

Division of labor, organizational coordination and market mechanisms in collective problem-solving

Luigi Marengo^{a,*}, Giovanni Dosi^{b,1}

^a *DSGSS, Università di Teramo, Loc. Colleparco, 64100 Teramo, Italy*

^b *Scuola Superiore S. Anna, P.zza Martiri della Libertà 33, 56127 Pisa, Italy*

Received 16 March 2003; accepted 10 March 2004

Available online 25 August 2005

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, determine which solutions are tried out and undergo selection. It is shown that if the design problem at hand is “complex” (in terms of interdependencies between the elements of the system), then a decentralized institutional structure is unlikely ever to 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.

© 2005 Elsevier B.V. All rights reserved.

JEL classification: C63; D21; D23; L22; O30

Keywords: Theory of the firm; Vertical and horizontal integration; Computational complexity

1. Introduction

One way to describe any economy or, for that matter, any economic organization, is as a huge ensemble of partially interrelated tasks and processes that, combined in certain ways,

* Corresponding author. Tel.: +39 050883343; fax: +39 050883344.

E-mail addresses: marengo@unite.it (L. Marengo), gdosi@sssup.it (G. Dosi).

¹ Tel.: +39 0861266661; fax: +39 0861266665.

produce “well-constructed” goods and services. It is a perspective that dates at least back 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 of Herbert Simon’s seminal works have explored the general structure of problem-solving activities of which technological search and economic production activities are just subsets (Simon, 1969). From different angles, several investigations from the team theory perspective have addressed the symmetric problem concerning coordination amongst multiple interrelated tasks (Marschak and Radner, 1972; Radner, 2000; Becker and Murphy, 1992). Finally, a flourishing literature has focussed on the “cognitive” characteristics of organizations (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 of at least 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 that concern the nature of problem-solving knowledge, addressed by more “cognitive” approaches to organizational analysis (Section 3). Next, by building on the discussion of some fundamental features of problem-solving (Section 4), a rather novel formalization of the decomposition 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, we draw some conclusions in Section 8.

2. Problem-solving tasks versus 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 (henceforth, TCE), albeit rather silent on the intra-organizational division of tasks, focuses on the latter issue, the *vertical* boundaries of organizations.

In a nutshell, TCE (Williamson, 1975, 1985; 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 markets¹ and predicts that whenever the working of the market price mechanism incurs costs that are higher than the corresponding costs of bureaucratic governance, then the latter can prevail on the grounds of higher allocative efficiency.

¹ “... in the beginning there were markets” (Williamson, 1975, p. 20).

In our view, however, this explanation does not tell the whole story, although it does indeed capture some determinants of the governance structure. Let us just mention three major difficulties of the theory that are crucial for the argument that follows.²

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 disintegration process as the industry grows in the expansion phase, and then re-integrate in the maturity phase, but often along integration profiles that 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 around. One could say that “at the beginning” there were hierarchies and then they partially disintegrated, giving rise to markets. Actually, it is the very process of division of labor, usually taking place within hierarchical organizations, that 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 (Perry, 1989), appear relatively more appropriate in describing the processes of growing vertical integration that 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 firm main activity 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 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 as well been repeatedly emphasized by scholars of the so-called “radical school” (Bowles, 1985; Marglin, 1974) who had in mind an opposite view of the world in which it is primarily the governance structure that determines the technology, not the other way around. Even without taking a position in this old debate, it seems hardly questionable that most of the division of labor processes 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 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 as 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”,

² For a broader critical appraisal of transaction costs theory see, for instance, Granovetter (1985) and Dow (1987).

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 certain conditions, select for the fittest structures, but only if the latter exist in the first place.³ Selection can account for the convergence of a population toward 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 the 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 the likelihood is that such a subset also includes the optimal structure.

3. “Cognitive” perspectives on organizations: some roots in the literature

As already mentioned in Section 1, a respectable tradition, dating at least back 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 in operational skills and opportunities for production mechanization, as argued in, e.g., 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.

With such a view, the procedural aspects of production processes and, dynamically, of learning processes are explicitly censured. This also includes the removal of any investigation of the sequences of operations that are “legal”, in the sense of being able ultimately to 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 loses 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 the “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

³ “. . . 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).

of Georgescu–Roegen comes to mind; while across the two centuries several authors of the Austrian school have kept the interest in the importance of the links between forms of economic organization and the patterns of knowledge distribution within society alive (Langlois, 1986; Morroni, 1992).

All this notwithstanding, it is fair to say that a new impetus to a procedural, knowledge-centered analysis of production and economic organization has occurred mostly 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,⁴ (ii) behavioral theories of organizations in general and firms in particular,⁵ (iii) evolutionary theories of economic change, with their emphasis on the process features of organizational knowledge and its partial embeddedness into organizational routines⁶ and (iv) capabilities and competencies-based views of firms.⁷

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 where problem-solving entities are nested.⁸

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 that 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 one hand, this division

⁴ See for instance Simon (1969, 1983).

⁵ To mention just the seminal works, see March and Simon (1958) and Cyert and March (1963).

⁶ See Nelson and Winter (1982) as well as Nelson (1981) and Winter (2004) more specifically on production theory, and Cohen et al. (1996) on routines.

⁷ See, among others, Teece et al. (1997), Dosi et al. (2000). Inspiring antecedents of this view are in the works of Penrose (1959) and Richardson (1972).

⁸ 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.

is necessary for boundedly rational organizations in order to reduce the dimensions 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 generated 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 that is explored and therefore the lower 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 of solutions to a given problem do possess the right decomposition of the problem itself.⁹

In particular, if the entities under selection are made up of many components that interact in a 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 that 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 force prevails.

4. Problem-solving: some special features

Problem-solving activities (which include, we repeat, most production and innovation activities) present some distinctive features that make them difficult to analyze with standard economic tools. First of all, they involve or are the outcome of a search in large combinatorial spaces of components, which must be closely coordinated. Interdependencies among such components are only partly understood and can only be locally explored through, for example, 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 size and shape, materials used, and so on, each of them, in turn, composed of many elements) whose interactions can only partly be expressed by general models and they have to be tested through simulation, prototype building and trial-and-error exercises, where tacit knowledge plays a key role.

⁹ One important caveat must be considered here: this paper assumes that there is a set of atomic components that cannot be further decomposed. This necessary analytical assumption does not allow us to capture another important advantage of division of labor, 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's sake, in the model that follows, there is in general no given lower bound to the process of in-depth hierarchical decomposition.

- The *solution to a difficult game* (e.g. solving a Rubik's cube or playing chess). An effective solution is a long sequence of moves, each of which is chosen out of a set of possibilities that 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) are only partly understood. Actually, understanding it fully would imply the knowledge of the entire game tree.¹⁰
- *Managing organizations* such as business firms. The latter are complex multi-dimensional bundles of routines, decision rules, procedures, incentive schemes and so on, whose interplay is largely unknown also to those who manage the organization itself—witness 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 inter dependencies, 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.¹¹ For instance, adding a more powerful engine could lower 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, for example, 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 that 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 the presence of strong inter dependencies, one cannot optimize a system by separately optimizing each element it is made of. Consider a problem that is made up of N elements and whose optimal solution is $x_1^* x_2^* \dots x_N^*$ while the current state is $x_1 x_2 \dots x_N$. In the presence of strong inter dependencies, it might well be the case that some or even all of the solutions of the $x_1 x_2 \dots x_i^* \dots x_N$ kind show a worse performance than the current one.¹²

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 that can be solved independently, 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.

¹⁰ In fact, one of the fundamental problems faced by human and artificial players is to build effective heuristics to evaluate the positions during the game without the knowledge of the entire tree.

¹¹ Similar aspects are present even in the simplest production technologies. Consider, for instance, team production as explored by Alchian and Demsetz (1972): two workers lifting a heavy load. Additional individual effort generally raises team production, but when the levels of effort applied by the two are disproportionate this might result in turning over the load, thus sharply decreasing team production.

¹² 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 the system's performance value.

It is important to remark that the introduction of any decentralized interaction mechanism, such as a competitive market for each component 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. Thus, interdependencies undermine the effectiveness of the selection process as a device for adaptive optimization and they introduce forms of path-dependency with lock-in into sub-optimal states that do not originate from the frictions and costs connected to the selection mechanism, but from the internal structure of the entities undergoing selection.

As Simon (1969) pointed out, since an optimal decomposition (i.e. a decomposition that divides into separate sub-problems all and only the elements that 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 normally be bound to design near-decompositions, that is decompositions that 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 system performance. However, near-decompositions involve a fundamental trade-off: on one hand, finer decompositions exploit the advantages of decentralized local adaptation, that is, the use of a selection mechanism for achieving coordination “for free” together with parallelism and adaptation speed. 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 adaptation speed), although they inevitably involve lock-in into sub-optimal solutions.

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 effectively drive 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 feedback from the environment concerning the fitness of the whole entity, only under conditions of near independence can the usual selection processes 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 reframing the same problem in a novel way. As shown by many case studies, major innovations often appear when various elements that were well known are re-combined and put together under a different perspective.¹³ Indeed, one can go as far as to say that it is the representation of a problem that determines its purported difficulty and that one of the fundamental functions of organizations is precisely to implement collective representations of the

¹³ An example is provided by Levinthal (1998), in his detailed account of the development of wireless communication technologies.

problems they face (Loasby, 2001). 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 suboptimal solutions can be avoided and better solutions can be discovered. Division of problem-solving labor is therefore very much a question of how the problem is represented.¹⁴ Needless to say, boundedly rational individuals cannot be innocently assumed to hold optimal representations.

Having given the foregoing qualitative intuitions, let us next develop a formal model that provides a precise measure of the above-mentioned tradeoffs.

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” that we generically call **components**, each of which can assume some number of alternative states. For the sake of 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 very easily be extended to the case of any finite number of states.

More precisely, we characterize a problem by the following elements:

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

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

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

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

In order to avoid technical complications, we assume, for the time being, that there exists only one configuration that is strictly preferred over 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 pair (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 decompose¹⁵ it into sub-spaces.

Let $I = \{1, 2, \dots, N\}$ be the set of indexes and let a **block**¹⁶ $d_i \subseteq I$ be a non-empty subset of it; we call the **size of block** d_i , its cardinality $|d_i|$. We define a **decomposition scheme**

¹⁴ A formal treatment of the properties of different representations in a particular class of problems can be found in Marengo (2003).

¹⁵ 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.

¹⁶ 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 that defines the set of modules into which a technological system or an organization are decomposed.

(or simply **decomposition**) of the problem (X, \succeq) a set of blocks:

$$D = \{d_1, d_2, \dots, d_k\}.$$

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 , 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 :

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

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:

$$x(d_j) \vee y(d_h) = z(d_j \cup d_h) \text{ where } z_v = \begin{cases} x_v & \text{if } v \in d_j \\ y_v & \text{otherwise} \end{cases}$$

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:

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

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.¹⁷

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

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

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.

¹⁷ As a first approximation, we assume that this sorting and selection mechanism is errorless and operates at no cost and without any friction.

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 that 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. 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 that 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 that 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 $D = \{d_1, d_2, \dots, d_k\}$, 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. $x^i \succeq x^j$,
2. $x^i_v = x^j_v \forall v \notin d_h$,
3. $x^i \neq x^j$.

Conditions 2 and 3 require that the two configurations differ only by components that 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:

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

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, in 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 optimum¹⁸ and let $Z \subseteq X$ with $x^0 \in Z$. We say that the problem (X, \succeq) is locally decomposable in Z by scheme D if $Z \subseteq \Psi(x^0, D)$. If $Z = X$, we say that the problem is globally decomposable by scheme D .¹⁹

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 it is the degenerate decomposition $D = \{\{1, 2, 3, \dots, N\}\}$ 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 irrespectively 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).

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

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

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 that is not the finest one, then finer decomposition schemes can usually be introduced that can reach the same local optimum more quickly.

¹⁸ We recall the assumption of uniqueness of the global optimum.

¹⁹ 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 $x^i \in X$.

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 component: $x_k^i(d_h) \neq x_k^j(d_h) \forall k = 1, 2, \dots, |d_h|$.

Proposition 2. Let $\Psi(x^\alpha, D) = \{x^\alpha, x^{\alpha+1}, \dots, x^{\alpha+m}\}$ be the ordered basin of attraction of a local optimum x^α (with $x^{\alpha+j} \succeq x^{\alpha+j+1} \forall j = 0, \dots, m - 1$). Define $\Psi^i(x^\alpha, D) = \Psi(x^\alpha, D) \setminus \{x^{\alpha+i+1}, x^{\alpha+i+2}, \dots, x^{\alpha+m}\}$ for $0 \leq i \leq m$. Let $d_{vi} \in D$ be the maximum size block(s) in D . Then, unless x^α and $x^{\alpha+1}$ totally differ for some maximum size block d_{vi} , 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 $d_v \in D$ with $|D| = |d_v|$. 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_v , then there exists a smaller decomposition D^i that is identical to D except that its largest block d_v can be split into two sub-blocks containing respectively the components for which $x^\alpha(d_v)$ and $x^{\alpha+1}(d_v)$ 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 $d_{vi} \in D$, it necessarily follows that x^α and $x^{\alpha+1}$ cannot totally differ for any of them. \square

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 that preserves optimality with certainty. Conversely, note that for decompositions which are finer than those of minimum size, whether or not 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.²⁰

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

The minimum size decomposition can be computed as follows:

1. start with the finest decomposition $D^0 = \{\{1\}, \{2\}, \dots, \{N\}\}$,
2. check whether or not $x^0 \in P\{x^i, D\} \forall x^i \ i = 1, 2, \dots, 2^N - 1$. If there is a path leading to the global optimum from every other configuration for decomposition D ; STOP,
3. if not, 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. A hypothetical ranking (where 1 is the rank of the most preferred) of configurations for $N = 3$

²⁰ 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.

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

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

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

Note that the worst local optimum has a larger basin of attraction²¹ because it covers all configurations except the global optimum itself. Thus, only a search that starts at the global optimum will (trivially) stop at the global optimum with certainty, while for the other initial configurations, a 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.

By using the notion of dominance (see footnote 19 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 decentralized search structure can 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 such case 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 that far looked for perfect decomposability, in the sense that we require that all blocks be optimized in a way

²¹ Kauffman 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 between components. In particular, he finds that as the span of interdependencies increases, the number of local optima also increases while the size of the basin of attraction of the global optimum shrinks.




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

totally independent from the others. In this way, we are sure to decompose the problem into perfectly isolated components that can be independently solved. This is, however, a very stringent 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, Fig. 1 shows that in Kauffman's NK random landscapes²² the above-described algorithm finds only decomposition schemes of size N or just below N even for very small values of K (that is, for highly correlated landscapes). In other words, a little bit of interdependence spread across the set of components immediately makes a system practically indecomposable.

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 that fully contain all interdependencies) but we might be content to find sub-problems that 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.²³ In other words, we construct **near-decompositions** that give a precise measure of the trade-off between

²² An 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 ex-post complexity measure (in terms of its decomposability) of the optimization problem on the resulting fitness landscape; small values of K usually generate landscapes that are not decomposable, but on the other hand, it is always possible that, even with very high values of K , the resulting landscape is highly decomposable.

²³ This procedure can also deal with the case of multiple global optima and thus we can now also drop the assumption of a unique global optimum.

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 rearrange all the configurations in X by descending rank $X = \{x^0, x^1, \dots, x^{2^N-1}\}$ where $x^i \succeq x^{i+1}$, and let $X_\mu = \{x^0, x^1, \dots, x^{\mu-1}\}$ with $0 \leq \mu \leq 2^N - 1$ 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 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 decomposition schemes for optimal decompositions.

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

Proposition 3. *Let D_μ be 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 $1 \leq |D_\mu| \leq N$. We still have to show that $|D_{\mu+1}| > |D_\mu|$ cannot happen: if this were 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 . \square

The latter proposition shows that higher degrees of decomposition and decentralization can be attained by giving up optimality and it 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 decomposable²⁴ (i.e. $|D| = 12$ for all of them). Fig. 2 displays the average size of the minimum size decomposition schemes for the 100 random problems as we vary the number μ of acceptable configurations. Fig. 2 shows that second-best solutions can be reached by search processes based upon finer decompositions (that is with more decentralized processes) that can find such solutions more quickly by exploiting coordination “for free” provided by the selection mechanism. In fact, when the size of the decomposition scheme decreases by one unit, the expected search time decreases by half.

²⁴ 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 minimum size decomposition schemes. Only those problems for which the size of the smallest decomposition schemes was 12 were used in these simulations. An alternative (and equivalent) method is to generate random NK landscapes *à la* Kauffman with $N = 12$ and a high K and then check that the resulting landscape is not decomposable, as it may happen that landscapes with a very high K may also be highly decomposable.

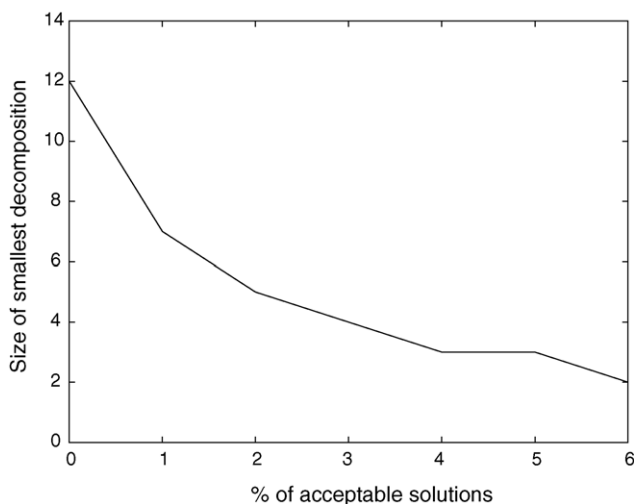


Fig. 2. Near decomposability.

7. Search speed and optimality: some consequences for organizational structures

The trade-offs outlined in the previous sections between decomposability, complexity reduction and search speed on one hand and asymptotic optimality on the other enable 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 Fig. 2 for randomly generated problems,²⁵ if second-best solutions are accepted we can reduce considerably the size of the decomposition schemes and the expected search time. 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 that are more decomposed and decentralized than what would be optimal given the interdependencies of the problem space. This property is shown in Figs. 3 and 4, that present the typical search paths on a non-decomposable problem of two search processes driven, respectively, by decompositions:

$$D1 = \{1, 2, \dots, 12\}$$

$$D12 = \{\{1\}, \{2\}, \dots, \{12\}\}$$

Fig. 3 shows the first 180 iterations in which the more decentralized structure (D12) quickly climbs the problem space and outperforms a search based on a coarser decomposition. If there were a tight selection environment, a more-than-optimally-decentralized

²⁵ In this figure and the following (with the exception of Fig. 5), we indicate on the vertical axis the rank of configurations re-parametrized between 0 (worst) and 1 (best).




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




Fig. 4. . . . After 3000 iterations.

organizational structure would quickly displace structure D1 that reflects the “true” decomposition of the underlying problem space.

However, the search process based on the finest decomposition quickly reaches a local optimum from where no further improvement can occur, while the process based on the coarser decomposition keeps searching and climbing slowly.²⁶ Fig. 4 shows iterations

²⁶ If we imagine that each iteration corresponds to a day, then the decentralized organization achieves a good fit with its environment within 2 or 3 weeks, while the centralized structure requires some years.




Fig. 5. Diversity of surviving decomposition schemes.

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 therefore tends 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 as “modular”, those characterized by blocks with strong interdependencies within blocks and much weaker (but non-zero) interdependencies between blocks; in these 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 generally 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 into.²⁷ This is easily shown by a simulation exercise where 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. The 10 worst performing organizations are selected out every 10 iterations and replaced by 10 new organizations where 5 are randomly generated and 5 have the same decomposition scheme of the best performing ones but are placed on a randomly chosen configuration.

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

²⁷ On this latter point, a similar result is obtained by Levinthal (1997).

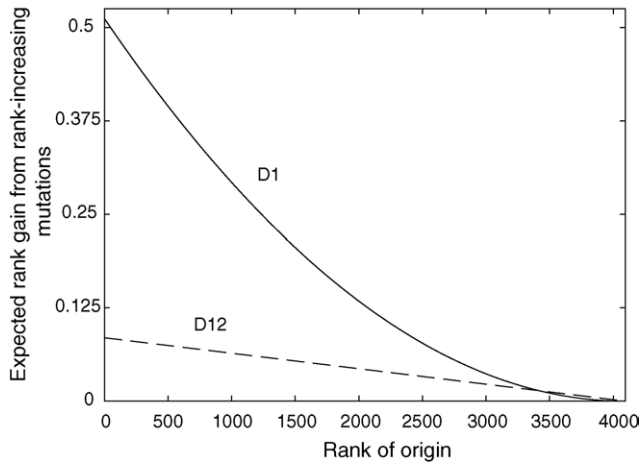


Fig. 6. Expected gain from rank-improving mutations for the finest (D12) and coarsest (D1) decompositions in a fully decomposable problem ($N=12$).

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

We have also run other simulations where, at given intervals, we have changed the current problem with one having exactly the same structure in terms of decomposability but different randomly generated orderings of configurations. This can be taken as a metaphorical proxy for environment volatility. For instance, consumers might have changing preferences for 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 ranks becomes more frequent, the population is entirely invaded by organizations characterized by coarser and coarser decompositions and, at the limit, by organizations that do not decompose at all. This robustly suggests that growing volatility has stronger consequences than does growing interdependence. The reason that this happens is shown in Figs. 6 and 7 that present, respectively, the expected improvements and the probability of improvement for searches based on the finest (D12) and coarsest (D1) decomposition schemes in a fully decomposable problem.²⁸ It is shown that, when starting from low rank configurations, a search based upon coarser decomposition has 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 because finer decompositions search only locally, and this, on average, cannot produce large improvements in fully-decomposable problems. 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 recover promptly.

²⁸ Figs. 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.




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

8. Conclusions

In this paper, we have presented a novel model of ‘Simonian’ ascendancy concerning the properties of the division of problem-solving labor and we have accounted for the properties of different institutional arrangements and in particular for different boundaries between undecomposed (in principle, organization-embodied) tasks and decomposed tasks (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? This paper, in fact, shows that this is only possible under some special and rather implausible conditions and that, on the contrary, the advantages of decentralization usually bear 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 . . . [that] appear with regularity also in human social organizations—e.g. business firms and government agencies—with their many-layered hierarchies of divisions, departments and sections” (Simon, 2002, pp. 598–599). However, our model also highlights the subtle trade-offs between decomposition levels, degrees of suboptimality of the achievable outcomes, and adaptation speed. Together, it casts strong doubts on the general validity of any “optimality-through-selection” argument in the sphere of organizational and technological designs and, more in general, of any “optimistic” view of market selection processes that can be found, for instance, in Alchian (1950) and Friedman (1953), as forces that substitute individual optimization with evolutionary optimization.

On more empirical grounds, the analysis of the foregoing trade-offs can provide a plausible mechanism that is compatible with the observed changing depth and profile of the integration of organizations along technology and industry life cycles. The results of our model are fully consistent with vast empirical evidence suggesting that new technologies develop in highly integrated organizations because of the need to control the strong interdependencies that 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 cost factors can very well be responsible at this stage for variations in the degree of integration), but the more that decentralization is pushed forward, the more unlikely it will be that new and better local optima can be discovered. There is an inevitable trade-off between decentralization and optimality that can hardly be avoided.

Finally, we suggested that organizations could actually play the even more fundamental role of building collective representations of the problems to be 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 of coordination in the presence of any form of division of cognitive and productive labor.

Acknowledgements

We are greatly indebted to Marco Valente for his comments on the subject and for his help in building and running the simulations reported in Section 7. We also thank Esben Andersen, Massimo Egidi, Koen Frenken, Yuri Kaniovski, Thorbjørn Knudsen, Scott Page, Corrado Pasquali, the late Herbert Simon and two anonymous referees for their useful remarks and suggestions. G.D. thanks the St. Anna School of Advanced Studies (project no. E6099GD) and the Italian Ministry of University and Research (project no. 2002.132415). L.M. gratefully acknowledges financial contributions from the NORMEC project (SERD-2000-00316), funded by the European Commission, Research Directorate, 5th Framework Programme. The authors are solely responsible for the opinions expressed in this paper, which are not necessarily those of the European Union or the European Commission.

References

- Alchian, A., 1950. Uncertainty, evolution and economic theory. *Journal of Political Economy* 58, 211–222.
- Alchian, A., Demsetz, H., 1972. Production, information costs and economic organization. *American Economic Review* 62, 777–795.
- Aoki, M., 2001. *Towards a Comparative Institutional Analysis*. MIT Press, Cambridge, Mass.
- Baldwin, C.Y., Clark, K.B., 2000. *Design Rules. The Power of Modularity*. MIT Press, Cambridge Mass.

- Becker, G.S., Murphy, K.M., 1992. The division of labour, coordination costs and knowledge. *Quarterly Journal of Economics* 107, 1138–1140.
- Bowles, S., 1985. The production process in a competitive economy: Neo-hobbesian and Marxian models. *American Economic Review* 75, 16–36.
- Coase, R.H., 1937. The nature of the firm. *Economica* 4, 386–405.
- Cohen, M.D., Burkhardt, R., Dosi, G., Egidi, M., Marengo, L., Warglien, M., Winter, S.G., 1996. Routines and other recurring action patterns of organizations: contemporary research issues. *Industrial and Corporate Change* 5, 653–699.
- Cyert, R.M., 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, 205–220.
- David, P.A., Greenstein, S., 1990. The economics of compatibility standards: an introduction to recent research. *Economics of Innovation and New Technology* 1, 3–41.
- Dosi, G., Nelson, R.R., Winter, S.G., 2000. *The Nature and Dynamics 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, 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, 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, 145–181.
- Langlois, R.N., 1986. *Economics as a Process: Essays in the New Institutional Economics*. Cambridge University Press, Cambridge.
- Langlois, R.N., Robertson, P.L., 1995. *Firms, Markets and Economic Change*. Routledge, London/New York.
- Leijonhufvud, A., 1986. Capitalism and the factory system, in economics as a process: essays in the new institutional economics (Langlois 1986), 203–223.
- 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, 217–247.
- Loasby, B.J., 1998. The organization of capabilities. *Journal of Economic Behavior and Organization* 25, 139–160.
- Loasby, B.J., 2001. Organizations as interpretative systems. *Revue d’Économie Industrielle* 97, 17–34.
- March, J.G., 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*. Routledge, Dordrecht, pp. 371–388.
- Marglin, S.A., 1974. What do bosses do? The origins and function of hierarchy in capitalist production. *Review of Radical Political Economics* 6, 33–60.
- Marschak, J., Radner, R., 1972. *Economic Theory of Teams*. Yale University Press, New Haven.
- Milgrom, P., Roberts, J., 1990. The economics of modern manufacturing. *American Economic Review* 80, 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, 1029–1064.
- Nelson, R.R., 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, 227–246.
- Page, S.E., 1996. Two measures of difficulty. *Economic Theory* 8, 321–346.
- Penrose, E., 1959. *The Theory of the Growth of the Firm*. Wiley, New York.
- Perry, M.K., 1989. Vertical integration: determinants and effects. In: Schmalensee, R., Willig, R. (Eds.), *Handbook of Industrial Organization*, vol. I. North Holland, Amsterdam, pp. 185–255.
- Radner, R., 2000. Costly and bounded rationality in individual and team decision making. *Industrial and Corporate Change* 9, 623–658.
- Richardson, G.B., 1972. The organization of industry. *Economic Journal* 82, 883–896.

- Riordan, M., Williamson, O.E., 1985. Asset specificity and economic organization. *International Journal of Industrial Organization* 3, 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, 587–599.
- Teece, D.J., Pisano, G., Shuen, A., 1997. Dynamic capabilities and strategic management. *Strategic Management Journal* 18, 509–533.
- Teece, D.J., Rumelt, R., Dosi, G., Winter, S.G., 1994. Understanding corporate coherence: theory and evidence. *Journal of Economic Behavior and Organization* 23, 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., Groves, T. (Eds.), *Adaptive Economic Models*. Academic Press, New York, pp. 73–118.
- Winter, S.G., 2004. Toward an evolutionary theory of production. In: Dopfer, K. (Ed.), *Principles of Evolutionary Economics*. Cambridge University Press, Cambridge, pp. 223–254.