

DECOMPOSABILITY AND MODULARITY OF ECONOMIC INTERACTIONS

Luigi Marengo, Corrado Pasquali and Marco Valente

Department of Economics
University of Trento
Via Inama, 5
I-38100 TRENTO, Italy
e-mail: lmarengo@gelso.unitn.it

Final Version – March 2001

Paper prepared for the Workshop "MODULARITY: UNDERSTANDING THE DEVELOPMENT AND EVOLUTION OF COMPLEX NATURAL SYSTEMS", Altenberg, Austria, October 26-29, 2000.

ACKNOWLEDGEMENTS: we thank Esben Andersen, Giovanni Dosi, Massimo Egidi, Koen Frenken, Yuri Kaniovski, Thorbjørn Knudsen, Daniel Levinthal, Paolo Legrenzi, Scott Page, and Herbert Simon for useful discussions, remarks and suggestions. Financial contribution from the project "Bounded rationality and learning in experimental economics" funded by the Italian Ministry of University and Scientific and Technological Research (MURST) is gratefully acknowledged.

ABSTRACT

Although never does the term "modularity" show up in economic textbooks, economic theory is built around a very strong idea of modularity: the entire economic life, as argued by orthodoxy, can be successfully managed by a system in which all information is encapsulated within individual "atomic" economic agents (individual consumers, individual workers, individual entrepreneurs) and coordination is entirely carried out within markets, whereby the only conveyed information concerns the relative scarcity of goods as reflected in prices. This view is epitomised by the so-called "Coase theorem" which states that, even in the presence of interdependencies, coordination can be successfully achieved – at least in principle - via market mechanisms, provided that the granularity of the system is fine enough to encompass all "atomic" economic entities and a proper market exists for each of them (for instance a good is not necessarily the right "grain" and we might need to split it into a multiplicity of independently negotiable rights).

Nevertheless the view of economic systems as organised around a minimal level of granularity is clearly confuted by a straightforward observation of the existence and importance of economic entities – such as business firms – whose grain is much coarser than the one prescribed by a theory which praises the virtues of decentralisation.

Even at a first glance, the aforementioned keywords revolve around some notion of "granularity" of the economic world: e.g. why do we have thousands of different firms rather than a unique and huge one? why are there firms producing cars rather than people buying wheels, windshields, carburettors and assembling them? or, even more radically, people buying raw iron and building cars from scratch? and how are new markets created? That is: why is economic life settled at the present level of granularity and aggregation?

We will propose and defend two theses: first, we will claim that the historical evolution of economic systems has created new entities and has settled upon a specific level of aggregation due to integration and disintegration processes and thus that these are the fundamental engine and cause of the world having settled at its actual grain. Secondly, we will ask what would happen if the "economic tape were run twice": should we observe again a multitude of firms, producers, consumers, markets and institutions or rather should we see something different? That is: has a god fixed a ultimate and necessary level of granularity of the world or is any level of aggregation possible and attainable?

In this respect, we will present a model of problem solving in which problems are inherently characterised by the presence of interdependencies. We then proceed to a definition of "complexity" of a problem in terms of its decomposability into smaller sub-problems that can independently solved and show that a decentralised market mechanism optimally works as a coordination mechanism only in those cases in which sub-problems are totally independent, in a very specific sense, from one another. In general however a trade-off exists between optimality and decentralisation.

1 Introduction

1.2 Overview and motivation

Talking about modularity and decentralisation in economics is a surprisingly difficult task and involves going into the very heart of the nature of institutions governing economic life. As reported e.g. in Hurwicz (1971), the discussion dates back at least to Plato defending central planning in the *Republic* and Aristotle warning against the dangers and disadvantages of collective ownership. Far from the presumption to directly join a discussion with these big names of Western thought, we present a model and some results dealing with how well different institutional settings, characterised by different degrees of decentralisation, perform as backgrounds for distributed problem solving.

The main focus of our work is on solving problems whose solution derives from coordinating a large number of interacting and interdependent entities which altogether contribute to forming a solution to the problem itself. The key issue and difficulty addressed here is the opacity of single entities' functional relations and the partial understanding of their context-dependent individual contributions in forming a solution to the problem at hand. Our model accounts for the relationships between problem complexity, task decentralisation and problem solving efficiency.

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 (Kauffman (1993)). Kauffman's approach to the exploration of a fitness landscape, however, does not necessarily fit well with the realm of social evolution. The main reason for this inadequacy is that social actors might well engage in adaptive walks based on a far richer and virtually huge class of algorithms other than single bit mutations alone considered by Kauffman.

In particular, following Simon (1983), 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 we have come to call a division of problem solving labour. Searches in the problem space based on one bit mutations amount to fully decompose the problem into its smaller components, while coarser decompositions correspond to mutating more bits together. In this work we study some general properties of decompositions and the effect of adopting mutational algorithms of any size on the exploration of a fitness landscape.

Imagine, for instance, a N -dimensional problem. Adopting a point-mutation algorithm corresponds to exploring the space of configurations by flipping elements' values one at a time. A point-mutation algorithm corresponds, in our view, to a maximally decentralised search strategy in which each of the N bits forming the problems will be given a value independently from all the other $N-1$ bits. In this sense the whole problem is (trivially) decomposed into the set of the

smallest-sized sub-problems. 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 at a time. Between the finest and the coarsest decompositions, all the other possible decompositions of the problem lie (i.e. all the possible algorithms of any cardinality), each corresponding to a different division of labour.

The division of problem solving labour, as we shall see in greater detail, to a large extent determines which solutions will be generated and then eventually selected. This leaves the possibility open that never will the optimal solution be generated and a fortiori selected, as it might well be that never it will be possible to reach it starting from a given decomposition of the problem. It thus turns out that while decomposing a problem is necessary in order to reduce the dimension of the search space, it also shapes and constrains a search process 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 believe that the main significance of our work is casting some further doubts on any “optimality through selection” argument from a specific point of view. Evolutionary arguments in economic theory have often taken a rather Panglossian form according to which the sole existence of, say, an organisational form or of a technological design can be reliably taken as a proof of its optimality. The view according to which market forces are always able to select away suboptimal types is a widely accepted idea in economic theory. It is also noteworthy that most of the historically grounded attempts to show the limits that selection encounters in the economic realm are grounded on a claim to some sort of weakness of selective pressures.

A famous example is Paul David's work on the persistence of the highly inefficient QWERTY keyboard in typewriters and computers. In his seminal work, David (1985) shows how a technological standard adopted for its efficiency in a given set of constraints (those imposed by mechanical typewriters, with the need to reduce the frequency of lever jams) was not displaced by standards which had become far superior when those constraints had totally disappeared (in electric typewriters and then computers, where the keyboard layout could be freely designed in order to make more easily accessible the most frequently struck keys). However, the case discussed by David deals with a situation in which the “optimal” configuration exists and is actually available as a possible choice but selective pressures are not strong enough to favour it over the others.

Quite on the contrary, our focus is on those cases in which the optimal alternative is not available at all as a consequence of a particular way of exploring a landscape and of a specific accessibility relation between solutions that holds as a consequence of the adopted search algorithms. In our approach, these considerations are highlighted by the fact that exploring a landscape according to

different decomposition patterns implies that the very geometry of the fitness landscape changes and a landscape that is very rugged when explored with a one bit algorithm might become smooth when explored with an algorithm of greater cardinality.

Many authors in the field of theoretical biology emphasise how many evolutionary asymmetries such as patterns and processes of phenotypic evolution, punctuated evolution, developmental constraints, homology and irreversibility do not naturally fit with evolutionary theories as implemented in the neo Darwinian paradigm as it is purely based on (heritable) diversity generation and selection.

An approach that we regard as particularly consonant with our own is that proposed by Stadler *et al.* (2000). In this essay, the authors emphasise the quest for a theory of accessibility between structures in terms of a precise theory of genotype-phenotype maps and claim that the aforementioned asymmetries are in some way rooted in the structure of the map itself and are nothing but “statements about the accessibility topology of the phenotype space” (Stadler et al. 2000). This way of looking at things, in a sense, relates to the idea that accessing or constructing a structure (say a phenotype or a technological design or a solution to a combinatorial problem) by means of a set of variational mechanisms is logically and empirically prior to any selective pressure and of its outcomes.

We thus take the perspective according to which asymmetries in selection and evolution are not solely dependent on selective pressures plus the structure of fitness landscapes but rather on accessibility relations between objects as defined and grounded on specific mutational operators that allow the exploration of a solution space by transforming objects into new objects. Once again, what can be reached from what is highly dependent on some notions of neighbourhood and distance which, as we shall precisely see, are prior and independent from any bare fitness consideration.

The problem of accessibility in current economic theory is approached (and radically solved) in a most clear way whose main effect is making the problem disappear. Basically, both for individual consumers and producers, solution spaces, that is the space of consumption and that of production bundles, are imagined to be uniformly, symmetrically and everywhere accessible so that a continuous path can always be imagined to exist connecting every and each point in those spaces. Expressed in this terms, our main focus is on seriously considering the tangled (though neglected) interaction between economic interaction as such and organisational structures and constrains adopted to complement ambiguous market signals.

1.2 A bird's eye view on economic theory

Lionel Robbins (1935) defined economic theory to be “the science which studies human behaviour as a relationship between given ends and scarce means which have alternative uses”.

Focussing on “scarcity” and “alternative uses” will give us the possibility of articulating a brief overview on the realm and character of economic theory and of its approach towards the problem of modularity.

That resources are not unlimitedly available to us and that limits to their utilisation force us to allocate them among mutually alternative uses are facts that surround each of us experience and everyday life. These two facts taken together imply social forms of competition for resources in order to endow oneself with the possibility of using them.

To the end that competition be mediated and efficiently organised, a system is needed that rules competition for resources and helps in efficiently allocating them to different possible utilisations. The “system” can be thought to be, for instance, an authoritarian one (such as in a military state or in a dictatorship) or a decentralised one such as a market.

Adam Smith emphasised that some form of societal or interpersonal organisation is needed in order to exploit the advantages of co-operation and social interaction. In particular, Smith observed that individuals are different, have different “talents” and, at the same time, skills and capacities of individuals in pursuing their ends increases with specialisation.

To this end, trade and division of labour are necessary so that a butcher is not forced to live his lifetime on meat and he can exchange meat for bread with a baker who, in turn, will get his own advantage from the exchange.

Economic theory has among its aims that of studying the properties of the outcomes resulting from different and alternative societal organisations. In order to tell a better system from a worse one, economists rely on the notion of Pareto efficiency: whatever we mean by “better” or “efficient”, we surely mean that a system or a particular allocation is better than another one if no individual perceives it as a worse one and at least one individual perceives it as a better one. That is: there is no allocation in which everyone would prefer to find himself in.

Leaving aside any reference to a (fairly big) set of assumptions that economic theory relies on, the very heart of the theory is the proof that efficiency so defined can be reached thanks to a specific societal system: the price system. The working of the system can be roughly described as follows. Goods are transferred and an income derives from selling them at given prices. Income is in turn utilised to buy other goods at given prices.

The general idea is that in a setting in which everyone tries to maximise its own welfare and utility, it naturally happens, given certain rather specific hypotheses, that the whole body of a society comes to a point in which a

remarkable degree of coherence shows among a number of different and sparse decisions to sell and buy different commodities. Coherence here means a state in which markets “clear” and supply equals demand for every commodity. To be slightly more precise: with a few (though very strong) assumptions on individual rationality it is possible to show the existence of an equilibrium. That is, for a given economy it is possible to find a vector of prices and an allocation to each individual such that the excess demand function of the economy is zero for every good.

Pressures and information coming from markets are what turn selfish behaviours into socially desirable outcomes. Both pressures and information are represented and implemented by the price system whereby prices reflect the relative scarcity of commodities. So, what happens according to this picture of the economic world is that economic interaction takes place through individual agents reacting (i.e. adjusting demand and supply) to quantitative information coming from the market.

The faith in the possibility and capacity for markets to reach a coherence state in which different and possibly conflicting actions are compensated is crucially based, on the one hand on strong notions of individual and collective rationality and, on the other hand, on a most clearly cut distinction between individual agents and institutional contexts. The latter (i.e. the rules and the “places” in which individual interaction lives) are assumed to be given before any interaction takes place and as either totally transparent to individuals or irrelevant at the level of the theory. In other words, given that a class of hypotheses are met on individual rationality, on collective rationality and on how the latter is grounded on the former, “equilibrium analysis” postulates that every dynamic trajectory of an economy leads to an equilibrium state.

As a matter of fact, an ex-ante definition of interaction structures and institutional contexts amounts to grounding markets' compensation possibilities on issues and factors that are independent from economic interaction itself. Not only: the sharp separation between economic agents and the contexts they interact in and, at the same time, the very fact that these be given before and independently from their interactions entails that agents are imagined as perfectly “adapted” to their environment and that, vice versa, interaction structures are perfectly suited, or either irrelevant, to any task at hand. The only relevant thing is the transparency of institutional contexts and settings and the fact that prices accurately reflect all the relevant information. This consideration will be the first cornerstone of our argument and the starting point of our analysis.

1.3 On the “granularity” of the economic world

Even from our rather sketchy description of economic life as depicted by orthodoxy, it is fairly evident how a strong and most peculiar notion of modularity

underlies the whole thing. The main tenet of this approach is indeed representing economic agents as autonomous and anonymous individuals which take decisions independently from one another and that interact only through the price system. So, all the relevant information is encapsulated within individual economic agents and coordination is achieved within markets by the use of prices, i.e. by pure selection.

One of the most fundamental questions asked by economic theory (and by welfare economics in particular) is about the extent to which perfect competition can lead to an optimal allocation of resources. Among the assumptions made by the theory in order to derive such results is that never must indivisibilities show up in consumption nor in production. The key problem here, and the main relevance of this point for our discussion, is that a decentralised coordination mechanism based on prices is no longer available when indivisibilities show up .

Classical results of welfare economics on the possibility of decentralisation heavily rely on an assumption of perfect decomposability of the underlying allocation problem. The proof of one of the most important results of welfare economics, the Second Welfare Theorem, critically depends on a separation argument and the same proof no longer holds in the presence of externalities, i.e. interdependencies which lead to social interactions not mediated through the market: in the presence of one of the most fundamental instances of non-separability.

In the picture of economic life described so far, every actor is an “island” for any given set of prices. In particular, never is his utility affected by choices made by other actors apart from those choices - i.e. decisions to buy or sell marketable goods - which have a direct reflection on prices. Externalities, on the other hand, are exactly those situations in which actions taken from others might well affect the outcome of ours but do not affect prices and thus cannot be coordinated through the action of markets. An ideal “orthodox” market ceases to work as perfectly as prescribed by the theory as soon as externalities are present. It is not by chance that those situations related to their presence are referred to as “market failures”.

This idea of a sort of perfect modularization of an economy is pushed to its extreme by the so-called Coase theorem (cf. Coase (1937)). In spite of its name Coase's theorem is rather a circular argument and states that if every single activity which affects agents' welfare can be exchanged and allocated in a perfectly competitive market then non-separability ceases to be a problem. For instance, consider an externality, say a factory which produces plastic buckets and in doing so pollutes the environment and negatively affects the welfare of people living nearby. The problem exists because while the socially optimal amount of buckets to be produced can be determined in the perfectly competitive market for buckets, on the contrary the socially optimal amount of pollution cannot be determined because a market for that does not exist. But if that is the problem the solution - at least in principle - can be very easy: create a market for pollution rights and then

they will be allocated in a socially optimal way as every other good¹. In order to have such a market we need to allow buckets and pollution to be allocated independently from each other, by establishing negotiable property rights on pollution rights and allow them to be traded in a competitive market separated from the one in which buckets are traded. In the language of modularity we could say that the problem of externalities arises because we are working with modules that are too large, thus the solution is to disassemble them and let market selection operate on finer units.

When interpreted in terms of degrees of correlation or, to use the biological terminology, in terms of degrees of epistasis, the Coase theorem asserts that, under a set of rather weak assumptions, in any situation the degree of interdependence can always be made minimal. The relevance of property rights in our discussion lies in the fact that with the modularization of property rights we modularise social interaction which then come to be mediated through the interface of voluntary exchange.

In particular, the theorem can be read as saying that every tangled situation can be transformed into one whose degree of granularity is the atomistic one prescribed by the theory, that is: a degree in which every atomic entity is properly encapsulated and bears no correlation with other entities. It is this degree that solely allows competitive markets to perfectly work as decentralised (atomistically)-modular mechanisms.

In a sense Coase's argument could pose an apparently odd but perhaps not so absurd question to biology: if selection is such a powerful device of coordination, adaptation and optimization, why do we observe so little of that? In particular, why is selection applied to rather large ensembles of modules and not to each module separately in order to exploit its power to optimize each module? In other words, why do we observe multi-functional complex living creatures and not much simpler entities specialised in single tasks and coordinated through selection?

Thus we are left with two analogous questions: in biology one could ask why multi-functional complex organisms exist, as in economics we do ask the question: why do firms exist? Is it just because, as Coase would say, there is a cost associated with the use of the price system, i.e. with the imperfections of the (market) selection mechanism (that Coase calls transactions costs)? And if transaction costs are the costs of a bad modularization, what could possibly go wrong with a totally atomistic modularization?

Adopting a Panglossian attitude (which is fairly common in Economics) Coase might well imagined to assert that in the absence of transaction costs, rights (modules) cut too thin or too coarse will quickly reassemble or disassemble into

¹ As Coase himself stresses, in many cases this solution holds only in principle. In reality there might exist many reasons for which the creation of such a market is difficult, costly or even impossible. These reasons are called transactions costs.

optimal bundles. The starting point of this kind of analysis would be imagining a completely atomistic, modular and market-based way of production, consumption and economic interaction (i.e. the finest possible decomposition of economic activities). Williamson (1975) well epitomises this view with the position according to which “in the origins there were markets”. A possible explanation for non separable organisations is what we might call the “non-monotonicity of marginal productivity”. An example is Alchian and Demsetz (1972) discussion on team-work which illustrates how some instances of coordination problems require kinds of transmission of information well beyond what can be sent through the interface of the price system.

We might thus argue that organisations arise as a non modular response to the fact and the need for non-price mediated interactions among different modules (whatever they be). An organisation is always a demodularisation and a repartitioning that limits the right of alienation from at least some rights of decision.

In addition, the view of economic systems as ideally or in principle organised around a minimal level of granularity is clearly confuted by a straightforward observation of the existence and importance of economic entities - such as business firms - whose grain is much coarser than the one prescribed by a theory which praises the virtues of decentralisation.

Actually, the most part of economic life occurs in forms and structures that go largely beyond the atomistic limits envisaged from orthodox economic theory. As recognised by many economists, a large part of economic life occurs in organisations composed of a large number of entities which are not regulated nor coordinated by the price system. Within organisations, other coordination mechanisms rule activities, such as hierarchies, power or authority.

Surprisingly enough, economic theory has not much to say on organisations as such. It rather assimilates them to the same action structure of individual agents and very little room is devoted in economic textbooks to an analysis of how, for instance, different organisational structures might differ in performance, efficiency or speed of adaptation. The most general questions with respect to the organisations/markets opposition are: what is that determines the boundaries between markets and organisations? And what is the real difference between what happens in a market and what happens in an organisation?

In this sense, Coase's main concern can be thought to be about the way negotiation can repair externalities and interdependencies. His argument is broader though and can be so reformulated: if nothing obstructs efficient bargaining, then people can negotiate until they reach Pareto efficiency. So, after all, it might well seem that the well functioning of a competitive market is a matter of finding the right "granularity", the level that ensures individual interaction to be as effective as to reach optimal outcomes. After all, the ultimate meaning of the Coasian argument is: coordination is always achievable via market mechanisms provided

that the "granularity" of the system is fine enough to encompass all atomic entities and a proper market exists (or is created) for all of them.

"Granularity" will be a big issue in our discussion. With this term, we refer, on the one hand, to the level of analysis adopted by economic theory (i.e. studying consumers' and firms' behaviour as individuals and postulating that each and every phenomenon occurring at the aggregate level can be traced back to individual actions and behaviours). On the other hand, we straightforwardly refer to economic reality and we ask why it has settled at the actual level of aggregation: why are there firms producing goods and not people buying raw materials and building the objects they need by themselves? Why are there a multitude of firms rather than a single huge one? And how are new markets created? In other words: why and how has economic life settled at the present level and degree of organisation?

With respect to these points, we will argue that the two main forces that drive the whole process are integration and disintegration processes. These two historical forces we regard to be the ultimate reasons for economic reality having settled into the present level of aggregation. We will further maintain that integration and disintegration processes are the two main economic forces that contribute to the creation of new economic entities. With respect to this point, it is worth pointing to the fact that even from a historical perspective (painting it with an extremely broad brush, of course) the evolutionary path undertaken by capitalist economies has been one in which new-born technologies, productive processes, commodities and whole industries and markets enter the scene in highly integrated settings (see Langlois and Robertson 1989 and Klepper, 1997 for a detailed discussion). It is only at successive stages that a finer grained organisation of, say, production takes place thus creating new markets, new specialised functions and further division of labour. Also note that Adam Smith himself referred to this theme stressing how one of the main reasons for the development of productive capacities is to be found in specialisation deriving from finer grained divisions of labour.

Causes and mechanisms by which new entities are created and enter the economic scene are not at the centre of economic theory as we know it. The problem was stressed by Joseph Schumpeter who pointed out that the real phenomenon that should surprise and interest economists is not how a firm is run, but rather how it got created. It thus seems that economics is affected by the same "object problem" (see Fontana and Buss, 1994) that plagues biology as we know it by the modern synthesis. This is not surprising as we realise that economic theory was born with a strong faith in the possibility of applying the main tools and abstractions of dynamical system theory to the domain of social interaction. This has driven this discipline to take the existence of its very objects and the nature of their dynamical couplings as given and immutable from the very start, thus committing itself to purely quantitative and non constructive analyses and theories. Transaction cost economics and principal-agent theories have partly tried to fill this gap but, we believe, missing a bunch of fundamental points.

Transaction cost economics has developed an explanation of vertical integration in strictly Coasian terms (actually, just expanding upon Coasian ideas). According to this explanation "in the beginning there were markets": there was full, atomistic modularity and integration phenomena, viewed as processes of assembling modules, took place whenever the working of the price system was bound to face comparatively higher costs than those associated with, say, bureaucratic governance. Firms exist because of greater allocative efficiency. This kind of explanation leave unexplained a class of most relevant phenomena and seems to be contradicted by a fairly large class of evident, though neglected, facts.

First of all it conflicts with the historical development of technologies and industries which have mostly developed according to a path going from initial states of high vertical integration to a progressive disintegration. It is also worth pointing out that it is the very process of division of labour that possibly creates opportunities for new markets to exist.

Secondly, it should be pointed out that both logical arguments and empirical evidence exist to support the view that markets, far from being an original state of nature, require some sets of conditions to be met in order to emerge, some of which are even determined by explicit forms of organisational planning.

Thirdly, it is usually implicit in many of transaction cost based approaches that efficiency can be considered as an explanation for existence. That is: proving the relative greater efficiency of an organisational form is in some way considered to be an explanation of its emergence. Actually, it might well be true that selection forces are strong enough to select fitter structures: but these have to exist to be selected: in Fontana and Buss terms, they must arrive before being selected. So, selective forces can and do account for a population's convergence to a specific form but not for the existence and emergence of such a form.

In addition, we know that when the entities subject to selection are entities whose internal structure and components present a strong degree of interdependence their selection landscape will be correspondingly rugged and uncorrelated thus possibly making selection forces unable to drive such entities to global optima. It then follows that selective forces may not be strong enough as to select sub-optimal structures out (hence providing an at least partial explanation for the persistent diversity of organisational forms).

We thus ask: can optimal organisational structures and division of labour patterns emerge out of decentralised interactions? We show this to be possible only in some very special conditions and, at the same time, we prove that decentralisation has usually an associated cost in terms of sub-optimality.

Our second fundamental question will be about the "necessary" character of the actual grain of the economic world. What would we see if we run twice the tape of economic history? Could other evolutionary paths have been taken by the development of economic history? Or, rather, there are some fundamental

properties that we would be bound to see at every run of the evolutionary tape? Is it true that evolution leads to increasing modular structures in the economic realm?

1.4 Our task ahead

In the rest of the paper we analyse some of the issues raised above by means of a very sketchy and abstract model. We analyse the properties of adaptive walks on N-dimensional fitness landscapes by entities which decompose the N dimensions of the problem into modules and adapt such modules independently of each other. Thus inter-modular coordination is performed by selection mechanisms which we assume to be "perfect" in the economic sense (i.e. not subject to any friction, inertia, or any other source of transaction cost).

In the next section we develop a methodology which determines, for any given landscape, its smaller "perfect" decomposition, that is the set of the smallest modules which can be optimised by autonomous adaptation with selection. Then, in section 3, we extend this methodology to near-decompositions, i.e. decompositions which, in a Simonian sense, isolate into separate modules only the most fitness-relevant interdependencies. We show that such decompositions determine, in general, a loss of optimality, but sharply increases the speed of adaptation. Section 4 uses such methodologies to build landscapes whose decompositions and near-decompositions are known and simulate competition among entities which search the landscape with algorithms based on different decompositions. This enables us to test the evolutionary properties of different decompositions.

Finally we provide some, still very tentative, hints to possible directions in which our model can be generalised in order to account for variable representations of the landscapes. There can hardly be any doubt that social institutions do face problems whose structure is not exogenously given, but they base themselves on socially constructed representations of the problem. Such representations are subject to change, through processes of social, organisational and technological innovation, which provide new representations of existing problem or new problems not considered before. Probably such changes in representations occur also in the biological realm, though on much longer time scale. In the final section we briefly discuss how our model could be extended to account for changing representations of the landscapes. We show that the representation of the structure of interdependencies can be modified at will, in particular highly interdependent worlds can be made fully decomposable by appropriate representations and viceversa. This also allows us to cast some doubts on the general usefulness of Kaufmann's K value as an appropriate measure of the complexity of a landscape and we present an example of a problem with a high K but fully decomposable.

2. Decomposition and coordination.

2.1 Problems and decompositions.

We assume that solving a given problem requires the coordination of N atomic elements, which we call generically components, each of which can assume some number of alternative states. For simplicity, we assume that each element can assume only two states, labelled 0 and 1. Note that all the properties presented below for the two-states case can be very easily extended to the case of any finite number of states.

More precisely, we characterise a problem by the following three elements:

1. the set of **components**: $\aleph = \{x_1, x_2, \dots, x_N\}$, with $x_i \in \{0,1\}$
2. the set of **configurations**: $X = \{x^1, x^2, \dots, x^{2^N}\}$ where a **configuration**, that is a possible solution to the problem, is a string $x^i = x_1^i x_2^i \dots x_N^i$
3. an **ordering** over the set of possible configurations: we write $x^i \geq x^j$ (or $x^i > x^j$) whenever x^i is weakly (or strictly) preferred to x^j .

In order to avoid some technical complications, we assume for the time being that there exists only one configuration which is strictly preferred to all the other configurations (i.e. a unique global optimum). This simplifying assumption will be dropped in section 4 below.

\geq

A **problem** is defined by the couple (X, \geq) , and solving it amounts to finding the $x^i \in X$ which is maximal according to \geq .

As the size of the set of configurations is exponential in the number of components, whenever the latter is large, the state space of the search problem becomes much too vast to be extensively searched by agents with bounded computational capabilities. One way of reducing its size is to decompose² it into sub-spaces:

² A decomposition can be considered as a special case of search heuristic: search heuristics are in fact ways of reducing the number of configurations to be considered in a search process.

Let $\mathfrak{S} = \{1, 2, 3, \dots, N\}$ be the set of indexes, and let a **block** $d_i \subseteq \mathfrak{S}$ be a non-empty subset of it, and let $|d_i|$ be the **size of block** d_i , i.e. its cardinality³.

We define a **decomposition scheme** (or simply **decomposition**) of the space \mathfrak{X} as a set of blocks:

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

Note that a decomposition does not have necessarily to be a partition (blocks may have non empty intersections).

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 :

$$x^j(d_k) = x_{k_1}^j x_{k_2}^j \dots x_{k_{|d_k|}}^j \quad \text{for all } k_i \in d_k$$

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 :

$$x^j(d_{-k}) = x_{k_1}^j x_{k_2}^j \dots x_{k_{N-|d_k|}}^j \quad \text{for all } k_i \notin d_k$$

Two block-configurations can be united into a larger block-configuration by means of the \wedge operator:

$$x(d_i) \wedge y(d_j) = z(d_i \cup d_j) \quad \text{where } z_h = x_h \text{ if } h \in d_i \text{ and } z_h = y_h \text{ otherwise.}$$

We can therefore write $x^j = x^j(d_k) \wedge x^j(d_{-k}) = x^j(d_{-k}) \wedge x^j(d_k)$ for any d_k .

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

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

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

- problems which can be successfully solved adopting the finest grained according to the scheme $D = \{\{1\}, \{2\}, \{3\}, \dots, \{N\}\}$ have minimum complexity, while a problem which cannot be decomposed has maximum complexity and it can only be searched extensively;

³ We intend to use intra block as a proxy for hierarchy or centralised organisation and inter block as a proxy for market or decentralised interaction.

- a problem which can be decomposed according to the scheme $D=\{\{1\},\{2\},\{3\},\dots,\{N\}\}$ can be solved optimally in linear time, while a problem which cannot be decomposed can be solved optimally only in exponential time;

- but, on the other hand, a problem which has not been decomposed can always be solved optimally while, as it will be shown below, a problem which has been decomposed according to the scheme $D=\{\{1\},\{2\},\{3\},\dots,\{N\}\}$ - or for that matter according to any scheme whose size is smaller than N - can be solved optimally only under some special conditions, which, as we will show, become generally - though with important exceptions - more and more restrictive as the complexity of the problem increases and 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.

2.2 Selection and coordination mechanisms.

We suppose that coordination among blocks in a decomposition scheme takes place through market-like selection mechanisms, i.e. there are markets which select at no cost and without any friction over alternative block-configurations.

More precisely, 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 configuration $x^h(d_k)$ for the same block, if:

$$x^h(d_k) \wedge x^j(d_{-k}) > x^j(d_k) \wedge x^j(d_{-k})$$

then $x^h(d_k)$ is selected and the new configuration $x^h(d_k) \wedge x^j(d_{-k})$ is kept in the place of x^j , otherwise $x^h(d_k) \wedge x^j(d_{-k})$ is discarded and x^j is kept.

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.

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

2.3 Selection and search paths.

A decomposition scheme is a sort of template which determines how new configurations are generated and can therefore be tested by market selection. 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 $D = \{d_1, d_2, \dots, d_k\}$, we define the following:

A configuration $x^i = x_1^i x_2^i \dots x_N^i$ is a **preferred neighbour** – or, shortly, a **neighbour** – of configuration $x^j = x_1^j x_2^j \dots x_N^j$ for a block $d_h \in D$ if:

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

Conditions 2 and 3 require that the two configurations differ only by components which belong to block d_h .

According to the definition, a neighbour can be reached from a given configuration through the operation of a single market selection mechanism.

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

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

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

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

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

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

A **search path** – or, shortly, a **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:

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

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

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

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

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:

$$\Phi(x^\delta, D) = x^\delta, x^{\delta+1}, x^{\delta+2}, \dots \quad \text{with } x^{\delta+i+1} \in B_h(x^{\delta+i}, d_h) \text{ and } d_h \in D$$

The following proposition states that reachability of local optima can be analysed by referring only to best-neighbour paths. This greatly reduces the set of paths we have to test in order to check for reachability.

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

Proof: by hypothesis x^δ belongs to the basin of attraction $\Psi(x^\alpha, D)$ of x^α . Let us order all the configurations in $\Psi(x^\alpha, D)$ by descending rank:

$$\Psi(x^\alpha, D) = \{x^\alpha, x^{\alpha+1}, x^{\alpha+2}, \dots\} \text{ with } x \geq x^{\alpha+1}$$

Now proceed by induction on x^δ . If $x^\delta = x^{\alpha+1}$ then, by definition, x^α must be a best-neighbour of x^δ for a block in D (in fact, by hypothesis, x^α does not have itself any strictly preferred neighbour). If $x^\delta = x^{\alpha+2}$ then either x^α is a best neighbour of x^δ or is not. In the latter case $x^{\alpha+1}$ must necessarily be a best-neighbour of x^δ . And so on...

Now let x^0 be the global optimum⁴ and let $Z \subseteq X$ be a subset of the set of configurations with $x^0 \in Z$, we say that the problem (X, \geq) is **locally decomposable** in Z by the scheme D if $Z \subseteq \Psi(x^0, D)$. If $Z = X$ we say that the problem is **globally decomposable**⁵ by the scheme D .

Thus, according to the previous proposition, if the problem is locally decomposable in Z there must exist a best-neighbour-path for decomposition D leading to the global optimum from every configuration in Z . If it is globally decomposable such a path must exist from every starting configuration.

Among all the decomposition schemes of a given problem, we are especially interested in those for which the global optimum becomes reachable from any starting configuration. One such decomposition always exists, and is the degenerate decomposition $D = \{\{1, 2, 3, \dots, N\}\}$ for which of course there exists only one local optimum and it coincides with the global one. But obviously we are interested in – if they exist – smaller decompositions 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 co-ordinated by market like selection, with the property that such selection processes can eventually lead to optimality from any starting condition. On the contrary, finer decompositions will 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.

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

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

Proof: we prove it by providing an example. Consider a problem whose globally optimum configuration is the string $x^0 = x_1^0 x_2^0 \dots x_N^0$ and whose second best

⁴ We remind the assumption of uniqueness of the global optimum.

⁵ A special case of global decomposability, which is generalised here, is presented in Page (1996) and is called dominance. In our terminology, a block configuration $x^h(d_k)$ is dominant when $x^h(d_k) \wedge x^j(d_{-k}) > x^j(d_k) \wedge x^j(d_{-k})$ for every configuration $x^j(d_{-k})$. That is when it is always preferred to all the other configurations of that block irrespective of the configuration of the rest of the string.

configuration is $x^j = x_1^j x_2^j \dots x_N^j$ where $x_h^j = |1 - x_h^0| \quad \forall h = 1, 2, \dots, N$. It is obvious that the global optimum can be reached from the second best only by mutating all components together, while any other mutation gives an inferior configuration.

The following proposition establishes a rather obvious but important property of decomposition schemes: as we climb into the basin of attraction of a local optimum for a decomposition D which is not the finest one, then finer decomposition schemes can be introduced which allow to reach the same local optimum.

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

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

Minimum size decomposition schemes can be found recursively with the following procedure: let us re-arrange all the configurations in X by descending rank $X = \{x^0, x^1, \dots, x^{2^N-1}\}$ where $x^i \geq x^{i+1}$.

The algorithm can be described informally⁶ as follows:

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

Let us finally provide an example for illustration:

⁶ The complete algorithm is quite lengthy to describe in exhaustive and precise terms. Its Pascal implementation is available from the author upon request.

Example: the following table contains an hypothetical ranking (where 1 is the rank of the most preferred) of configurations for $N=3$

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

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

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

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

Note that the worst local optimum has a larger basin of attraction⁷ as it covers all possible configurations except the global optimum itself. Thus, only a search which starts at the global optimum will (trivially) stop at the global optimum itself with certainty, while for 4 initial configurations search might end up in either local optima (depending on the sequence of mutations) and for the remaining 3 initial configurations search will end up at the worst local optimum with certainty.

Using the notion of dominance (cf. Page (1996)) it is possible to find out that the only dominant block-configuration is actually the globally optimum string itself, corresponding to the degenerate decomposition scheme of size 3 $D=\{\{1,2,3\}\}$. So apparently no decentralised search structure allows to locate always the global optimum from every starting configuration.

Actually this is not true: also the decomposition scheme $D=\{\{1,2\},\{3\}\}$ allows decentralised selection to climb up to the global optimum. For instance if we start from configuration 111 we can first locate 011 (using block $\{1,2\}$) then 010 (using block $\{3\}$) and finally 100 (again with block $\{1,2\}$); or alternatively we can locate 110 (using block $\{3\}$) and 100 (with block $\{1,2\}$). It can be easily

⁷ Kauffman (1993) provides some general properties of one-bit-mutation search algorithms (equivalent to our bit-wise decomposition schemes) on string fitness functions with varying degrees of interdependencies among components. In particular, he finds that as the span of interdependencies increases, the number of local optima increases too, while the size of the basin of attraction of the global optimum shrinks.

verified that the same blocks do actually “work” for all other starting configurations. The algorithm just presented will find this decomposition.

3. Near-decomposability.

When building a decomposition scheme for a problem we have looked so far for perfect decomposability, in the sense that we require that all blocks can be 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 very stringent a requirement: even when interdependencies are rather weak, but diffused across all components, we easily tend to observe problems for which no perfect decomposition exists. For instance in Kauffman’s NK landscapes (cf. Kauffman (1993)), already for such small values of K as 1 or 2 - that is for highly correlated landscapes - the above described algorithm finds only decomposition schemes of size N or just below N .

We can soften the requirement of perfect decomposability into one of near-decomposability: we do not want the problem to be decomposed into completely separated sub-problems, i.e. sub-problems which fully contain all interdependencies, but we want 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 solutions⁸. In other words we construct “near-decompositions” which give a precise measure of the trade-off between decentralisation and optimality: higher degrees of decentralisation and market coordination, and therefore higher speed of adaptation, can be obtained at expenses of the optimality of the solutions which can be reached.

Let $X_\mu = \{x^0, x^1, \dots, x^{\mu-1}\}$ with $0 \leq \mu \leq 2^N - 1$ be the set of the **best μ 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 **μ -decomposition** for the problem.

⁸ This procedure allows to deal also with the case of multiple global optima and thus we can now drop the assumption of a unique global optimum.

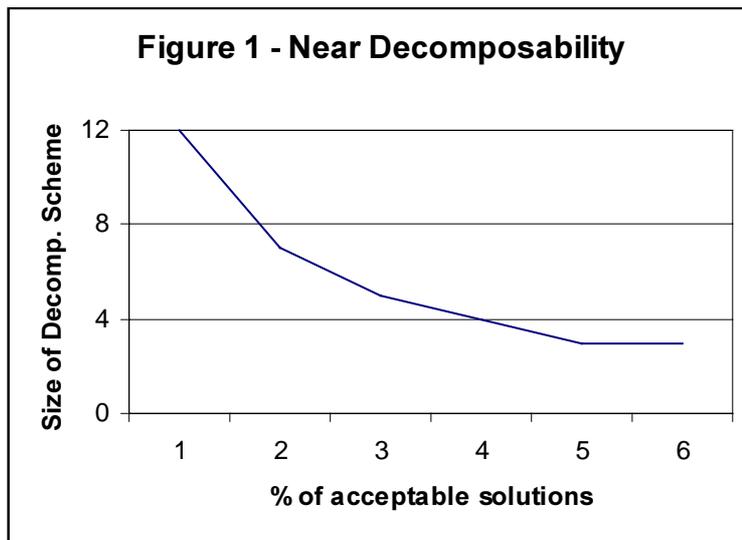
μ -decompositions of minimum size can be found algorithmically with a straightforward generalisation of the above algorithm which computes minimum size decompositions schemes for optimal decompositions.

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

Proposition 4: if D_μ is a minimum size μ -decomposition, then $|D_\mu|$ is monotonically weakly decreasing in μ .

Proof: if $\mu=2^N-1$ then X_μ includes all configurations and it is trivially reachable for any decomposition, including the finest $D_\mu=\{\{1\},\{2\},\{3\},\dots,\{N\}\}$ with $|D_\mu|=1$. If $\mu=1$ X_μ includes only the global optimum, thus the size of the minimum size decomposition is $1 \leq |D_\mu| \leq N$. We still have to show that it cannot be $|D_{\mu+1}| > |D_\mu|$: if this was the case X_μ could not be reached from $X_{\mu+1}$ for decomposition D_μ , but this contradicts the assumption that X_μ is reachable from any configuration in X for decomposition D_μ .

The latter proposition shows that higher degrees of decomposition and decentralisation can be attained by giving up optimality and allows to provide a precise measure for this trade-off. In order to provide an example we generated random problems of size $N=12$ all characterised by $|D|=12$ (i.e. they are not decomposable). The figure below shows the sizes of the minimum size decomposition schemes as we vary the number μ of acceptable configurations (average on 100 random landscapes).



4. Speed and accuracy of search: some consequences for organisational structures.

The trade-off outlined in the previous section between decomposability, reduction of complexity and speed of search on one side and optimality on the other, enables us to discuss some interesting evolutionary properties of various organisational structures competing in a given problem environment. The properties and algorithms analysed in section 3 allow us to build problems which can be decomposed with any decomposition scheme decided by the modeller. We can thus run simulations where various organisational structures compete in finding solutions to a problem whose characteristics are entirely controlled by the experimenter.

In this section we briefly discuss how such simulations have been built and the main results they produced⁹. First of all, in order to reduce the space of possible decompositions we have supposed that only decompositions which are partitions of the problem into sub-problems of the same size are possible and that only organisational structures which fulfil this constraint are viable. For instance, if $N=12$ (as in most simulations) only the following 6 decompositions, named after their size, are possible:

⁹ For reasons of space we can only give a short summary of all the simulations. Programs and detailed results can be obtained directly from the author upon request.

D1= {{1},{2},{3},{4},{5},{6},{7},{8},{9},{10},{11},{12}}

D2= {{1,2},{3,4},{5,6},{7,8},{9,10},{11,12}}

D3= {{1,2,3},{4,5,6},{7,8,9},{10,11,12}}

D4= {{1,2,3,4},{5,6,7,8},{9,10,11,12}}

D6= {{1,2,3,4,5,6},{7,8,9,10,11,12}}

D12= {{1,2,3,4,5,6,7,8,9,10,11,12}}

Problems characterised by one of these decomposition schemes can be created and populations of agents, each of which characterised by one of these decomposition schemes, compete in a simple selection environment to find better solutions. Such competition works as follows:

1. a problem is randomly generated whose minimum size decomposition scheme is one of the 6 possible;
2. a population of agents is created: each agent is characterised by one of the 6 possible decomposition schemes and is located in a randomly chosen configuration (normally we have used populations of 180 agents, 30 for each possible decomposition scheme);
3. each agent picks a randomly chosen block and by mutating at least one and up to all bits within such a block generates a new configuration: if the latter is preferred to the previous one the agent moves to it, otherwise stays put;
4. at given time intervals, all the agents are ranked: the ones located on the worst configurations are deleted from the populations and substituted by copies of the agents located on the best configurations. Such copies inherit the same decomposition scheme of the parent, but are positioned on a different, randomly chosen, configuration.

Thus we have a selection environment in which decompositions compete and are reproduced from an initial population in which 1/6 of the decompositions are the “right” ones and the others are wrong.

The main results can be summarised as follows. First of all, it is not necessarily the “right” decomposition which invades the population: as the size of the right decomposition becomes big enough, agents characterised by decompositions which are finer than the right one tend to prevail. In fact, only agents with the right decomposition can find the global optimum wherever they start from, but their search process is very slow and can be invaded by agent which cannot reach the global optimum, but do indeed reach good local optima relatively fast. Indeed potential optimisers can die out before they reach even good solutions. This result is even stronger in problems that we could define “modular”, i.e. characterised 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.

In general, simulations show two persistent features, which are present in all but the most simple (highly decomposable) problems: a persistent sub-optimality and a persistent diversity of agents, both in terms of the configurations achieved and in terms of the decomposition schemes which define them. This is the outcome of the multiplicity of local optima which characterise environments with high degrees of interdependencies (cf. Levinthal (1997)). Simulations tend therefore to support the view that heterogeneity and sub-optimality of organisational structures can indeed be a persistent feature of organisational evolution.

We have also run other simulations in which we have, at given intervals, changed the current problem with one having exactly the same structure in terms of decomposability, but with different, randomly generated, orders. This can be taken as a condition of uncertainty: for instance consumers still have changing preferences over a stable set of characteristics. Interestingly enough, it turns out that even with totally decomposable problems, as the change of the order becomes more frequent, the population is entirely invaded by agents characterised by coarser and coarser decompositions, and at the limit by agents which do not decompose at all. It seems therefore that growing uncertainty has similar consequences of growing interdependency.

Finally, two more points are worth considering. First of all, higher degree of decomposition allow the selection process to work effectively with less underlying variety. Blocks of size k can be optimised by selecting upon 2^k types of individuals: as k grows the variety requirement becomes stronger and stronger and thus less plausible.

Secondly, it must be pointed out that decentralised market coordination mechanisms can indeed exploit the advantages of parallelism and increase the speed of adaptation, but such parallelism can prevent important reductions of complexity in systems composed by nested or overlapping subsystems¹⁰. Consider as an example the extreme case of a problem which can be decomposed by nested blocks as in $D = \{\{1\}, \{1,2\}, \{1,2,3\}, \dots, \{1,2,3, \dots, N\}\}$. The size of such a decomposition scheme is N and thus markets working in parallel will face a problem of maximum complexity. However the problem would have minimum complexity if it were solved according to the following sequence: first block $\{1\}$ can be optimised, then block $\{1,2\}$ can exploit the optimal configuration of block $\{1\}$ and optimise only the second component, and so on... Note that, in order to exploit such a reduction of complexity, a precise unique sequence must be followed: it seems quite unlikely that it could spontaneously emerge starting from isolated markets working in parallel.

We already mentioned that there exists a trade-off between speed of increments and optimality, in that highly decomposed strategies increment quickly but are frequently bound to get stuck in local optima. However, there is also

¹⁰ A similar property is measured by Page (1996) as the “ascent size” of a problem.

another interesting finding: in the very early stages, when simulated agents find themselves in very poor areas of the landscape they are *slower* than non decomposed strategies. This is due to the fact that their moves out of “wells” have a limited range as compared to non decomposed strategies. This result suggests an interesting interpretation of the historical development of real world market organizations. Typically, when a new product is just invented a vertically integrated firm is the most common producer. Later a process of vertical disintegration brings to introduce market interactions between producers who limit their activities to portions of the whole production process.

5. 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 sometimes more importantly – 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 than searching possibilities generated within a given representation: decentralisation can be increased if more “powerful” representations are built.

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

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

\blacktriangleright is a **preference relation** over possible configurations.

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 \blacktriangleright is a “subjective” one which does not necessarily coincide with the “objective” one \geq .

We suppose that agents do not know the “objective” problem, but only its representation and therefore that it is the space defined by such representation which is being searched with a given decomposition scheme.

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

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

The next two propositions 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 6) and/or the preference relation (proposition 7) we can transform any problem into one of minimum complexity. Comparing proposition 5 with propositions 6 and 7 we are thus led to conclude that acting on the representation can be a more powerful problem solving strategy than acting on the solution algorithm for a given representation.

Proposition 6: given any problem (X, \geq) , it admits an encoding Ξ such that it can be solved optimally with the decomposition scheme of minimum complexity $D=\{\{1\},\{2\},\{3\},\dots,\{n\}\}$.

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

Consider the mapping $\Phi:X\rightarrow\mathbb{N}$ from the set of configurations into the set of non-negative integers so defined:

$$\Phi(x_0)=0; \Phi(x_1)=1; \dots; \Phi(x_{2^N})=2^N \quad \text{with } x_{2^N} \geq x_{2^{N-1}} \geq \dots \geq x_1 \geq x_0$$

Define now the encoding $\Xi^*(x_i)=bin_N[\Phi(x_i)]$ where bin_N is a function which maps an integer into a string of length N which gives its binary encoding (filling with 0's the missing bits). It is now very easy to verify, because of the properties of binary encoding (i.e. every mutation from 0 to 1 always produces a higher number), that Ξ^* is an encoding which satisfies proposition 6.

Proposition 7: given any problem (X, \geq) and any encoding Ξ , there exist a preference relation \triangleright such that the problem can be solved optimally with the decomposition scheme of minimum complexity $D=\{\{1\},\{2\},\{3\},\dots,\{n\}\}$.

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

Any preference relation such that $l_1 l_2 \dots l_i^* \dots l_n \triangleright l_1 l_2 \dots l_i \dots l_n$
 $\forall l_i \neq l_i^* \quad i = 1, 2, \dots, n$
satisfies proposition 7.

There is a different and more interesting sense in which problem representations can be dealt with: it concerns the nature and representation of the underlying structure of interdependencies. In Kauffman's NK model interdependencies (epistatic interactions) are assumed to have a random nature, reflecting our ignorance on how in the biological realm epistatic interactions within the genome do actually map into phenotypic characteristics. This allows Kauffman to derive statistical properties of the population of random landscape, but these properties do not necessarily reflect the properties of landscapes which are of interest in some other domain in which we know more about the nature of interdependencies. Short, for the time being, of a systematic analysis of the mapping between the nature of interdependencies and the structure of the resulting landscape, in the following we provide an example of a large family of problems in which, in spite of the existence of strong interdependencies, there exists a "natural" complete representation of the problem which makes it fully decomposable.

The example we are going to develop is represented by puzzles or games which can be expressed in the form of deterministic transitions between states (for a more extensive treatment of these puzzles and for some graph theoretic results which are equivalent to ours see Egidi (2000)). States are characterised by a full description of the configurations of the game (layout of cards in a solitaire, positions of colours in a Rubik cube, etc.) including, in games which involve more than one player, a variable which simply indicates whose turn is to move. We assume that the game does not involve any random component in the transitions: it might indeed involve hidden information and randomness with respect to the initial configuration, as it usually happens in cards games, but a given move applied to a given configuration must always cause the same transition.

Such puzzles can be fully described by:

1. A finite **set of states** $S \cup W$, where $S = \{s_1, s_2, \dots, s_n\}$ is the set of non-terminal states in which some move has to be performed and $W = \{w_1, w_2, \dots, w_m\}$ the set of terminal or "winning" states, i.e. those in which the puzzle is successfully solved or the game is won and no further moves should be performed. A subset $B \subseteq S$ of the possible initial states (for instance chess has only one starting state, while in the Rubik cube all non-winning configurations can be starting ones). We consider only games in which W can be reached from every state in B .
2. A finite **set of moves** $M = \{m_1, m_2, \dots, m_k\}$
3. A **transition matrix** $T = [t_{ij}]$ of dimension $n \times k$ in which the element $t_{ij} \in S \cup W$ is the state which is reached from state i when move j is performed. Some moves can be illegal at some states (i.e. violate the rules of the game), thus the matrix T may contain empty cells. Alternatively, illegal moves can be modelled by introducing an "illegal state" s_0 which is reached from any other state when an illegal move is performed and which terminates the game (no further moves are allowed). We call $S_0 = \{s_0\} \cup S$

A complete and deterministic playing **program** P is a string of length n which specifies one and only one move for all the non-winning states:

$$P = m_{i_1} m_{i_2} \dots m_{i_n} \text{ with } m_{i_j} \in M$$

Given a starting state $s_i \in B$, the program P determines a unique sequence of states which can end up in:

1. one of the winning states w_j , in this case we also compute the length of the sequence of moves which the program has performed to reach w_j (we call it N_moves);
2. the illegal state s_0 if an illegal move has been performed by the program;
3. an infinite loop if the sequence visits twice a state which is neither winning nor illegal.

We call **performance** of the program P , with s_j as a starting state, the couple of integers (N_fail, N_moves) . N_fail is set to 1 if the program fails to reach one of the winning states or loops. N_moves is the number of moves that a program takes to reach one of the winning states.

If we compute the performance of the program for all the possible starting states, we can compute the global performance of P as the couple

$$\left(\sum_{s_i \in B} N_fail_{s_i}, \sum_{s_i \in B} N_moves_{s_i} \right).$$

It is natural to assume the following lexicographic **preference order** on programs:

we say that program P_i is strictly preferred to program P_j and write $P_i \succ P_j$ if:

- a) $N_fail(P_i) < N_fail(P_j)$
- or
- b) $N_fail(P_i) = N_fail(P_j)$ and $N_moves(P_i) < N_moves(P_j)$

we say that P_i and P_j are indifferent and write $P_i \approx P_j$ if:

- a) $N_fail(P_i) = N_fail(P_j)$
- and
- b) $N_moves(P_i) = N_moves(P_j)$

In this way we introduce a complete ranking over the space of all k^n programs, we refer to it as the **landscape of complete programs**.

Optimal programs: we can partition S into subsets which are connected to W by no less than $1, 2, \dots, h$ moves¹¹.

Such a partition can be constructed by including in a set S_1 all and only the states from which W can be reached with a single move, in a set S_2 all and only the states from which S_1 can be reached with a single move, and so on...

More formally, the partition can be constructed recursively:

```

1 - build  $S_1 = \{s_i\}$  for all the  $s_i$  such that there exists a  $t_{ij} \in W$ ,
 $\forall j = 1, 2, \dots, k$ 
2 - set  $h = 1$ 
3 - repeat
     $h = h + 1$ 
    build  $S_h = \{s_i\}$  for all the  $s_i \in S \setminus S_1 \setminus \dots \setminus S_{h-1}$  such that there exists a
 $t_{ij} \in S_{h-1}$ ,  $\forall j = 1, 2, \dots, k$ 
    until  $S \setminus S_1 \setminus S_2 \setminus \dots \setminus S_h = \emptyset$ 

```

the resulting h is the length of the optimal program(s).

Decomposing and searching the space of complete programs.

In this section we show that the landscape of complete programs as defined above:

¹¹ Of course for every state which is connected to W there must exist a path to W of at most $n-1$ moves (i.e. which at most visits all the non-winning states), thus $h \leq n$.

1. is generally highly non-correlated, i.e. characterised by strong epistatic interdependencies in the sense of Kauffman (1993) (i.e., Kauffman's jargon, its K is relatively high);
2. has always cover size h and ascent size 1 in the sense of Page (1996);
3. is always fully decomposable, in the sense defined in the previous sections, i.e. one-bit mutation with selection always locates the global optimum.

The first statement can be easily checked numerically by building puzzles as defined above and computing the correlation structure of the resulting landscapes of complete programs. We found that using Weinberger's autocorrelation measure (cf. Weinberger (1991) and Kauffman (1993), p. 63), such landscapes have very low correlation, that is the same correlation values of NK landscapes whose K is very close to $N-1$.

For the following two statements we provide instead a formal proof in propositions 8 and 9.

Proposition 8: the landscape of complete programs for a puzzle (S,W,M,T) has ascent size 1 and cover size h (as defined in Page (1996)).

Proof: consider the partition $S=S_1 \cup S_2 \cup \dots \cup S_h$ introduced above and consider a state $s_i \in S_1$. Since $s_i \in S_1$ there must exist a move m_i^* which connects s_i directly to W , therefore mutating the move in the i -th position into m_i^* always improves the performance of a program regardless the configuration of the rest of the program. This implies that the hyperplane $### \dots m_i^* \dots ###$ must be a dominant one.

Consider now the set S_2 and suppose that $s_j \in S_2$ can be connected directly to s_i , then the hyperplane $## \dots m_{i_j}^* \dots m_i^* \dots ##$ must also be dominant, and so on until the subset S_h is connected to S_{h-1} . Thus if h is the length of the optimal program(s), the landscape of complete programs has a minimum size cover h and ascent size 1.

Proposition 9: the landscape of complete programs for a puzzle (S,W,M,T) is always fully decomposable, i.e. does not have local but not global optima with respect to all possible one-bit mutations.

Proof: this proposition can be proven as a corollary of proposition 8. Assume for simplicity, but without loss of generality, that there exist a unique globally optimum program:

$$P^* = m_{i_1}^* m_{i_2}^* \dots m_{i_n}^*$$

We show that for any other program $P = m_{i_1} m_{i_2} \dots m_{i_n}$ with at least one $m_{i_j} \neq m_{i_j}^*$ there always exists at least one mutation which gives a program of higher rank.

Consider again the partition $S = S_1 \cup S_2 \cup \dots \cup S_h$, if P is such that at least one state in S_1 is not directly connected to W , then the performance of P can be improved by making one mutation in such a way as to connect it directly to W , regardless the current state of the rest of the program. If instead all states in S_1 are already directly connected to W , then we can repeat the same argument for the connections from S_2 to S_1 , and so on until we connect states in S_h to the states in S_{h-1} .

In order to better understand the implications of proposition 9, consider the case in which the current program P is not the optimal one. We may have that some moves in P , in spite of not being optimised, have no performance-improving mutation available, which amounts to saying that these moves have non-linear epistatic links with other moves. Proposition 9 shows that nevertheless there always exists at least one other move in the program P which can undergo a performance improving mutation and, moreover, sooner or later these mutations elsewhere in the program will make a performance improving mutation possible also for those moves for which currently it does exist.

References

- Alchian, A. and Demsetz, H. (1972), Production, information costs and economic organization, *American Economic Review*, vol. 62, pp. 772-795.
- Coase, R.M. (1937), The Nature of the Firm, *Economica*, vol. 4, pp. 386-405.
- David, P.A. (1985), Clio and Economics of QWERTY, *American Economic Review*, vol. 75, pp.332-337.
- Egidi, M. (2000), Biases in Organizational Behaviour, CEEL, University of Trento, mimeo.
- Fontana, W. And Buss, L. (1994), The arrival of the fittest: toward a theory of biological organization, *Bulletin of Mathematical Biology*, vol. 56, pp. 1-64.
- Frenken, K., Marengo L. and M. Valente (1999), "Interdependencies, nearly-decomposability and adaptation" in Th. Brenner (ed.), *Computational Techniques for Modelling Learning in Economics*, Kluwer Academics, pp.145-165.
- Hurwicz, L. (1971), Centralization and decentralization in economic processes, in: Eckstein, A. (ed.) *Comparison of economic systems*, Berkeley, University of California Press.
- 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.
- Langlois, R. and Robertson, P. (1989), Explaining vertical integration: lessons from the american automobile industry, *The Journal of Economic History*, vol. 49, pp. 361-375.
- Levinthal, D. (1997), Adaptation on rugged landscapes, *Management Science*, vol. 43, pp. 934-950.
- Page, S.E. (1996), Two measures of difficulty, *Economic Theory*, vol. 8, pp.321-346.
- Robbins, L. (1932) *An Essay on the Nature and Significance of Economic Science*, London, Mac Millian
- Simon, H. (1969), *The Sciences of the Artificial*, Cambridge Mass., MIT Press.
- Simon, H. (1983), *Reason in Human Affairs*, Stanford, Stanford University Press.
- Simon, H. (1991), Organizations and markets, *Journal of Economic Perspectives*, vol 5, pp. 25-44.

Smith, A. (1976), *An inquiry into the nature and causes of the wealth of nations*, Oxford, Clarendon.

Stadler, B., P. Stadler, G. Wagner and W. Fontana (2000), *The topology of the possible: Formal spaces underlying patterns of evolutionary change*, Santa Fe NM, Santa Fe Institute Working Paper #00-12-070.

Weinberger (1991), A more rigorous derivation of some properties of uncorrelated fitness landscapes, *Journal of Theoretical Biology*, vol.134, pp.

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