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

Via Carducci, 40 - I-56127 PISA (Italy)
Tel. +39-050-883-341 Fax +39-050-883-344
Email: lem@sssup.it Web Page: http://lem.sssup.it

LEM
Working Paper Series

The structure of problem-solving knowledge and the structure of organisations

Luigi Marengo [*]
Giovanni Dosi [†]
Paolo Legrenzi [°]
and
Corrado Pasquali [*]

[*] *University of Trento, Trento, Italy*
[†] *St.Anna School of Advanced Studies, Pisa, Italy*
[°] *State University of Milan, Italy*

1999/09 **July 1999**

ISSN (online) 2284-0400

Final Draft **July 2000**

THE STRUCTURE OF PROBLEM-SOLVING KNOWLEDGE AND THE STRUCTURE OF ORGANISATIONS

Luigi Marengo*

Giovanni Dosi**

Paolo Legrenzi***

Corrado Pasquali****

* University of Trento, Italy, lmarengo@gelso.unitn.it

** St. Anna School of Advanced Studies, Pisa, Italy, gdosi@sssup.it

*** State University of Milan, Italy

**** University of Trento, Italy, corrado@black.gelso.unitn.it

We want to thank Marco Valente for very valuable discussions and help in running the simulations. The participants in the SCANCOR conference "The Roots and Branches of Organizational Economics", Stanford University, September 1999, and in particular Jim March and Roy Radner, and two anonymous referees have offered insightful comments. The usual disclaimer applies.

We gratefully acknowledge financial support by the Dynacom Project (European Union, TSER, DGXII) and by the Italian Ministry of University and Scientific and Technological Research (research project “Bounded Rationality and Learning in the Experimental Approach”).

1. Introduction

This work is meant to contribute to the analysis of organisations as repositories of problem-solving knowledge and of the ways the latter co-evolves with governance arrangements. These themes find their roots in at least four complementary streams of literature, namely, first, the studies – of "Simonian" ascendancy – of problem-solving activities, and of the structure of knowledge they entail (cf. Simon, 1981); second, diverse investigations on comparative performance of diverse organisational architectures and related patterns of distribution of information and division of labour (cf. Radner, 1992 and Aoki, 1988); third, evolutionary theories of the firm, in particular with their emphasis on the knowledge content of organisational routines (Nelson and Winter, 1982, Cohen et al., 1996); fourth, and largely overlapping, competence-based views of the firm (Winter, 1988, Dosi, Nelson and Winter, 2000).

Broadly in tune both with analyses of organisations, mainly inspired by Herbert Simon, as sort of imperfect problem-solving arrangements as well as a good deal of evolutionary theories of the firm, we censor in a first approximation any explicit incentive compatibility issue among organisational members. Rather, we focus upon the ways different patterns of division of labour shape and constrain search processes in high dimensional problem spaces. Examples of such search processes are all those problems requiring the coordination of a large number of interdependent "elements" whose functional relationships are, to a good degree, opaque to the organisational members themselves[1].

Here by "elements" we mean elementary physical acts – such as moving one piece of iron from one place to another – and elementary cognitive acts – such as applying inference rules. Relatedly, problem-solving can be straightforwardly understood as combination of elementary acts leading to a feasible outcome (e.g. the design and production of an engine, the discovery and testing of a chemical compound, etc.).

In this perspective, we present a quite general formal framework enabling the exploration of the problem-solving properties of diverse patterns of division of labour and routine-clustering practices, ranging over a continuum that notionally spans from totally decentralised market-like mechanisms to fully decentralised coordination processes. Not surprisingly, the complexity of the problem-solving tasks bears upon the performance outcomes of different organisational arrangements. In a broad and somewhat impressionistic definition, which we shall refine below, by “complex problems” we denote high dimensional problems whose solution requires the coordination of interdependent components whose functional relations are only partly understood. Designing a

[1] By that censorship of the double nature of organisations as both problem-solvers and mechanisms of governance of potentially conflicting interests we clearly fall short of the "grand research program" sketched in Coriat and Dosi (1998) whereby evolutionary and competence-based theories of the firm begin to take on board incentive alignment issues. In the present "first approximation", however, we feel well justified by the still rudimentary state of knowledge-centred investigations of organisational arrangements, especially when compared with nearly pathological theoretical refinements on hyper-rational incentive compatible schemes that no one will ever observe on earth.

complex artefact, establishing a sequence of moves in a game or designing a multi-agent organisational structure are all instances of "complex problems".

Consider the case of designing an aircraft. This will require the coordination of a large number of different elements such as engine type and power, wing size and shape, materials used, etc. whose interactions have to be tested through simulations, prototype building and maybe some other forms of trial and error processes. At the end, an “effective” solution to the problem (i.e. a properly flying aircraft) will be one in which a large set of traits and characteristics have been coordinated in ways which turn out to be compatible with each other. Note that, for instance, adding a more powerful engine might imply a decreasing overall performance if other components (wings, etc.) are not properly adjusted and tuned with that change. Playing chess is not too different a case: a winning solution is a long sequence of moves each of which is chosen out of a set of possibilities large enough to make an extensive search unfeasible for any boundedly rational agent. Even in this case, the key point is the opaqueness of the relations among such moves in the sequence (a notionally optimal strategy might involve, for instance, castling at a given time of the game but the same castling, as a part of some sub-optimal, but otherwise effective, strategy, could turn out to be a losing move...).

Organisations such as business firms generally face a similar class of problems. Indeed, they can be represented as complex multi-dimensional bundles of routines, decision rules, procedures, incentive schemes whose interplay can be hardly assumed to be perfectly known also to those who manage the very same organisation (witness all the problems, unforeseen consequences and unexpected feed-backs emerging whenever managers try to promote organisational changes: cf. March and Simon (1993) for a classic treatment of the subject). So, for example, introducing some routines, practices, or incentive schemes which have proven superior in another organisational context, could prove harmful in a context where other elements are not appropriately tuned (more on this issue, from different angles, in the chapters by Fujimoto, Coriat and Levinthal in Dosi, Nelson and Winter (2000)).

The main underlying issue is that functional relations among components (e.g. elementary cognitive and practical acts) are only partly understood while the contribution of each of the component to the overall solution depends, to various degrees, on the state assumed by other components: hence the need to explore a large combinatorial space of components possibly by a trial and error process. At the same time, these very characteristics of problem-solving search tend to jeopardise a fruitful use of global information - i.e. information which derives from some global performance measurement – for "good" adjustments at a local level.

The problem here is the one of whether and under what conditions it is possible to achieve optimal or nearly optimal solutions through small local and incremental adjustments. This problem has long interested especially computer scientists (all search algorithms face this kind of problems) and biologists (the very small rates of mutation observed in the biological realm can justify models of evolution whereby only one gene at a time gets mutated). Biologists in particular have produced some easy and highly suggestive models (see in particular Kauffman's "NK model" of fitness landscapes, cf. Kauffman (1993)) which show that if the entities subject to evolutionary selection pressure are somewhat "complex" entities, i.e. made up of many non-linearly interacting components, then local incremental search combined with selection generally to highly sub-optimal and path-dependent evolutionary paths. These biological models seem to have many relevant implications also when the "complex" entities under scrutiny are social organisations such as business firm (cf. Levinthal (1997)).

This paper can be considered as a generalisation of Kauffman's argument as it presents a model which determines the extent to which a problem space can be decomposed into smaller sub-problems which can be solved independently or quasi independently without affecting (or affecting only within given limits) the possibility of finding optimal or at least good solutions.

In the following, we shall first introduce some technical notions on problem-solving and organisations as problem-solving arrangements (section 2). Section 3 presents our basic model and discusses a few of its generic properties, while section 4 provides some examples of different decomposition schemes and related solution patterns. Finally, in section 5 we draw the main conclusions and outline some possible directions for further research.

2. On the nature of organisational problem-solving: some introductory notes

In problems whose solution involves the exploration of high dimensional spaces, agents endowed with limited computational capabilities and with a limited knowledge of the interdependencies can explore only a subset, possibly very small, of solutions. Even if we assume that the selection mechanism which selects among alternative solution works perfectly (i.e. without delays, inertia, errors or "transaction costs" as economists would say), the outcome of selection is bound by the set of solutions which are produced by the generative mechanism. It may well be the case that optimal or even “good” solutions will never be generated at all and thus that they will never be selected by any selection mechanism whatsoever.

The problem here is that strong interdependencies create a large set of local optima in the search space. Marginal contributions rapidly switch from positive to negative values, depending on which value is assumed by other components. As a consequence, the presence of strong interdependencies prevents the possibility of reaching optimal solutions by simply adopting an optimal value for each of the components a problem is made of. It is thus possible that, given a n-dimensional problem whose current state is $a_1, \ldots, a_n$ and whose optimal solution is $a^*_1, \ldots, a^*_n$ some or even all of the solutions of the form $a_1, \ldots a^*_i, \ldots a_n$ have a worse performance than the currently adopted one. Also note that if each of the a_i's was traded on a competitive market with prices reflecting their revealed productivity, notionally superior resources a^*_i would never be hired as their marginal productivity is negative. As a consequence - that we will largely expand upon in the following sections - it might well be the case that the optimal solution will never be generated and thus never selected.

The issue of interdependencies and of how they shape search processes in a space of solutions is also faced in Kauffman's NK model of selection dynamics in biological domains with heterogeneous interdependent traits.

NK models the evolution of systems composed of a number of elements which locally and interdependently contribute to the global features of the system they collectively constitute. The model was originally intended to capture the evolutionary dynamics of organisms (i.e. systems) as described by sets of genes (i.e. elements). The formal structure of the model and the idea behind it, are however general enough to allow its application to realms and domains different from molecular biology (cf. Levinthal (1997), Westhoff *et al.* (1996)).

In NK's terminology, a system is described by a string composed of different *loci* each referring to one of the elements that compose the system. In the aforementioned aircraft example, for instance, we might well imagine the aircraft represented by a string in which each *locus* refers to one of the aircraft elements (i.e. wings, engine, body…).

The whole thing can be shortly explained by referring to the N and the K in the name of the model.

The N in the name NK refers to the number of elements or *loci* one is considering, that is: to the dimension of the problem at hand. Each element can assume one out of a set of different states (called *alleles* in biology). That is: each element can be assigned a value representing, for instance, a specific feature being present or not or, which specific shape or feature one chooses for a given element. The number of possible strings (i.e. of different configurations a system can be in) is called the *possibility space* of a system.

The K, on the other hand, refers to the number of interdependencies between different elements. These are usually called *epistatic correlations* and they do describe the inner structure of the system in terms of the number of elements that each *locus* is interdependent with. In particular, K describes how the contribution of each element to the system is dependent not only on its own value, but on other elements' values as well. A system whose K value is 2, for instance, is a system in which each element contribution is dependent on the values assumed by two other elements. The two limit cases are that of K being equal to 0 (i.e. each element contribution is solely dependent on its own state and there are no interdependencies in the system) and that of K being equal to N-1 (i.e. each element contribution is dependent on the state assumed by every other element in the system).

We can then easily imagine that each state assumed by a system (i.e. each assignment of values to its elements) is assigned a measure of, say, effectiveness with respect to a given task; let us call this a fitness measure. Now, the distribution of fitness values of all possible states is called a *fitness landscape*. A fitness landscape is thus a way of mapping a system' states onto their relative fitness values and it constitutes a representation of a possibility space along with the fitness values of each possible string in it.

The point we'll be interested in is the exploration of a fitness landscape, that is: the search for better (i.e. fitter) configurations in the space of all possible configurations.

Given that a system's actual configuration is a string in which every *locus* has been assigned a value, the way to explore a landscape and to test the fitness values of other configurations corresponds to changing some elements' values thus moving from one configuration to a new one and, consequently, from one point of the landscape to a new (and possibly higher) one.

The value assumed by K is a key point with respect to the shape of a fitness landscape and, consequently, to its exploration. Indeed, being K=0 the contribution of each element to overall fitness will be independent from every other element and, consequently, a change in a single element's value will leave unchanged the contribution of all the other N-1 elements. It then follows that the whole landscape will look very "smooth" and configurations that are similar with respect to values assumed by their elements will also have similar fitness values. The highly correlated structure of the fitness landscape can be effectively exploited by local and incremental search processes. On the contrary, as K increases the landscape will be increasingly rugged and points that are close in the landscape will no longer have similar fitness values.

What is most relevant to our point is the fact that the more a landscape is rugged the less (locally) informative is its exploration and the less is the degree of "correlatedness" of different configurations.

However, Kauffman's approach to the exploration of a fitness landscape does not necessarily fits well with the realm of social evolution. The main reason for this inadequacy is that social actors might well explore a fitness landscape by the application of a far richer class of algorithms than one-bit mutational ones studied by Kauffman and grounded on the laws of genetics. Actually, a social agent (be it an individual or an organisation) can possibly adopt many kinds of problem solving strategy and search algorithms of virtually any cardinality. The very notion of locality and neighbourhood search is not clearly defined in social realms but is itself a product of how individuals and organisations represent the problems they try to solve.

The relevance of this point will be made evident by the following considerations. The notion of fitness landscape is indeed centred upon two ideas: a function that assigns a fitness value to each element of the space of configurations and a metric defined in that space which reflects a measure of distance between two different configurations. Once an algorithm is defined that transforms a configuration into another one, the notion of distance between two configurations is defined as the minimum number of applications of the algorithm needed in order to transform a configuration into the other. In this way the set of neighbours of a configuration is defined in terms of what can be reached from it. In the case of mutational algorithms of cardinality one, the set of neighbours of a configuration is defined as the set of configurations that are a single step away from it[2]. It is then evident that a change in the search algorithm will result in a change in the geometry of the landscape: a landscape might be extremely rugged when defined on one-bit mutations but very smooth when defined on mutations of higher cardinality.

[2] "Single step" is a *naive* formulation of the notion of unitary Hamming distance. Hamming distance denotes the number of *loci* at which corresponding bits take on different values.

We have thus considered possible ways of exploring a fitness landscape other and more general than single-bit mutations. In particular, following Simon (1981), our focus is on those problem solving strategies which decompose a large problem into a set of smaller sub-problems that can be treated independently, by promoting what might be called a division of problem solving labour. Imagine, for instance, a N-dimensional problem: i.e. a problem which is composed of N different bits or traits. A one-bit mutational algorithm correspond to a maximally decentralized decomposition in which each of the N bits will be given a value independently of all the other N-1 bits. In this sense the whole problem will be decomposed into the smallest possible sub-problems. This search strategy is a potentially very quick one as it will require at most nN steps, given that each of the N components can be in one out of n possible states. Unfortunately, as already hinted and as showed by Kauffman, the existence of interdependencies among components will prevent the system to reach the global optimum (i.e. the optimal solution) but only the local optimum whose basin of attraction contains the initial configuration of the problem. On the other hand, the same problem could be left totally undecomposed and a search algorithm might be adopted which explores all the N dimensions of the problem. This strategy corresponds to mutating up to all the N components of the problem and it leads with certainty to the global optimum by examining all the n^N possible configurations. In between there are all the other possible decomposition of the problem, each corresponding to a different division of labour.

According to this view, the division of problem solving labour to a large extent determines which solutions will be possible to generate and then select. As we have suggested, decompositions are necessary in order to reduce the dimension of the search space, but, at the very same time, they also constrain search processes to a specific sub-space of possible solutions thus making it possible for optimal solutions not to be ever generated and for systems to be locked into sub-optimal solutions. We thus believe that our arguments cast some serious doubts on any view of markets relying on a “optimality-through-selection” assumption or on any view which assumes market forces as capable of substituting individual rationality with evolutionary optimisation. At the same time, our route to the analysis of organisational forms, is in some important respects the symmetric opposite to the one explored by a good deal of contemporary accounts of agency problems, focussing, as they do, on the identification of efficient incentive mechanisms for coordination on the implicit but momentous assumption that optimal problem solving governance structures and optimal search heuristics are naturally in place from the outset.

Of course, only a planner who has perfect knowledge of all the interdependencies, a Laplacian God for instance, could implement a perfect set of signals and incentives mechanisms. Limiting ourselves to more earthly and Simonian rather than Laplacian situations, we investigate the case of boundedly rational and computationally limited problem solvers who are forced to decompose problems which are not fully decomposable by separating only the more “fitness-relevant” interdependencies, while other still persist across subproblems. By separating interdependent elements into different sub-problems, a problem whose complexity is far beyond available computational resources is reduced to smaller subproblems which can be handled, ruling away in any case the possibility of reaching with certainty the optimal solution. According to this softening of the requirements on decompositions' grain, we will show how “near-decompositions” can be constructed which, when considered as satisficing rather than optimal structures, can significantly improve search times and attain a higher degree of decentralization at the expenses of optimality of the solutions. In order to study some evolutionary properties of populations of agents endowed with different decompositions of the same problem, we also run a class of simulations, whose results will be presented and discussed in section 4.

Our last point is concerned with representational issues related with problem solving. Up to now, we have only considered that a problem representation is exogenously given to agents. Actually, amore cognition-oriented perspective might be adopted that amounts to considering how agents trying to solving a problem can derive much of their possibilities of success from the adoption of different and subjective perceptions of the very same problem. A reasonable hypothesis is that agents' search might well take place not on the “objective” landscape, but in a space constructed as a subjective representation of it. In this sense, changing representations might turn out to be an even more powerful problem solving strategy than decomposition.

According to this more cognitive-oriented view, we imagine agents' representations of a problem to be actually grounded on the objective landscape but, notwithstanding that, to be a simplified version of it (possibly of lower dimensionality). As noted by Gavetti and Levinthal (2000), cognitive representations tend to be simplified caricatures of the actual landscape as a result of their lower (perceived) dimensionality which, in turn, also results in a reduction of its apparent (perceived) degree of connectivity. In particular, we show that both acting on the encoding of a problem and on the ordering relation on solutions, every problem can be transformed in one of minimal complexity.

One of the fundamental aspects of problem solving procedures concerns representations of the problems itself. As a couple of us argued in more detail elsewhere (Dosi, Marengo and Fagiolo (2000)) – well in tune with a vast literature from cognitive science – a crucial step concerns the very processes through which individual agents and structured collectives of them i.e. organisations interpret the environment wherein they operate by means of inevitably imperfect and possibly fuzzy cognitive categories, causal links amongst environmental variables, conjectures on action/payoff relations, etc..

In the language of the analysis below, all this is captured through subjective decompositions and thus also subjective landscapes of problem solving search which might only bear vague similarities (if at all) with the “true” structure of the problem – as a God-like observer would be able to see it. Indeed section 3 that follows largely takes the latter God-like perspective in order to identify some general characteristics of problem solving, in terms of e.g. task decomposability, dimensionality of search spaces, attainability of optimal solutions, properties of search landscape, speed of convergence. In section 4 we begin to progressively “come down to earth” and account for systematic gaps between subjective and “true” (God-like) representations of the search landscapes.

An important intermediate step of the analysis focuses upon problem-solving set-ups characterised by cognitively bounded representations which nonetheless maintain some isomorphism with the "objective" search landscape. Broadly in line with Gavetti and Levinthal (2000), we begin by assuming agents' representations of a problem to be grounded on the "true" landscape, although as a simplified version of it, of lower dimensionality, entailing also a reduction of its perceived degree of connectivity amongst subproblems/elementary tasks.

The subsequent and even more difficult analytical challenge ought to push much further the exploration of systematic gaps between actual representations/search landscapes for both individuals and organisations, on the one hand, and this God-like nature. Some conjectures are put forward in section 5. Remarkably, our model shows that, with striking generality, operating on both the encoding of a problem and on the ordering relation on solutions, every problem can be transformed into one of minimal complexity. In a profound sense, the degrees of complexity of any problem are in the minds of the beholders. All that however implies stringent demands for the theorist in its descriptive mode to be disciplined by empirical evidence on the nature of cognitive structures, learning processes, patterns of organisational coordination etc.

3. The Basic Model: Problem Representation, Decomposition and Coordination

3.1 Basic definitions.

Let us assume that the solution of any given problem requires the coordination of N atomic elements, which we call generically components, each of which can assume a number of alternative states. For simplicity, we assume that each element can assume only two states, labelled 0 and 1. Note that all the procedures and results presented below can be very easily extended to the case of any finite number of states and also, with some complications, of numerable alternative states.

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

The set of **components**: $S=\{s_1,s_2,.....,s_N\}$, where $s_i \in \{0,1\}$

A **configuration** or possible solution to the problem is a string $x_i = s_1 s_2.....s_N$

The set of possible **configurations**: $X=\{x_1,x_2,.....,x_p\}$ where $p=2^N$

An **ordering** $\geq$ over the set of possible configurations[3]: we write $x_i \geq x_j$ whenever x_i is (weakly) preferred to x_j. For simplicity we assume that $\geq$ is anti-symmetric, i.e. that $x \geq y$ and $y \geq x$ implies $x=y$ (this assumption will be dropped later on.)

Thus, a **problem** is defined by the couple $(X, \geq)$.

As the size of the set of configurations is exponential in the number of components, whenever the latter is relatively large, the state space of the search problem becomes too vast to be

[3] Of course such an ordering could be substituted, where applicable, by a real-valued fitness function $F: X \rightarrow \Re$.

extensively searched by agents with bounded computational capabilities. One way of reducing its size it to decompose it into sub-spaces:

Let I={1,2,3,.......,N} be the set of indexes, and let a **block** d $\subseteq$ I be a non-empty subset of it, and let |d| be the **size of the block** d, (i.e. its cardinality).

In what follows, we shall confine ourselves to a particular class of search algorithms. Given that a search algorithm is a map which transforms a state into another according to some preliminary encoding (i.e. S: L $\rightarrow$ L), we shall consider algorithms which are:

a) **climbing:** (i.e. they move from a state to another if and only if this has a higher evaluation):

S(x) = x* iff x* ≥ x

S(x) = x otherwise.

b) **mutational:** (i.e. they can be characterised in terms of sets of bits they can mutate). So, given our index set I = {1, 2,..., N}, every such algorithm can be characterised by the subset of such positions it can mutate, that is: d $\in$ $\{2^I \setminus\varnothing\}$. According to this view, the size |d| of an algorithm can be imagined to be the number of bits it can mutate (i.e. the cardinality of its defining subset of indexes.)

We define a **decomposition scheme** (or simply **decomposition**) of the set X as a set of blocks:

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

A decomposition scheme is therefore a decomposition of the N dimensional space of configurations into sub-spaces of smaller dimension, whose union returns the entire problem (but they do not necessarily form a partition of it).

Given a configuration x_j and a block d_k, we call **block-configuration** $x_j(d_k)$ the substring of length $|d_k|$ containing the components of configuration x_j belonging to block d_k :

$$\boldsymbol{x_j(d_k) = s_{k_1} s_{k_2} .... s_{k_{|d_k|}}} \quad \text{for all } \boldsymbol{k_i \in d_k}$$

We also use the notation $x_j(d_{-k})$ to indicate the substring of length N-$|d_k|$ containing the components of configuration x_j not belonging to block d_k :

$$\boldsymbol{x_j(d_{-k}) = s_{k_1} s_{k_2} .... s_{k_{N-|d_k|}}} \quad \text{for all } \boldsymbol{k_i \notin d_k}$$

We can thus indicate the configuration x_j also by using the notation $x_j(d_k) \mid x_j(d_{-k})$

We define the **size of a decomposition scheme** as the size of its largest defining block:

$$sz(D) = \max \{|d_1|,|d_2|,....,|d_k|\}$$

A decomposition scheme and its size are important indicators of the complexity of the algorithm which is being employed to solve a problem:

1) a problem decomposed according to the scheme D={{1},{2},{3},.....,{N}} has been reduced to the union of sub-problems of minimum complexity, while a problem which has not been decomposed, i.e. whose decomposition scheme is D={{1,2,3,....., N}}, is a problem of maximum complexity because it can only be searched extensively (i.e. intuitively, there is no local feedback for search);

2) a problem which has been decomposed according to the scheme D={{1},{2},{3},.....,{N}} is being solved in linear time (in N steps), while a problem which has not been decomposed can only be solved in exponential time (in 2^N steps);

3) on the other hand, a problem which has not been decomposed can always been solved optimally (though in exponential time) while, as it will be shown below, a problem which has been decomposed according to the scheme D={{1},{2},{3},.....,{N}} - or for that matter according to any scheme whose size is smaller than N – can be solved optimally only under some special path-dependent conditions (which, as we will show, become more and more restrictive as the size of the decomposition scheme decreases).

Thus there is a trade-off between complexity and optimality for which we will provide a precise measure in the following.

For sake of theoretical simplicity, assume that coordination among blocks in a decomposition scheme takes place through decentralised market-like selection mechanisms which select at no cost and without any friction over alternative block-configurations.

In particular, assume that the current configuration is x_j and take block d_k with its current block-configuration $x_j(d_k)$. Consider now a new block-configuration $x_h(d_k) \neq x_j(d_k)$, if:

$$x_h(d_k) | x_j(d_{-k}) \geq x_j(d_k) | x_j(d_{-k})$$

then $x_h(d_k)$ is selected and the new configuration $x_h(d_k) | x_j(d_{-k})$ is retained in the place of x_j, otherwise $x_h(d_k) | x_j(d_{-k})$ is discarded and x_j is retained.

It might help to think in terms of a given structure of division of labour (the decomposition scheme), with firms or workers specialised in the various segments of the production process (a single block) and competing in a market which selects those firms or workers whose characteristics give the highest contribution to the overall production process.

On the other hand the coordination of components held together in a block has an implicit cost, in terms of search speed growing exponentially, with the size of the block.

We can now analyse the properties of decomposition schemes in terms of their capacities to generate and select better configurations.

3.2 Selection and search paths.

A decomposition scheme is a sort of template which determines how new configurations are generated and tested. In large search spaces in which only a very small subset of all possible configurations can be tested, 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 neighbourhood of the given one, where neighbours are new configurations obtained by changing some (possibly all) components within a given block.

Given a decomposition scheme $D=\{d_1,d_2,\ldots,d_i,\ldots,d_k\}$, we define, following Marengo (1999):

A configuration $x^i=s_1s_2.....s_N$ is a (**preferred-**) **neighbour** of configuration $x^j=z_1z_2.....z_N$ for a block $d_h\in D$ iff:

1) $x^i \geq x^j$ and

2) $s_v=z_v$ if $v\notin d_h$

i.e. the two configurations can differ only by components which belong to the block d_k.

According to the definition, a neighbour can be reached from a given configuration through the operation of a single elementary environment (i.e. a single "market").

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

The[4] **best neighbour**[5] $B_i(x, d_i)$ of a configuration x for block d_i is the most preferred configuration in the set of neighbours:

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

A configuration x is a **local optimum** for block d_k if $H_i(x, d_h)=\varnothing$

By extension from single blocks to entire decomposition schemes, one can give the following definitions:

$H(x,D)=\bigcup_{i=1}^{k} H_i(x,d_i)$ is the **set of neighbours** of configuration x for decomposition scheme D

A configuration x is a **local optimum** for the decomposition scheme D if $H(x,D)=\varnothing$

A (**search-**) **path** $P(x^\delta,D)$ from a configuration x^δ and for a decomposition scheme D is a sequence, starting from x^δ, of neighbours:

$P(x^\delta,D)=x^\delta, x^{\delta+1}, x^{\delta+2},\ldots$ where $x^i\in H(x^{i-1},D)$

A configuration y is **reachable** from another configuration x and for decomposition D if there exist a $P(x,D)$ such that $y\in P(x,D)$

Suppose configuration x^0 is a local optimum for decomposition D, we call **basin of attraction** $\Psi(x^0,D)$ of x^0 for decomposition D the set of all configurations from which x^0 is reachable:

$\Psi(x^0,D)=\{y \text{ such that } \exists P(y,D) \text{ with } x^0\in P(y,D)\}$

---

[4] The assumption that ≥ be anti-symmetric guarantees the uniqueness of the best neighbour.

[5] A special case is when a block-configuration is always a best neighbour for any starting configuration. This case is called dominance of a block configuration and is examined by Page (1996).

A **best-neighbour-path** $\Phi(x^{\delta},D)$ from a configuration x^{δ} and for a decomposition scheme D is a sequence, starting from x^{δ}, of best neighbours:

$\Phi(x^{\delta},D)=x^{\delta}, x^{\delta+1}, x^{\delta+2},\ldots$ where $x^{i}=B_{h}(x^{i-1},d_{h})$ for $d_{h}\in D$

Proposition 1: if x^0 is a local optimum for decomposition D and it is reachable from x^i, then there exist a best-neighbour-path leading from x^i to x^0.

Proof: cf. Appendix.

Proposition 1 states that reachability of local optima can be analysed by referring only to best-neighbour paths. This greatly reduces the set of paths one has to test in order to check for reachability. It is worth emphasising however that the result is fundamental for the God like analysis of the objective problem solving structure but must not be taken as saying much about problem-solving complexity by any cognitively bounded agent.

The following proposition establishes a rather obvious but important property of decomposition schemes: as one climbs within the basin of attraction of a local optimum, finer decomposition schemes can usually be introduced which allow to reach the same local optimum.

Proposition 2: let $\Psi(x^0,D)=\{x^0,x^1,.....,x^{\delta}\}$ be the ordered basin of attraction of local optimum x^0, and define $\Psi^{i}(x^0,D)=\Psi(x^0,D)\backslash\{x^i,x^{i+1},.....,x^{\delta}\}$ for $0<i\leq\delta$. Then if $D\neq\{\{1\},\{2\},\{3\},.....,\{N\}\}$ there exist an i such that for $\Psi^{i}(x^0,D)$ a decomposition $D'\neq D$ can be found with $sz(D')<sz(D)$.

Proof: cf. Appendix

Among all the (possibly many) decomposition schemes of a given problem, one is obviously interested in those whereby the global optimum becomes reachable from any starting configuration. One such decomposition always exist, and is the degenerate decomposition $D=\{\{1,2,3,....., N\}\}$ for which of course there exist only one local optimum and it coincides with the global one. But one is equally 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 decentralised selection, with the further requirement that such selection processes can eventually lead to optimality from any starting condition. Note that, finer decompositions do not in general (unless the starting configuration is "by luck" within the basin of attraction of the global optimum) allow decentralised selection processes to optimise.

Let us re-arrange the configurations in X by descending order $X=\{x^1,x^2,.....,x^p\}$ where $x^i \geq x^{j+1}$. The algorithm can be described informally[6] as follows:

1. start with the finest decomposition $D^0=\{\{1\},\{2\},\{3\},.....,\{N\}\}$
2. check whether there is a best-neighbour path leading to x^1 from x^i, for $i=2,3,\ldots 2^N$, 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.

3.3 Near-decomposability.

When building a decomposition scheme for a problem in the foregoing "perfect cognition" vein, one has looked for perfect decomposability, in the sense that one required that all blocks be optimised in a totally independent way from the others. In this way we are guaranteed to decompose the problem into perfectly isolated components (in the sense that each of them can be solved independently). This is indeed very stringent a requirement: even when interdependencies are rather weak, but diffused across all components, one easily tends to observe problems for which no decomposition exists. For instance in Kauffman's NK landscapes, already for small values of K such as 1 or 2 - that is for highly correlated landscapes - the algorithm described above finds only decomposition schemes of size N or just below that value.

Let us soften the requirement of perfect decomposability into one of near-decomposability: one does not want the problem to be decomposed into completely separated sub-problems, (i.e. sub-problems which fully contain all interdependencies) but wants sub-problems to contain only the most "relevant" interdependencies while less relevant ones can persist across sub-problems. In this way, optimising each sub-problem independently will not necessarily lead to the global optimum, but to one of the "best" solutions[7]. In the rigorous meaning defined below, one may construct sorts of "near-decompositions" which give a precise measure of the trade-off between decentralisation and optimality: higher degree of decentralisation together with higher speed of adaptation, vs. the optimality of the solutions which can be ultimately reached.

Let $X_\mu=\{x^1,x^2,.....,x^\mu\}$ with $1\leq\mu\leq 2^N$ be the set of the **best** m **configurations**.

We say that X_μ is **reachable** from a configuration x and for a decomposition D if there exist at least one $y\in X_\mu$ such that y is reachable from x.

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 m**-decomposition** for the problem.

[6] 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.

[7] If solutions were not only ordered but a value was attributed to them, we could easily express all the following in terms of approximations to the global optimum.

m**-decompositions of minimum size** can be found algorithmically with a straightforward generalisation 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, namely:

Proposition 3: if D_μ is a minimum size μ-decomposition, then $sz(D_\mu)$ is monotonically weakly decreasing in μ.

Proof: cf. Appendix.

The general intuition of this proposition is indeed that higher degrees of decentralisation can be attained at the price of giving up the search for globally optimal solutions.

3.4 Problem-solving with changing representations.

So far we have supposed that the "structure" of the problem, i.e. the representation of the space to be searched is exogenously given and cannot be manipulated. But, as already mentioned, problem-solving does not only involve search in a given space but also – and probably more important – a re-framing of the problem itself. In this section we put forward a very preliminary investigation of the properties of problem representations using the toolbox developed in the previous sections. In particular, we show that changing representations can generally be a more powerful problem-solving strategy then searching possibilities generated within a given representation: decentralisation can be increased if more "powerful" representations are built.

One needs some further definitions.

A **representation** of the problem $(X,\geq)$ is a pair (Ξ,➢) where:

$\Xi: X \rightarrow L$ is an **encoding** of the problem, which maps configurations into words of a language L;

➢ is a **preference relation** over possible words in such a language.

We assume that L is made of all and only the words (strings) of a fixed length n over a binary alphabet: $L=\{l;\ l\in\{0,1\}^n\}$. We also assume that the encoding Ξ is a one-to-one mapping, i.e.:

1) $\Xi(x_i)\neq\Xi(x_j) \quad \forall i\neq j$

2) $\exists l\in L;\ l=\Xi(x_i) \quad \forall x_i\in X$

The preference relation ➢ is a "subjective" one which does not necessarily coincide with the "objective" one $\geq$.

We further assume, moving somewhat closer to a descriptive theory, that agents do not know the "objective" problem, but only some imperfect representation of it: this is also defining the space which is being searched with a given decomposition scheme.

Falling short of any exhaustive exploration of the role of representation dynamics (i.e. ultimately of the dynamics of cognitive categories and of individual and collective knowledge) we highlight the paramount role of representation dynamics in problem solving search.

As a preliminary to a much deeper investigation which is still to be undertaken, in the following we just mention three benchmark propositions which together point to the fact that representations can be very powerful search tools. This hints to a possible line of inquiry which considers the construction of shared representations as one of the main functions accomplished by an organisation.

Proposition 4: every problem $(X, \geq)$ admits a representation $(\Xi, \succ)$ and a decomposition scheme $D(\Xi, \succ)$ which can solve it.

Proof: cf. Appendix.

Together, the two propositions which follow, claim instead that the complexity of a problem, its decomposability and the time required to solve it depend on its representation. In fact, by modifying the encoding (proposition 5) and/or the preference relation (proposition 6) we can transform any problem into one of minimum complexity. Considering them together, one is led to the important conclusion that acting on representations tends to be a more powerful problem-solving strategy than acting on the solution algorithm for a given representation.

Proposition 5: given any problem $(X, \geq)$, it admits an encoding Ξ such that it can be solved optimally with the decomposition scheme of minimum complexity $D=\{\{1\},\{2\},\{3\},.....,\{N\}\}$.

Proof: cf. Appendix.

Proposition 6: given any problem $(X, \geq)$ and any encoding Ξ, there exist a preference relation $\succ$ such that it can be solved optimally with the decomposition scheme of minimum complexity $D=\{\{1\},\{2\},\{3\},.....,\{n\}\}$.

Proof: cf. Appendix.

Given the foregoing general properties of problem structures and problem solving, let us put to work such a formal machinery and illustrate some applications and refinements.

4. Problem Decomposition and Attainable Solution Patterns: Some Examples

4.1 Decomposability of a simple coordination problem.

Let us consider a coordination problem involving three binary elements. The set of possible configurations is given by:

$X = \{x_1 , x_2, \ldots, x_n\}$ with $x_i = s_1, s_2, s_3$ (where $\forall j$: $s_j \in \{0, 1\}$)

Two different payoffs structures are considered (where, of course, only the ordinal value matters):

First case

The payoff matrix is given by:

s_1=0

	s_3=0	s_3=1
s_2=0	2,2	-2,-2
s_2=1	-2,-2	3,3

s_1=1

	s_3=0	s_3=1
s_2=0	1,1	-3,-3
s_2=1	-3,-3	4,4

Following the methodology presented above, it can be shown that the problem is not decomposable. However, if one computes the degree of quasi decomposability of the same problem according to various “satisficing degrees”, it can be easily shown that an attainable solution emerges which is accepted that is not the optimal one but allows full decomposability and selects among the best three Pareto superior solutions (i.e. (4, 4,), (3, 3), (2, 2)) the problem becomes a fully decomposable one.

Second case

The payoff matrix is given by:

$s_1=0$

	$s_3=0$	$s_3=1$
$s_2=0$	1,1	-2,-2
$s_2=1$	-2,-2	3,3

$s_1=1$

	$s_3=0$	$s_3=1$
$s_2=0$	2,2	-3,-3
$s_2=1$	-3,-3	4,4

Indeed the problem can be shown to be decomposable according to the scheme: D={{1},{2,3}}. The first element's separability stems from the fact that Nash equilibria for $s_1=1$ are dominant with respect to the corresponding Nash equilibria for $s_1=0$. This is a sufficient condition for separability, even if $s_1=1$ is not a dominant strategy.

Also for this case, problem's decomposability has been computed according to various degrees of satisficingness. Similarly to the previous case, if we accept that the problem be less that optimally solved and one of the three best solutions (i.e. (4, 4), (3, 3), 2, 2)) be accepted, then the problem becomes decomposable (that is: s_2 and s_3 become separable as well.)

4.2 Competing organisational structures

Of course, the measures of decomposability presented so far ought to be taken as sorts of theoretical benchmarks which however may not be generally assumed to be known by boundedly rational agents (actually, finding the minimum size decomposition scheme of a problem is computationally more complex than solving the problem itself), but we can assume that agents search adaptively the landscape with conjectural search strategies based on a given hypothesis on the decomposition of the search space. What are then the evolutionary properties of populations of agents (organisational structures) which compete on the basis of search strategies based on conjectural decompositions?

In order to provide some preliminary answers to this questions, we simulated[8] a population of agents, where each agent is characterised by one out of a limited set of decomposition strategies. We

[8] All simulations have been developed using the simulation platform called "LSD" (Laboratory for Simulation Development) (Valente, 1998) which provides a programming environment where simulations can be easily run also by inexperienced computer users. The interested reader can find downloadable code, user's manuals, more details of these

let them compete in an environment defined by some simple rules of selection: worst scoring agents are removed from the population and replaced with "copies" of the best scoring agents. Copies are new agents that inherit the parent'*s* decomposition scheme, but explore the landscape starting from a different – randomly assigned - point.

Agents start from a random point in the landscape and explore it according to their individual decomposition scheme. Each agent is defined by a given decomposition of the system into sub-systems and generates new points by choosing randomly one of the sub-systems, and mutating its bits (possibly all of them) randomly.

For instance, consider an agent who follows the decomposition strategy D={{1,2,3},{4,5,6,},{7,8,9},{10,11,12}}. Thus, the search space is decomposed into four subspaces of equal size. In order to generate a new point, the agent chooses randomly one of these 4 sub-spaces, then some (possibly all) of the bits in the chosen schema are mutated. The fitness value of the new point is observed: if such a fitness is higher than the one of its current position the agent moves to it, and the latter becomes the new starting point, otherwise the agent remains where it is. These steps are iterated. Thus, what differentiates agents is only the way they generate new points to be tested, i.e. their search strategy, which in turn is determined by their decomposition of the search space.

Of course there is a huge number of possible decompositions, but in order to restrict such a number[9] we imagine that agents "know" that only some "well-behaved" decompositions are possible, in particular we imagine that blocks of an admissible decomposition must have all the same dimension and that they form a partition of the search space. We are thus left with only six possible decomposition and, correspondingly, six types of agents, named after the dimension of the sub-problems into which they decompose the problem:

Agent type 1: D={{1},{2},{3},{4},{5},{6},{7},{8},{9}{10},{11},{12}}
Agent type 2: D={{1,2},{3,4},{5,6},{7,8},{9,10},{11,12}}
Agent type 3: D={{1,2,3},{4,5,6},{7,8,9},{10,11,12}}
Agent type 4: D={{1,2,3,4},{5,6,7,8},{9,10,11,12}}
Agent type 6: D={{1,2,3,4,5,6},{7,8,9,10,11,12}}
Agent type 12: D={{1,2,3,4,5,6,7,8,9,10,11,12}}

We first checked whether agents whose decomposition perfectly reflects the structure of the underlying landscape are actually able to find its global optimum irrespectively of the initial conditions. To test this hypothesis we built five kinds of random landscapes with a given structure, determined by the following minimum size decomposition schemes:

---

and other simulations and programs to run simulations with different parameters and settings at the site: http://www.business.auc.dk/~mv/Lsd1.1/Intro.html

[9] As we want to study the evolutionary properties of optimal decompositions, we have to compare them with some "reasonable" ones, this is the reason why we restrict so much the number of possible decompositions.

Landscape type 1: D={{1},{2},{3},{4},{5},{6},{7},{8},{9}{10},{11},{12}}

Landscape type 2: D={{1,2},{3,4},{5,6},{7,8},{9,10},{11,12}}

Landscape type 3: D={{1,2,3},{4,5,6},{7,8,9},{10,11,12}}

Landscape type 4: D={{1,2,3,4},{5,6,7,8},{9,10,11,12}}

Landscape type 6: D={{1,2,3,4,5,6},{7,8,9,10,11,12}}

Landscape type 12: D={{1,2,3,4,5,6,7,8,9,10,11,12}}

For each type of landscape we let 180 artificial agents, 30 for each type, evolve without any selection. Agents are initialized at a random point of the landscape, and then they pursue their search strategy as described above (i.e. they choose one sub-problem and mutate randomly at least one bit in it). To avoid the effect of lucky initial conditions (of course any agent can find the global optimum if it starts “close” enough to it) we periodically “shake” the population: when fitness values have settled, we reposition all existing agents in randomly chosen points, from which they have to start again their search.

In figure 1 we report the average fitness values for each class on a landscape of type 4.

Figure 1: Average Fitness of Classes of Agents – Landscape type 4

Title:
Window .f.new1.f.plots
Creator:
Tk Canvas Widget
Preview:
This EPS picture was not saved
with a preview included in it.
Comment:
This EPS picture will print to a
PostScript printer, but not to
other types of printers.

Average fitness values of each class are always non-decreasing (because agents either increase their fitness or stay put) except the sudden drops they incur every 1000 iterations, when agents are randomly repositioned. Note that the average fitness of agents of types 1, 2 and 3 quickly stops growing, because all agents in those classes become locked in local optima (some of them may actually reach the global optimum if their starting point is close enough to it, but this effect tends to be cancelled out by the periodic repositioning of agents). All agents in class 4 increase quickly their fitness and set on the global maximum. Agents in class 6 grow less quickly than agents in the lower classes; moreover they cannot all reach the global optimum, since some of them are trapped in local optima. This is because at each mutation they can indeed jump on a wider range, but their decomposition does not allow to identify with certainty the global maximum. Agents in class 12 use a purely random search, having the possibility to jump on every point at each mutation. This allows them to grow slowly but steadily, and in the end all agents in this class may also be able to find the global maximum. However this strategy is extremely slow, and is clearly outperformed, in speed of convergence, by the one used by agents of class 4.

Equivalent results have been found for landscapes of types 2, 3, 6 and 12. In all exercises one observes the same property: strategies using decompositions not including the one corresponding to the minimum size decomposition scheme are bound to be trapped in local optima. Strategies using decompositions which include the minimum one but are bigger do always reach the global optimum but do so slowly. Search strategies corresponding to the minimum decomposition scheme clearly outperform every other strategy.

However, it is remarkable that correct decomposition strategies might not always prevail when nested in competitive environments characterised by some form of selection. In fact, while they are able to always locate the global optimum with certainty, the time required might be so long that they are actually eliminated by the selection mechanism. In order to test this proposition, we ran a set of simulations in a set-up very similar to the above one (180 agents, 30 for each strategy initially located in points of the landscape randomly chosen) however agents are not "re-positioned". Moreover, every 10 mutations they are ranked according to their fitness: the 30 worst scoring agents are eliminated from the population and replaced by copies of 30 agents chosen among the surviving ones.

Figure 2 shows the average number of copies over 10 different simulation runs on each of the landscapes considered (the averages are meant to avoid possible influences of particular initial conditions).

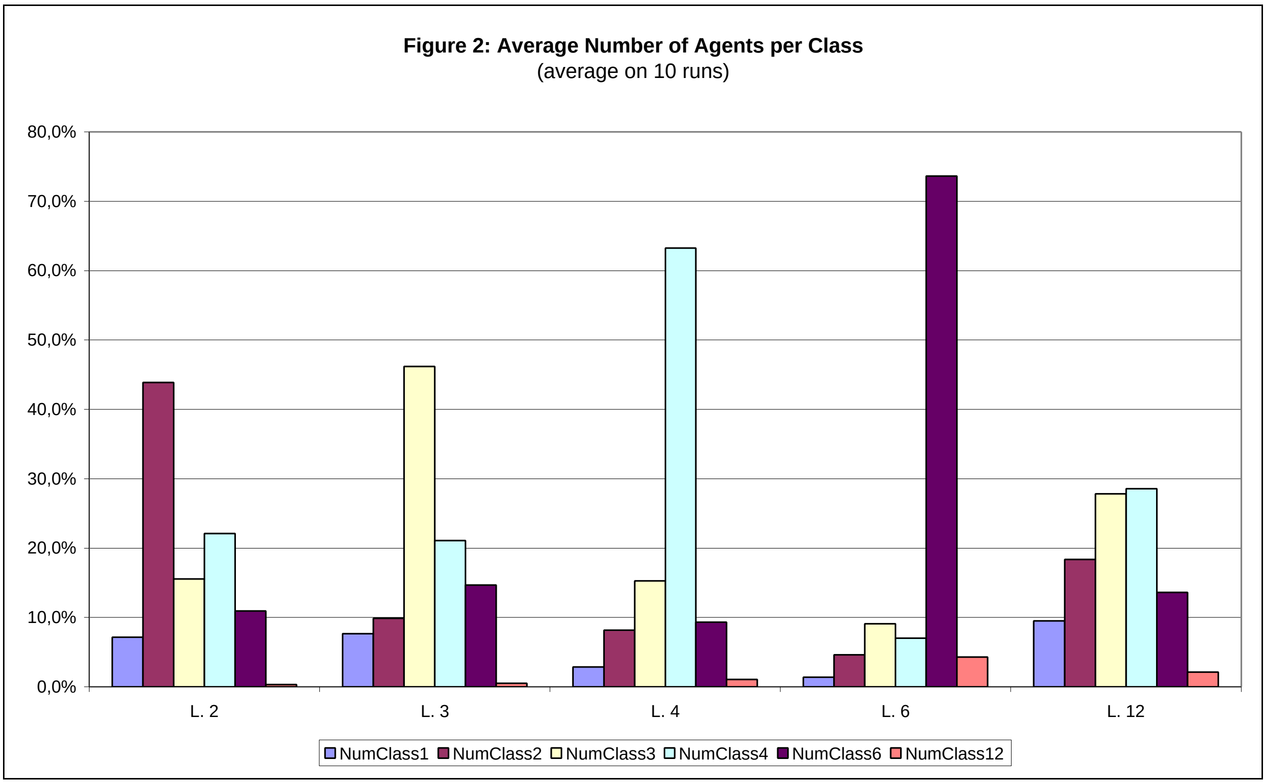

As it might have been expected, for landscapes of types 2, 3, 4 and 6, agents whose decomposition perfectly reproduce the structure of the landscape tend to dominate the population. On the contrary, on landscapes of size 12 this does not happen, as agents of type 12, in spite of being the only ones always able to find the global optimum, are displaced by “simpler” strategies which tend be locked into local optima, but reach them relatively quickly.

Finally, we have considered a population of agents competing in the exploration of a landscape whose structure, defined in terms of decomposability, is given but is however “deformed” at regular time intervals. This means that, even if the structure of the landscape stays basically the same, the position of peaks can change.

Again, we consider a population of 180 agents evolving according to the same adaptation and selection rules described above.

Our results show very robustly that as the frequency of landscape’*s* deformation increases, the population tends to be dominated by agents whose decomposition degree is smaller. In the limit, even in the case of totally decomposable landscapes, the population is rapidly dominated by non-decomposing agents.

4.3 Increasing Returns to Specialization

We have already emphasized that the formal framework presented here is well suited to capture some general problem-solving properties of alternative patterns of division of labour. A straightforward refinement entails the introduction of some form of Smithian increasing returns to specialization.

In order to do that, let us consider a class of coordination problems whose elements are in the set $S=\{s_1,s_2,\ldots,s_{12}\}$ and whose configuration set is $X=\{x_1,x_2,\ldots,x_{4096}\}$ with, as usual, $x_i=s_1s_2\ldots s_{12}$ where $s_j\in\{0,1\}$. Alike the foregoing exercises, we consider a population of 180 agents, evenly distributed among the same 6 types and evolving according to the same death and birth process. In addition, we model increasing returns to specialization by assuming that when an agent finds itself located on a local optimum relative to its decomposition, the value of that optimum itself increases at a rate *q* that is inversely proportional to the dimension of the adopted decomposition. It follows that type 1 agents (i.e. those which tend to decompose more) are subject to lock in effects on suboptimal peaks. In other words, as *q* increases, those agents with a greater attitude towards decomposition tend to dominate the whole population (no matter the problem's decomposability degree.)

Simulations show that as the rate *q* increases, even for problems characterized by low levels of decomposability, the population tends to be more and more dominated by agents which employ search heuristics founded on high levels of decompositions.

With respect to these results, our conjecture is however that the key factor in inducing increasing returns to specialization is to be explained in terms of the absence of "necessary" limits to the division and specialization process. In other words, we can imagine a division of labour process as a tree diagram of unbounded depth in which, say, nodes at level *i* can be reached only after parents nodes at level *i-1* have been reached. The main point is thus not quite that of improving efficiency in accomplishing those tasks associated with a single node, but rather that of increasing the depth and width of the division of labour by progressively sub-dividing tasks. However the methodology presented here is unable to capture these aspects as it assumes a lower bound to the division of labour (our constituent components) which cannot be further subdivided.

4.4 Partial Representations and Collective Problem-solving

A different set of preliminary simulations considers instead a population of agents who hold only a partial representation of the overall problem, in the sense that they control only a limited number of elements involved.

Let us consider again the usual 12-dimensional problem. That is: $S=\{x_1,x_2,\ldots,x_{12}\}$, $X=\{x_1,x_2,\ldots,x_{4096}\}$ with, as usual, $x_i=s_1s_2\ldots s_{12}$ where $s_j\in\{0,1\}$.

We let $I=\{1,2,3..,12\}$ be the index set. Every agent A_j is defined by a subset of the index set $A_j\subseteq I$ and it is characterized by a dimension $|A_j|$ (that denotes the agent's "perspective", i.e. the (limited) number of dimensions it controls.)

The decomposition methodology defined above allows the definition of:

1. the optimal structure (i.e. that which minimizes agents' dimension) that permits parallel solution;
2. the optimal structure that permits sequential solution ("simpler" than the previous one when blocks forming the decomposition are overlapping).

We have thus carried out simulations of random organisations. That is, a configuration is randomly selected and an agent is then randomly chosen that, restricted to its controlled dimensions, makes some mutations thus generating a new configuration. Should this new configuration be fitter than the starting one this is assumed as a new starting configuration. Iterating this kind of procedure, a problem space can be explored not only in terms of different decomposition strategies but in terms of different "cognitive" (i.e. relative to complexity and competencies) perspectives as well.

Our preliminary results show that agents characterised by coarser decompositions get bigger average payoffs, as they have lower probabilities of being locked into local optima, as they have the possibility of moving through ampler steps.

5. Heuristics and decomposability in individual and collective problem-solving: some conclusions and directions for further research.

This paper is our first building block of a research project which tries to achieve a better understanding of organisations as 'collective, imperfect, inferential machines' to extend a definition that one of us used earlier in connection with individual judgements and decisions (Legrenzi, Girotto and Johnson Laird (1994)).

In this work we have suggested that organisations – through their very structure – implement collective mechanisms which perform complex inferential tasks by combining simple individual heuristics. In particular we have shown how the division of labour determines which collective solutions are generated and tested on the grounds of simple individuals trial-and-error search heuristics. We may thus interpret our results as a step towards a better understanding of the relations that bind together the cognitive and the governance roles of organizational routines.

More specifically, in our perspective diverse organizational forms map into diverse a) problem representations, b) problem decompositions, c) task assignment, d) heuristics for and boundaries to exploration and learning, e) mechanisms for conflict resolution over interests but also over alternative cognitive frames and problem interpretations. With respect to these dimensions, one might think, at one extreme, of an archetype involving complete, hierarchical representations, precise task assignment, quite tight boundaries to exploration and, if all that works, no need for expost conflict resolution. The opposite extreme archetype might be somewhat akin a university departments, with a number of representations at least as high as the number of individuals, fuzzy decompositions and conflict resolution rules, little task assignments and loose boundaries to exploration (more on this in Dosi, Hobday and Marengo (2000)). The formal apparatus suggested above allows indeed a quite general account of the problem-solving properties of different structural forms and the diverse instantiations they entail of the fundamental tradeoffs between decentralization, search costs and "quality" of the ensuing solutions.

More speculatively, let us conclude by suggesting some possible deeper isomorphisms in structures of problem-solving knowledge, inferential mechanisms and learning, which goes well beyond the acknowledgement that organisations rather than curbing individual decision biases tend indeed to amplify them (for discussions, cf. March (1994) and Dosi and Lovallo (1995)). Even at a superficial look one finds strong similarities between concepts used here and elsewhere with reference to organisational analysis and those used in relation to individual cognition and problem-solving (cf. Legrenzi and Girotto (1996)). For example, the notions of 'problem-restructuring' inherited from *gestalt* psychology is highly germane to the idea developed above that different 'decomposition' of a problem entail different 'divisions of labour' in the search space. Or, another example is the similarity one finds between 'routine expertise' in the literature on 'expert knowledge' (Holyoak and Spellman (1993)) and routines without further qualification in organisational analysis (Nelson and Winter (1982), Cohen *et al.* (1995)). Needless to say, the fundamental task is to go beyond sheer analogy and explore, via both formal explorations and experimental exercises, what insights can be gained for the analysis of organisations as systems of distributed, but relatively structured, evolving, problem-solving knowledge.

In particular, we submit the conjecture that there is much more than a metaphorical analogy between problem-decomposition in collective organisations and decompositions and other heuristics they typically employ in their cognitive endeavours. At a formal level, decompositions are a particular kind of search heuristics, that is sets of rules for search that limit the space of configurations to be generated and tested. In these abstract terms, decompositions and heuristics have the very same function: that of reducing the size of the search space. Actually, we can imagine an agent who does not explicitly decompose the search space, but who has a very "powerful" (i.e. effective) heuristic that can reduce, to a significant degree, the space of configurations to be tested, for instance by drawing powerful inferences from the already tested solutions. The space can be possibly reduced to one whose magnitude is similar to that of an agent capable of full decomposition. In other words, a decomposition is, after all, nothing but a particular instance of heuristic.

At least equally interestingly, there is a striking similarity between problem-decomposition we have formalised so far and the experimental evidence from cognitive psychology on the 'naive' problem-solving heuristics which individuals typically display[10].

Finally note that, as argued at the beginning, on purpose we meant to explore as a first approximation a conflict-free, incentive-free fiction of organisations, in order to focus precisely on their knowledge dimension. However, the formal machinery presented here allows fruitful links with those other neglected dimensions. In particular, as already hinted above particular knowledge decomposition easily relates with credit assignment decomposition of organisational structures (i.e., roughly speaking, who should be blamed or rewarded for what) and by the same token, with issues of incentives, control and power over tasks assignments..

But, of course, all this is well beyond the foregoing, highly preliminary, work, largely aimed at presenting some basic building-blocks of a 'constructive' theory of organisations as repositories of knowledge.

[10] For instance the so-called "conservative focussing" heuristic, which has been shown to be the typical heuristic in concept formation (cf. Bruner, Goodnow and Austin (1956)), is nothing but a full decomposition of the search space. Actually, experimental evidence shows that "concepts" which are highly decomposable are "easy" to apprehend by human subject, while concepts which are not decomposable are much more difficult and usually human subjects fail to discover them.

Appendix – Proofs of Propositions

In this appendix we provide the proofs of the propositions in the text, reported from Marengo (1999).

Proof of Proposition 1: by hypothesis x^i belongs to the basin of attraction $\Psi(x^0,D)$ of x^0. Let us order all the configurations in $\Psi(x^0,D)$ by descending rank: $\Psi(x^0,D)= =\{x^0,x^1,.....,x^\delta\}$ with $x^i \geq x^{i+1}$. If $x^i=x^1$ then, by definition, x^0 must be a best-neighbour of x^i. If $x^i=x^2$ then either x^0 is a best neighbour of x^i or is not. In the latter case x^1 must necessarily be a best-neighbour of x^i. And so on by induction…

Proof of Proposition 2: If i=1 x^0 is trivially reachable from x^0 itself for all decompositions, including the finest $D=\{\{1\},\{2\},\{3\},.....,\{N\}\}$.

Proof of Proposition 3: if $\mu=2^N$ then X_μ includes all configurations and it is trivially reachable for any decomposition, including the finest $D_\mu=\{\{1\},\{2\},\{3\},.....,\{N\}\}$ with sz=1. If $\mu=1$ X_μ includes only the global optimum, thus the size of the minimum size decomposition is $1 \leq sz(D_1) \leq N$. We still have to show that it cannot be $sz(D_{\mu+1})>sz(D_\mu)$: 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

Proof of Proposition 4: given that we are considering finite problems, this proposition is trivial. Consider in fact a representation $(\Xi,\succ)$ where Ξ is completely free and $\succ$ has the only constraint of preserving the same global maximum of the fitness function. Clearly the decomposition scheme $D=\{\{1,2,3,....., n\}\}$ will find such a global maximum in 2^n steps.

Proof of Proposition 5: we prove the proposition by constructing an encoding which has such a property for a generic problem.

Consider a mapping $\Phi:X\rightarrow N$ from the possible solution into the set of non-negative integers such as:

$\Phi(x_0)=0, \Phi(x_1)=1,.........., \Phi(x_{2^n})=2^n$

where $x_{2^n} \geq \ldots \geq x1 \geq x0$

Define now the encoding $\Xi^*(x_i)=bin(\Phi(x_i))$ where bin is a function which maps an integer into its binary encoding. It is now very easy to verify, because of the properties of binary encoding, that Ξ^* is an encoding which satisfies proposition 2.

Proof of Proposition 6: we prove also this proposition by construction. Let us call *x** the point corresponding to the global maximum of the fitness function, and let Ξ(*x**)=l* be its representation, with $l^* = l_1^* l_2^* \ldots.. l_n^*$

Any preference relation such that $l_1 l_2 \ldots. l_i^* \ldots l_n \succ l_1 l_2 \ldots. l_i \ldots l_n \quad \forall l_i \neq l_i^* \quad i = 1,2,\ldots\ldots, n$ satisfies proposition 9.

References:

Alchian, A. and H. Demsetz (1972), Production, Information Costs and Economic Organisation, *American Economic Review*, vol. 62, pp. 777-795

Aoki, M. (1988), *Information, Incentives and Bargaining in the Japanese Economy,* Cambridge, Cambridge University Press.

Arrow, *K*.J. (1974), *The limits of organisation,* New York, Norton.

Bruner, J.*S*., J. Goodnow and G. Austin (1956), *A Study of Thinking*, New York, Wiley.

Chandler A.D., P. Hagstrom and O. Solvell (eds.) (1998) , *The Dynamic Firm. The role of technology, strategy, organisation and regions*, Oxford, Oxford University Press

Coase, R.M. (1937), The Nature of the Firm, *Economica*, vol. 4, pp. 386-405.

Cohen M., R. Burkhart, G. Dosi, M. Egidi, L. Marengo, M. Warglien, *S*. Winter, and B. Coriat (1995), Routines and other recurring action patterns of organisations : Contemporary Research Issues, *Industrial and Corporate Change*, vol. 5, pp. 653-98.

Coriat, B. and G. Dosi (1998), Learning how to govern and learning how to solve problems, in Chandler *et al.* (1998).

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*, vol. 5, pp. 205-20.

Dosi, G. (1988), Sources, Procedures and Microeconomic Effects of Innovation, *Journal of Economic Literature*, vol. 26, pp. 1120-1171.

Dosi G., M. Hobday and L. Marengo (2000), *Problem-Solving Behaviours, Organisational Forms and the Complexity of Tasks*, Pisa, St. Anna School of Advanced Studies, LEM working paper 2000/06

Dosi G. and D. Lovallo (1996), Rational Entrepreneurs or optimistic Martyrs ? Some considerations on Technological regimes, corporate Entries and the evolutionary Role of Decision Biases, in R.

Garud, P. Nayyar and Z. Shapira (eds.), *Technological Foresights and Oversights*, Cambridge, Cambridge University Press..

Dosi G. and F. Malerba (1996) : *Organisation and Strategy in the Evolution of the Enterprise*, London, MacMillan.

Dosi G. and L. Marengo (1994), "Toward a Theory of Organisational Competencies", in R.W. England (ed.), *Evolutionary Concepts in Contemporary Economics*, Ann Arbor, Michigan University Press, pp. 157-178.

Dosi, G., L. Marengo and G. Fagiolo (2000), Learning in Evolutionary Environments, in K. Dopfer (ed.), *The Evolutionary Foundations of Economics*, Cambridge, Cambridge University Press.

Dosi G., R. Nelson and *S*.G. Winter (2000) (eds.), *The Nature and Dynamics of Organisational Capabilities*, Oxford, Oxford University Press.

Dow, G.*K*. (1987), The function of authority in transaction cost economics, *Journal of Economic Behavior and Organisation*, vol. 8, pp. 13-38.

Freeman, C. (1982), *The Economics of Industrial Innovation*, 2nd edition, London, Frances Pinter.

Frenken, K., Marengo L. and M. Valente (1998)*, Modelling Decomposition Strategies in Design Spaces as Complex Fitness Landscapes. Implications for the Economics of Technological Change*, mimeo. Dept. Of Economics, University of Trento.

Gavetti, G., Levinthal D. (2000), Looking Forward and Looking Backward: Cognitive and Experiential Search, *Organization Science*, forthcoming.

Grossman *S*. and O. Hart (1986), The costs and benefits of ownership: A theory of vertical and lateral integration, *Journal of Political Economy*, vol. 94, pp. 691-719.

Groves, Th., R. Radner and S. Reiter (1987) (eds.), *Information, Incentives and Economic Mechanisms. Essays in Honor of Leonid Hurwicz,* Minneapolis, University of Minnesota Press.

Holyoak, K.J. and B.A. Spellman (1993), Thinking, *Annual Review of Psychology*, vol. 44, pp. 265-315.

Johnson-Laird, P.N., Legrenzi, P. (1992), Science as heuristic search, *International Studies in the Philosophy of Science*, vol. 6, pp. 35-40.

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

Klepper, S. (1997), Industry life cycles, *Industrial and Corporate Change*, vol. 6, pp. 145-181.

Laffont, J.J. and D. Martimort (1998), *Mechanism design with collusion and correlation*, Toulouse. GREMAQ working paper.

Laffont J.J. and J. Tirole (1986), The dynamics of incentive contracts, *Econometrica*, vol. 94, pp. 1153-1175.

Legrenzi, P., Girotto, V. Johnson-Laird, P.N. (1994). Focussing in reasoning and decision making, in P.N. Jonhson, E. Shafir (Eds.), *Reasoning and Decision Making* , Cambridge MA, Blackwell.

Legrenzi, P., e Girotto , V. (1996). Mental models in reasoning and decisions making, in A. Garham e J. Oakhill (ed.), *Mental models in cognitive science*, Hove, Erlbaum-Psychology Press, pp.95-118.

Legrenzi, P., G. Politzer and V. Girotto (1996), Contract proposals. A sketch of a grammar, *Theory and Psychology*, vol. 6, pp. 247-65.

Levinthal, D. (1997), Adaptation on rugged landscapes, *Management Science*, vol. 43, pp. 934-950.

Levinthal D. (1998), The slow pace of rapid technological change: gradualism and punctuation in technological change, *Industrial and Corporate Change*, vol. 7, pp. 217-247.

March, J.G. (1991), Exploration and exploitation in organisational learning, *Organisation Sciences*, vol. 2, pp. 71-87.

March J.G. (1994), *A Primer on Decision Making. How decisions happen*, New York, Free Press.

March, J.G. and H.A. Simon (1993), *Organisations*, New York, Wiley, 2nd ed..

Marengo L. (1996) : Structure, competencies and learning in an adaptive model of the firm, in G. Dosi and F. Malerba (1996).

Marengo, L. (1999), *Decentralisation and market mechanisms in collective problem-solving*, Dept. of Economics, University of Trento, mimeo.

McGuire C.B. and R. Radner (1986) (eds.), *Decision and Organisation,* Minneapolis, University of Minnesota Press.

Milgrom, P. and J. Roberts (1990), The Economics of Modern Manufacturing, *American Economic Review*, vol. 80, pp. 511-528.

Milgrom, P. and J. Roberts (1992), *Economics, Organisation and Management,* Englewood Cliffs, Prentice Hall.

Nelson R.R. and *S*.G. Winter (1982), *An Evolutionary Theory of Economic Change*, Cambridge MA, Harvard University Press.

Page, S.E. (1996), Two measures of difficulty, *Economic Theory*, vol. 8, pp.321-346.

Pavitt, K. (1999), *Technology, Management and Systems of Innovation*, Cheltenham, Edward Elgar.

Radner R. (1992), Hierarchy: The economics of managing, *Journal of Economic Literature*, vol. 30, pp. 1382-1415.

Richardson, G. (1972), The organisation of industry, *Economic Journal*, vol. 82, pp. 883-896.

Riordan, M. and O. Williamson (1985), Asset specificity and economic organisation, *International Journal of Industrial Organisation,* vol. 3, pp. 365-378.

Simon, H. (1981), *The Sciences of the Artificial*, Cambridge Mass., MIT Press. (2nd edition).

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

Simon, H.A. (1991), Organizations and markets, *Journal of Economic Perspectives*, vol. 5, pp. 25-44.

Stigler, G. (1951), The division of labor is limited by the extent of the market, *Journal of Political Economy*, vol. 59, pp. 185-193.

Teece, D.J., G. Pisano and A. Shuen (1994), Dynamic Capabilities and Strategic Management, *Strategic Management Journal,* 18, pp. 509-533.

Teece D.J., R. Rumelt, G. Dosi, and S.G. Winter (1994) : Understanding Corporate Coherence: Theory and Evidence, *Journal of Economic Behavior and Organization*, vol. 23, pp. 1-30.

Tirole J. (1986) , Hierarchies and bureaucracies: On the role of collusion in organizations, Journal of Law, Economics and Organization, vol. 2, pp. 181-214.

Valente, M. (1998), Laboratory for simulation development – Lsd, www.business.auc.dk/~mv/Lsd1.1

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

Williamson, O.E. (1985), *The Economic Institutions of Capitalism*, New York, Free Press.

Winter, *S*.G. (1988), On Coase, competence and the corporation, *Journal of Law, Economics and Organisation*, vol. 4, pp. 181-197.