is "by chance" within the basin of attraction of the global optimum).

The following proposition shows that there are problems which are globally decomposable only by the degenerate decomposition $D = \{\{1, 2, 3, \ldots, N\}\}$.

Proposition 1 *There exist problems which are globally decomposable only by the degenerate decomposition $D = \{\{1, 2, 3, \ldots, N\}\}$*

Proof: we prove the statement by providing an example. Consider a problem whose unique global optimum is configuration $x^0 = x_1^0 x_2^0 \ldots x_N^0$ and whose second best configuration is $x^1 = x_1^1 x_2^1 \ldots x_N^1$ where $x_i^1 = |1 - x_i^0| \, \forall i = 1, 2 \ldots, N$. It is obvious that the global optimum can be reached from the second best only by mutating all the N components together.

---

[16] We remind the assumption of uniqueness of the global optimum.

[17] A special case of decomposability, which is generalized here, is presented in Page (1996) and is called dominance. In our terminology, a block configuration $x^j(d_h)$ is dominant when $x^j(d_h) \vee x^i(d_{-h}) \succeq x^i$ for every configuration $x^i \in X$.

The next proposition establishes a rather obvious but important property of decomposition schemes. As one climbs into the basin of attraction of a local optimum for a decomposition D which is not the finest one, then finer decomposition schemes can usually be introduced which allow to reach more quickly the same local optimum.

For this proposition we need an additional definition: given a decomposition scheme D we say that two configurations x^i and x^j totally differ with respect to block $d_h \in D$ if the corresponding block configurations $x^i(d_h)$ and $x^j(d_h)$ differ in every components: $x_k^i(d_h) \neq x_k^j(d_h) \ \forall k = 1, 2, \ldots, |d_h|$.

Proposition 2 *Let $\Psi(x^\alpha, D) = \{x^\alpha, x^{\alpha+1}, \ldots, x^{\alpha+m}\}$ be the ordered basin of attraction of a local optimum x^α (with $x^{\alpha+j} \succeq x^{\alpha+j+1} \ \forall j = 0, \ldots, m - 1$). Define $\Psi^i(x^\alpha, D) = \Psi(x^\alpha, D) \setminus \{x^{\alpha+i+1}, x^{\alpha+i+2}, \ldots, x^{\alpha+m}\}$ for $0 \leq i \leq m$. Let $d_{\nu_i} \in D$ be the block(s) of maximum size in D. Then, unless x^α and $x^{\alpha+1}$ totally differ for some maximum size block d_{ν_i}, there exists a $0 < i \leq m$ such that the set $\Psi^i(x^\alpha, D)$ admits a decomposition D^i with $|D^i| < |D|$.*

Proof: Suppose for simplicity that D contains a unique maximum size block $d_\nu \in D$ with $|D| = |d_\nu|$. If the local optimum x^α and the second best of its basin of attraction (with respect to D) $x^{\alpha+1}$ do not totally differ with respect to d_ν, then there exists a smaller decomposition D^i which is identical to D except that its largest block d_ν can be split into two sub-blocks containing respectively the components for which $x^\alpha(d_\nu)$ and $x^{\alpha+1}(d_\nu)$ differ and those for which they do not. By construction x^α is reachable from $x^{\alpha+1}$ for D^i and $|D^i| < |D|$ and therefore $i = 1$ satisfies the proposition. If there are multiple maximum size blocks $d_{\nu_i} \in D$ it is necessary that x^α and $x^{\alpha+1}$ do not totally differ for any of them.

Among all the possible global decompositions of a problem, those of minimum size are especially interesting: in fact they set a lower bound to the degree of decentralization which preserves optimality with certainty. Conversely, note for decompositions which are finer than those of minimum size whether the optimal solution will ever be generated and thus selected depends on the initial condition.

Minimum size decomposition schemes can be found recursively with the procedure informally described in the following[18]:

Let us re-arrange all the configurations in X by descending rank $X = \{x^0, x^1, \ldots, x^{2^N - 1}\}$ where $x^i \succeq x^{i+1}$.

The minimum size decomposition can be computed as follows:

1. start with the finest decomposition $D^0 = \{\{1\}, \{2\}, \ldots, \{N\}\}$

---

[18]The complete algorithm is quite lengthy to describe in exhaustive and precise terms. Its Pascal and C++ implementations are available from the authors upon request.

2. check whether $x^0 \in P(x^i, D)$ $\forall x^i$ $i = 1, 2, \ldots, 2^N - 1$, i.e. if there is a path leading to the global optimum from every other configuration for decomposition D, if yes STOP

3. if no, build a new decomposition D^1 by union of the smallest blocks for which condition 2 was violated and go back to 2.

Let us illustrate it with an example.

Example: An hypothetical ranking (where 1 is the rank of the most preferred) of configurations for N=3:

CONFIGURATIONS	RANKING
100	1
010	2
110	3
011	4
001	5
000	6
111	7
101	8

If search proceeds according to the decomposition scheme $D = \{\{1\}, \{2\}, \{3\}\}$, there exist two local optima: 100 (which is also the global optimum) and 010. The basins of attraction of the two local optima are respectively:

$$\Psi(100) = \{100, 110, 000, 111, 101\}$$

$$\Psi(010) = \{010, 110, 011, 001, 000, 111, 101\}$$

Note that the worst local optimum has a larger basin of attraction[19] as it covers all configurations except the global optimum itself. Thus, only a search which starts at the global optimum will (trivially) stop at the global optimum with certainty, while for the other initial configurations search might end up in either local optima (depending on the sequence of mutations) or even (in three cases) in the worst local optimum with certainty.

Using the notion of dominance (cf. Page (1996) and footnote 17 above) it is possible to establish that the only dominant block-configuration is actually the globally optimum string itself, corresponding to the degenerate decomposition scheme $D = \{\{1, 2, 3\}\}$. Thus apparently no

---

[19]Kauffman (1993)provides some general properties of one-bit-mutation search algorithms (equivalent to our bit-wise decomposition schemes) on string fitness functions with varying degrees of interdependencies among components. In particular, he finds that as the span of interdependencies increases, the number of local optima increases too, while the size of the basin of attraction of the global optimum shrinks.

decentralized search structure allows to always locate the global optimum from every starting configuration.

Granted that, can one find some alternative decompositions allowing for partly decentralized search processes yielding global optima? In our example, one of such cases occurs with the decomposition scheme $D = \{\{1, 2\}, \{3\}\}$. For instance if one starts from configuration 111, one can first locate 011 (using block $\{1, 2\}$) then 010 (using block 3) and finally 100 (again with block $\{1, 2\}$); or alternatively one can locate 110 (using block 3) and 100 (with block 1,2). It can be easily verified that the same blocks do actually "work" for all other starting configurations. The algorithm just presented will find this decomposition.

6 Near decomposability

When building a decomposition scheme for a problem we have looked so far for perfect decomposability, in the sense that we require that all blocks can be optimized in a totally independent way from the others. In this way we are guaranteed to decompose the problem into perfectly isolated components which can be solved independently. This is however very stringent a requirement: even when interdependencies are rather weak, but diffused across all components, one easily tends to observe problems for which no perfect decomposition of size smaller than N exists. For instance figure 1 shows that in Kauffman's NK random landscapes[20], already for very small values of K - that is for highly correlated landscapes - the above described algorithm finds only decomposition schemes of size N or just below N. In other words, a little bit of interdependence spread across the set of components makes immediately a system practically indecomposable.

One can soften the requirement of perfect decomposability into one of near-decomposability: one no longer requires the problem to be decomposed into completely separated sub-problems, i.e. sub-problems which fully contain all interdependencies, but one might just be content to find sub-problems which contain the most "relevant" interdependencies while less relevant ones can

---

[20]A NK random fitness landscape is similar to our definition of problem except that instead of a preference relation, a real valued fitness function $F : X \mapsto \mathbb{R}$ is defined as an average of each component's fitness contribution. The latter is a random realization of a random variable uniformly distributed over the interval $[0, 1]$ for each possible configuration of the K-size block of the other components with which each component interacts (Kauffman 1993). Note however that Kauffman's K is not a good measure of ex-post complexity – in terms of its decomposability – of the optimization problem on the resulting fitness landscape: small values of K usually generate landscapes which are not decomposable, but, on the other side, it is always possible that even with very high values of K the resulting landscape is highly decomposable.

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[21]. 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 higher speed of adaptation, are likely to be obtained at the expenses of the asymptotic optimality of the solutions which can be reached.

Let us re-arrange all the configurations in X by descending rank $X = \{x^0, x^1, \ldots, x^{2^N-1}\}$ where $x^i \succeq x^{i+1}$, and let $X_\mu = \{x^0, x^1, \ldots, x^{\mu-1}\}$ with $0 \le \mu \le 2^N - 1$ 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 exist a configuration $x^i \in X_\mu$ such that $x^i \in P(y, D)$.

We call 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 decompositions schemes for optimal decompositions.

The following proposition gives the most important property of minimum size μ-decompositions:

Proposition 3 *Let D_μ is a minimum size μ-decomposition for problem $(X, \succeq)$ then $|D_\mu|$ is monotonically weakly decreasing in μ.*

Proof: if $\mu = 2^N - 1$ the set X_μ includes all configurations and it is trivially reachable for all decompositions, including the finest with $|D_\mu| = 1$. If $\mu = 1$ then X_μ includes only the global optimum and therefore the size of the minimum size decomposition is $1 \le |D_\mu| \le N$. We still have to show that it cannot happen that $|D_{\mu+1}| > |D_\mu|$: if this was the case X_μ could not be reached from $X_{\mu+1}$ for decomposition D_μ, but this contradicts the assumption that X_μ is reachable from any configuration in X for decomposition D_μ.

The latter proposition shows that higher degrees of decomposition and decentralization can be attained by giving up optimality and provides a precise measure for this trade-off. As an example, we generated 100 random problems of size $N = 12$ all characterized by not being decomposable[22] (i.e. $|D| = 12$

---

[21]This procedure allows to deal also with the case of multiple global optima and thus we can now drop also the assumption of a unique global optimum.

[22]Random problems are generated in a straightforward way: we generate random rankings of all the binary strings of size $N = 12$ and then compute – using the algorithm presented above – their decomposition schemes of minimum size. Only those problems for which the size of the smallest decomposition schemes was 12 were used in this simulations. An alternative (and equivalent) method is to generate random NK landscapes *à la* Kauffman (1993) with $N = 12$ and a high K and then check that the resulting landscape is not decomposable, as it may happen that also landscapes with a very high K may be highly decomposable.

for all of them). Figure 2 displays the average sizes of the minimum size decomposition schemes for the 100 random problems as we vary the number μ of acceptable configurations. Figure 2 shows that second best solutions can be reached by search processes based upon finer decompositions, that is with more decentralized processes, which can find such solutions more quickly by exploiting coordination "for free" provided by the solution mechanism. In fact, when the size of the decomposition scheme decreases of one unit the expected search time decreases by half.

7 Speed and optimality in search: some consequences for organizational structures

The trade-offs outlined in the previous sections between decomposability, reduction of complexity and speed of search, on the one hand, and asymptotic optimality on the other, enables us to discuss some interesting evolutionary properties of various organizational structures competing in a given problem-solving environment.

Let us consider the properties of near-decompositions. As illustrated in figure 2 for randomly generated problems[23], if second best solutions are accepted it is possible to have considerable reductions of the size of decomposition schemes and of the expected time of search. This outlines that the organizational structure sets a balance in the trade-off between speed and optimality of search and adaptation. 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 optimal given the interdependencies of the problem space. This property is shown in figures 3 and 4, which present the typical search paths on a non-decomposable problem of two search processes driven respectively by decompositions:

$$D1 = \{1, 2, \ldots, 12\}$$
$$D12 = \{\{1\}, \{2\}, \ldots, \{12\}\}$$

Figure 3 shows the first 180 iterations in which the more decentralized structure (D12) quickly climbs the problem space and outperforms search based on a coarser decomposition. If there was a tight selection environment, more than optimally decentralized organizational structure would quickly displace the structure D1 which reflects the "true" decomposition of the underlying problem space.

However the search process based on the finest decomposition quickly reaches a local optimum from which cannot make any further improvement,

---

[23]In this and the following figures (with the exception of figure 5) on the vertical axis we indicate the rank of configurations re-parametrized between 0 (worst) and 1 (best).

while the process based on the coarser decomposition keeps searching and slowly climbing. Figure 4 shows iterations between 3000 and 3800, where the finest decomposition is still locked-into the local optimum it reached after very few iterations, while the coarsest one slowly reaches the global optimum (normalized to 1). Strong selective pressure tends therefore to favor organizational structures whose degree of decentralization is higher than what would be optimal from a mere problem-solving perspective.

This result is even stronger in problems that we could define "modular", i.e. characterized by blocks with strong interdependencies within blocks and much weaker - but non-zero - interdependencies between blocks: in this problems higher levels of decompositions can be achieved at lower costs in terms of sub-optimality.

Another important property concerns micro ("idiosyncratic") path-dependencies of organizational forms and their long-term persistence. If finer-than-optimal decompositions tend to emerge and to spread because of their "transient" evolutionary advantages, then one will in general observe also long-term diversity in the population of organizations in terms of (i) the decomposition they are based upon; (ii) the problem solutions they implement; and (iii) the local peaks they settle in[24]. This is easily shown by a simulation exercise in which we model a simple selection environment in which we generate 100 organizations characterized by a randomly generated decomposition and a random initial string and let them search a randomly generated indecomposable problem. Every 10 iterations the worst performing 10 organizations are selected out and replaced by 10 new organization among which 5 are randomly generated and 5 have the same decomposition scheme of the best performing ones but are placed on a randomly chosen configuration.

Figure 5 plots the number of diverse organizational forms at every iteration. Initially, diversity does indeed sharply decrease because of selective pressure but then stabilizes on numbers consistently and persistently higher than 1.

A very similar trend describes the number of surviving different configurations, which reflect the fact that the population of organizations settles onto several local peaks of similar value.

We have also run other simulations in which, at given intervals, we have changed the current problem with one having exactly the same structure in terms of decomposability, but with different, randomly generated, orderings of configurations. This can be taken as a metaphorical proxy for volatility of the environment. For instance consumers might have changing preferences over a given set of characteristics, or, on the production side, relative input prices may change. Interestingly enough, it turns out that even with totally decomposable problems, as the change of the orderings becomes more frequent, the population is entirely invaded by organizations characterized by coarser

---

[24]On this latter point a similar result is obtained by Levinthal (1997).

and coarser decompositions, and at the limit by organizations which do not decompose at all. This robustly suggests that growing volatility has stronger consequences then those of growing interdependence. The reason why this happens is shown in figures 6 and 7 which present, respectively, the expected improvements and the probability of improvement for searches based upon the finest (D12) and coarsest (D1) decomposition schemes in a fully decomposable problem[25]. It is shown that, when starting from low rank configurations, search based upon coarser decomposition has both a higher probability of finding a better configuration and when such a better configuration is found, its expected rank is higher for coarser decompositions. This is due to the fact that finer decompositions search only locally and in fully decomposable problems this on average cannot produce large improvements. When the problem space is highly volatile – though always fully decomposable – sooner or later every organization will fall into an area of very "bad" configurations from which coarser decompositions have a higher chance to promptly recover.

8   Conclusions

In this work we have presented a novel model of 'Simonian' ascendancy concerning the properties of the division of problem solving labor and also to account for the properties of different institutional arrangements and in particular for different boundaries between un-decomposed (in principle organization-embodied) tasks, and decomposed ones (possibly coordinated via market-like mechanisms but also via mechanisms based on the interaction of quasi-independent units within simple organizations).

The issue is basically one of organizational (and technological) design: can optimal organizational structures (or optimal technological designs) emerge out of decentralized local interactions? The paper in fact shows that this is possible only under some special and rather implausible conditions and that, on the contrary, the advantages of decentralization bear usually a cost in terms of sub-optimality.

The results are largely consistent with Simon's general proposition suggesting that "... near decomposability is an exceedingly powerful architecture for effective organization ... [which] appear with regularity also in human social organization – e.g. business firms and government agencies – with their many-layered hierarchies of divisions, departments and sections ..." (Simon 2002, pp. 598-99). However our model also highlights the subtle trade-offs between levels of decomposition, degrees of suboptimality of the achievable outcomes, and speed of adaptation. Together, it casts strong doubts on the general

---

[25]Figures 6 and 7 refer to the fully decomposable search space given by the binary numbers between 0 and $2^N - 1$. But the same qualitative results are obtained for any kind of fully decomposable search space.

validity of any "optimality-through-selection" argument in the space of organizational and technological designs and, more in general, of any "optimistic" view of market selection processes (expressed for instance by Alchian (1950) and Friedman (1953)) as forces which substitute individual optimization with evolutionary optimization.

On more empirical grounds, the analysis of the foregoing trade-offs can help explain the changing depth and profile of integration of organizations along technology and industry life cycles. The results of our model are well consistent with an empirical story according to which new technologies develop in highly integrated organizations because of the need to control for the strong interdependencies which characterize difficult problems. Market-like decentralized mechanisms, it has been argued, do not provide appropriate signals in this early "problem-solving" phase, because they do not - except in very simple problems - allow for the coordination of interdependent elements. As search proceeds and a local peak (a set of standards in the techno-organizational design problem) is selected, the degree of decentralization can be greatly increased in order to allow for fast climbing of this peak (and indeed transaction costs factors can very well be responsible at this stage for variations of the degree of integration), but the more decentralization is pushed forward, the more unlikely it will be that new and better local optima can be discovered. There is an unavoidable trade-off between decentralization and optimality which can hardly be escaped.

Finally we suggested that organization could actually play the even more fundamental role of building collective representations of the problems to solved, and that such representations could act as frames within which the division of labor takes place inside and across organizations. Consider such a proposition as the beginning of a promising line of inquiry concerning the crucial role of organization in the construction of collectively shared representations - fundamental ingredients, as such, of coordination in presence of any form of division of cognitive and productive labor.

References

ALCHIAN, A. (1950) "Uncertainty, evolution and economic theory", *Journal of Political Economy*, **58** pp. 211–222.

ALCHIAN, A. and DEMSETZ, H. (1972) "Production, information costs and economic organization", *American Economic Review*, **62** pp. 777–795.

AOKI, M. (2001) *Towards a comparative institutional analysis*, MIT Press, Cambridge, Mass.

BALDWIN, C. Y. and CLARK, K. B. (2000) *Design Rules. The power of modularity*, MIT Press, Cambridge Mass.

BECKER, G. S. and MURPHY, K. M. (1992) "The division of labour, coordination costs and knowledge", *Quarterly Journal of Economics*, **107** pp. 1138–1140.

BOWLES, S. (1985) "The production process in a competitive economy: Neo-Hobbesian and Marxian models", *American Economic Review*, **75** pp. 16–36.

COASE, R. H. (1937) "The nature of the firm", *Economica*, **4** pp. 386–405.

COHEN, M. D., BURKHART, R., DOSI, G., EGIDI, M., MARENGO, L., WARGLIEN, M. and WINTER, S. G. (1996) "Routines and Other Recurring Action Patterns of Organizations: Contemporary Research Issue", *Industrial and Corporate Change*, **5** pp. 653–699.

CYERT, R. M. and MARCH, J. G. (1963) *A Behavioral Theory of the Firm*, Prentice Hall, Englewood Cliff, NJ.

DAVID, P. A. (1994) "Why are institutions the 'carriers of history'? Path dependence and the evolution of conventions, organizations and institutions", *Structural Change and Economic Dynamics*, **5** pp. 205–220.

DAVID, P. A. and GREENSTEIN, S. (1990) "The economics of compatibility standards: An introduction to recent research", *Economics of Innovation and New Technology*, **1** pp. 3–41.

DOSI, G., NELSON, R. R. and WINTER, S. G. (2000) *The Nature and Dynamic of Organizational Capabilities*, Oxford University Press, Oxford.

DOW, G. K. (1987) "The function of authority in transaction cost economics", *Journal of Economic Behavior and Organization*, **8** pp. 13–38.

FRIEDMAN, M. (1953) *Essays in Positive Economics*, University of Chicago Press, Chicago.

GRANOVETTER, M. (1985) "Economic action and social structure: The problem of embeddedness", *American Journal of Sociology*, **91** pp. 481–510.

KAUFFMAN, S. A. (1993) *The Origins of Order*, Oxford University Press, Oxford.

KLEPPER, S. J. (1997) "Industry life cycles", *Industrial and Corporate Change*, **6** pp. 145–181.

LANGLOIS, R. N. (1986) *Economics as a process: Essays in the new institutional economics*, Cambridge University Press, Cambridge.

LANGLOIS, R. N. and ROBERTSON, P. L. (1995) *Firms, Markets and Economic Change*, Routledge, London and New York.

LEIJONHUFVUD, A. (1986) "Capitalism and the factory system", in: Langlois (1986), pp. 203–223.

LEVINTHAL, D. (1997) "Adaptation on rugged landscapes", *Management Science*, **43** pp. 934–950.

LEVINTHAL, D. (1998) "The slow pace of rapid technological change: Gradualism and punctuation in technological change", *Industrial and Corporate Change*, **7** pp. 217–247.

LOASBY, B. J. (1998) "The organization of capabilities", *Journal of Economic Behavior and Organization*, **25** pp. 139–160.

LOASBY, B. J. (2000) "Organizations as interpretative systems", Unpublished Manuscript.

MARCH, J. G. and SIMON, H. A. (1958) *Organizations*, Wiley, New York.

MARENGO, L. (2003) "Problem complexity and problem representation: some implications for the theory of economic institutions", in: RIZZELLO, S., ed., *Cognitive Developments in Economics*, pp. 371–388, Routledge, Dordrecht.

MARGLIN, S. A. (1974) "What do bosses do? The origins and function of hierarchy in capitalist production", *Review of Radical Political Economy*, **6** pp. 33–60.

MARSCHAK, J. and RADNER, R. (1972) *Economic Theory of Teams*, Yale University Press, New Haven.

MILGROM, P. and ROBERTS, J. (1990) "The economics of modern manufacturing", *American Economic Review*, **80** pp. 511–528.

MORRONI, M. (1992) *Production process and technical change*, Cambridge University Press, Cambridge.

NELSON, R. R. (1981) "Research on productivity differences", *Journal of Economic Literature*, **19** pp. 1029–1064.

NELSON, R. R. and WINTER, S. G. (1982) *An evolutionary theory of economic change*, The Belknap Press of Harvard University Press, Cambridge, MA.

PAGANO, U. (1992) "Organizational equilibria and production efficiency", *Metroeconomica*, **43** pp. 227–246.

PAGE, S. E. (1996) "Two measures of difficulty", *Economic Theory*, **8** pp. 321–346.

PENROSE, E. (1959) *The Theory of the Growth of the Firm*, Wiley, New York.

PERRY, M. K. (1989) "Vertical integration", in: SCHMALENSEE, R. and WILLIG, R., eds., *Handbook of Industrial Organization*, volume I, North Holland, Amsterdam.

RADNER, R. (2000) "Costly and bounded rationality in individual and team decision making", *Industrial and Corporate Change*, **9** pp. 623–658.

RICHARDSON, G. B. (1972) "The organization of industry", *Economic Journal*, **82** pp. 883–896.

RIORDAN, M. and WILLIAMSON, O. E. (1985) "Asset specificity and economic organization", *International Journal of Industrial Organization*, **3** pp. 365–378.

SIMON, H. A. (1969) *The Sciences of the Artificial*, MIT Press, Cambridge, MA.

SIMON, H. A. (1983) *Reason in Human Affairs*, Stanford University Press, Stanford.

SIMON, H. A. (2002) "Near decomposability and the speed of evolution", *Industrial and Corporate Change*, **11** pp. 587–599.

TEECE, D. J., PISANO, G. and SHUEN, A. (1997) "Dynamic capabilities and strategic management", *Strategic Management Journal*, **18** pp. 509–533.

TEECE, D. J., RUMELT, R., DOSI, G. and WINTER, S. G. (1994) "Understanding Corporate Coherence: Theory and Evidence", *Journal of Economic Behavior and Organization*, **23** pp. 1–30.

VINCENTI, W. (1990) *What Engineers Know and How They Know It: Analytical Study from Aeronautical History*, The John Hopkins University Press, Baltimore.

WILLIAMSON, O. E. (1975) *Markets and Hierarchies: Analysis and Antitrust Implications*, Free Press, New York.

WILLIAMSON, O. E. (1985) *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*, MacMillan, London.

WINTER, S. G. (1975) "Optimization and evolution in the theory of the firm", in: DAY, R. H. and GROVES, T., eds., *Adaptive Economic Models*, Academic Press, New York.

Winter, S. G. (forthcoming) "Toward and Evolutionary Theory of Production", in: Dopfer, K., ed., *Principles of Evolutionary Economics*, Cambridge University Press, Cambridge.

Figure 1: Size of minimum decomposition schemes for random NK problems. (N=12)

Figure 2: Near decomposability.

Figure 3: Fitness values for search processes with finest (D12) and coarsest(D1) decompositions (N=12). First 180 iterations. . .

Figure 4: . . . after 3000 iterations.

Figure 5: Diversity of surviving decomposition schemes.

Figure 6: Expected gain from rank-improving mutations for the finest (D12) and coarsest (D1) decompositions in a fully decomposable problem. (N=12)

Figure 7: Probability of rank-improving mutations for the finest (D12) and coarsest (D1) decompositions in a fully decomposable problem. (N=12)