

**KNOWLEDGE DISTRIBUTION AND COORDINATION
IN ORGANIZATIONS: ON SOME SOCIAL ASPECTS OF
THE EXPLOITATION VS EXPLORATION TRADE-OFF**

Luigi Marengo
*Department of Economics,
Trento University, Trento, Italy*

RR-94-13
November 1994

Reprinted from the *Revue Internationale de Systémique*,
Vol. 7, No. 5, 1993, pp. 553 à 571

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
Laxenburg, Austria

Research Reports, which record research conducted at IIASA, are independently reviewed before publication. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work.

Reprinted, with permission from Dunod Publishers, from the *Revue Internationale de Systémique*, Vol. 7, No. 5, 1993, pp. 553 à 571.
Copyright © Afcet Gauthier-Villars

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage or retrieval system, without permission in writing from the copyright holder.

Printed by Novographic, Vienna, Austria.

Systems Analysis of Technological and Economic Dynamics

This new research project at IIASA is concerned with modeling technological and organizational change; the broader economic developments that are associated with technological change, both as cause and effect; the processes by which economic agents – first of all, business firms – acquire and develop the capabilities to generate, imitate, and adopt technological and organizational innovations; and the aggregate dynamics – at the levels of single industries and whole economies – engendered by the interactions among agents which are heterogeneous in their innovative abilities, behavioral rules and expectations. The central purpose is to develop stronger theory and better modeling techniques. However, the basic philosophy is that such theoretical and modeling work is most fruitful when attention is paid to the known empirical details of the phenomena the work aims to address: therefore, a considerable effort is put into a better understanding of the ‘stylized facts’ concerning corporate organization routines and strategy; industrial evolution and the ‘demography’ of firms; patterns of macroeconomic growth and trade.

From a modeling perspective, over the last decade considerable progress has been made on various techniques of dynamic modeling. Some of this work has employed ordinary differential and difference equations, and some of it stochastic equations. A number of efforts have taken advantage of the growing power of simulation techniques. Others have employed more traditional mathematics. As a result of this theoretical work, the toolkit for modeling technological and economic dynamics is significantly richer than it was a decade ago.

During the same period, there have been major advances in the empirical understanding. There are now many more detailed technological histories available. Much more is known about the similarities and differences of technical advance in different fields and industries and there is some understanding of the key variables that lie behind those differences. A number of studies have provided rich information about how industry structure co-evolves with technology. In addition to empirical work at the technology or sector level, the last decade has also seen a great deal of empirical research on productivity growth and measured technical advance at the level of whole economies. A considerable body of empirical research now exists on the facts that seem associated with different rates of productivity growth across the range of nations, with the dynamics of convergence and divergence

in the levels and rates of growth of income in different countries, with the diverse national institutional arrangements in which technological change is embedded.

As a result of this recent empirical work, the questions that successful theory and useful modeling techniques ought to address now are much more clearly defined. The theoretical work described above often has been undertaken in appreciation of certain stylized facts that needed to be explained. The list of these ‘facts’ is indeed very long, ranging from the micro-economic evidence concerning for example dynamic increasing returns in learning activities or the persistence of particular sets of problem-solving routines within business firms; the industry-level evidence on entry, exit and size-distributions – approximately log-normal; all the way to the evidence regarding the time-series properties of major economic aggregates. However, the connection between the theoretical work and the empirical phenomena has so far not been very close. The philosophy of this project is that the chances of developing powerful new theory and useful new analytical techniques can be greatly enhanced by performing the work in an environment where scholars who understand the empirical phenomena provide questions and challenges for the theorists and their work.

In particular, the project is meant to pursue an ‘evolutionary’ interpretation of technological and economic dynamics modeling, first, the processes by which individual agents and organizations learn, search, and adapt; second, the economic analogs of ‘natural selection’ by which interactive environments – often markets – winnow out a population whose members have different attributes and behavioral traits; and, third, the collective emergence of statistical patterns, regularities, and higher-level structures as the aggregate outcomes of the two former processes.

Together with a group of researchers located permanently at IIASA, the project coordinates multiple research efforts undertaken in several institutions around the world, organizes workshops and provides a venue of scientific discussion among scholars working on evolutionary modeling, computer simulation and non-linear dynamical systems.

The research will focus upon the following three major areas:

1. Learning Processes and Organizational Competence.
2. Technological and Industrial Dynamics
3. Innovation, Competition, and Macrodynamics

Preface

This paper is a theoretical contribution to the work being carried out on the topic “Learning Processes and Organizational Competence”, one of the research areas of the Systems Analysis of Technological and Economic Dynamics project at IIASA.

The paper tackles one of the key problems in understanding the dynamics of organizational capabilities: that of coordination of distributed knowledge, and, in particular, it focuses on the trade-off between diversity and commonality of knowledge in organizational learning processes. The methodology employed is rather new and seems particularly promising for this kind of analysis. It is a simulation model based on Classifiers Systems, a methodology developed in artificial intelligence, with the purpose of modeling learning as an adaptive process of category formation. The organization is represented as a collection of adaptively learning agents, linked to each other by information flows which define the very structure of the organization (hierarchical vs. horizontal information flows). Organizational learning is seen therefore as a co-evolutionary process in which the characteristics and directions of information flows set a particular balance between decentralized autonomous learning and the social construction of shared conventions and routines.

In spite of the highly preliminary and exploratory nature of the model, simulation results cast a new light on the role that the structure of the organization plays in defining its capabilities to adapt to environments characterized by different degrees of stationarity and predictability.

Giovanni Dosi

Leader

Systems Analysis of Technological
and Economic Dynamics Project

**KNOWLEDGE DISTRIBUTION AND COORDINATION
IN ORGANIZATIONS: ON SOME SOCIAL ASPECTS OF
THE EXPLOITATION VS EXPLORATION TRADE-OFF (*)**

Luigi MARENGO¹

Abstract

This paper puts forward a preliminary investigation of the relationship between the distribution of knowledge and the capability of learning and adapting to changing environmental conditions in organizations. The main focus of the paper is on the trade-off each organization faces between commonality of knowledge, which enables coordination, and diversity of knowledge, which on the contrary favours learning and discovery of new ways of doing things. By means of a simulation model the paper compares the performance, in terms of coordination and learning, of different organizational designs, characterized by the way in which knowledge is distributed among the members of the organization and by the way coordination is achieved through centralized or decentralized coordinating devices.

Résumé

Cet article propose une analyse préliminaire de la relation entre la distribution du savoir et la capacité d'apprentissage et d'adaptation des organisations face aux changements de leur environnement. L'article focalise sur l'arbitrage qui existe pour chaque organisation entre l'aspect commun du savoir, permettant la communication et la diversité des savoirs favorisant l'apprentissage et les nouvelles opportunités de modalités d'action. En utilisant un modèle de simulation nous examinons la performance en termes de coordination et d'apprentissage des

1. Dipartimento di Economia, Università di Trento, Via Inama, 1, 38100 Trento, Italy,
Tel: +39-461-882201, Fax: +39-461-882222, email: marengo@itncisti.bitnet

(*) I would like to thank Giovanni Dosi and Massimo Egidi for helpful discussions. Preliminary versions of this paper were presented at the seminar "Institutional Change and Network Evolution", Stockholm, June 16-18, 1993, and at the E.M.O.T. Workshop, Strasbourg, September 23-25, 1993. Comments made there, especially by R. Cowan and P. Llerena greatly contributed to improvements.

différentes conceptions organisationnelles caractérisées par la manière dont le savoir est distribué parmi les membres de l'organisation et par la manière dont opère la coordination.

I. INTRODUCTION

In every kind of economic organization, be it an organization *stricto sensu*, or a network of organizations and/or individuals, or a market, there exists a trade-off between commonality and diversity of knowledge. Sharing a common and homogeneous knowledge basis is a necessary condition for agents to communicate and coordinate their actions. But, on the other side, if all the members of an organization were sharing exactly the same body of individual knowledge no form of collective learning from each other would be possible and the organization would ultimately lose its capability of learning and adapting to new environmental conditions.

This is clearly a typically evolutionary argument: collective adaptation and learning require diversity (mutation) but also mechanisms which guarantee the necessary overall coherence (selection). Ultimately, each economic organization can be considered as an evolutionary system which implements a particular balance between mechanisms of variation and mechanisms of selection on what constitutes the organizational knowledge basis.

In economic organizations, this trade-off between commonality and diversity of knowledge is also strictly connected to the trade-off between exploitation and exploration (*cf.* March, 1991): organizations always face the dilemma between concentrating their resources on the exploitation of the knowledge which is already available to them and the exploration of new possibilities. Both exploitation and exploration are necessary for the survival of an organization. Without exploration of new possibilities, the organization would find itself trapped into sub-optimal states and would eventually become ill-adapted to changing environmental conditions. But organizations which devote all their resources to the exploration of new possibilities will face too high a degree of risk, and even in case of successful discoveries they fail to exploit the knowledge they acquire and will systematically perform worse than followers and imitators.

March stresses the importance of the social context in which organizational learning takes place. A "distinctive feature of the social context... is the mutual learning of an organization and the individuals in it. Organizations store knowledge in their procedures, norms, rules, and forms. They accumulate

such knowledge over time, learning from their members. At the same time, individuals in an organization are socialized to organizational beliefs" (March, 1991, p. 73).

Within organizations, these two processes normally coexist and interact at different levels: one of the strengths of organizations is their capability of flexibly combining procedures for selection and procedures for innovation. Fast-learning and slow-learning individuals and departments can coexist. Innovation itself can become a largely routinized process, though uncertain in its outcome. Learning by doing can add exploratory value to normally exploitative activities.

Mutual learning and the distribution of knowledge are fundamental factors which determine an organization's balance between the processes of exploration and exploitation. A high degree of differentiation of knowledge among the members of an organization increases the total amount of knowledge possessed by the organization. But differentiation makes coordination more difficult and ultimately can inhibit the social exploitation of this broad knowledge basis. On the contrary, a body of organizational knowledge which is commonly shared by all the members facilitates coordination but reduces the scope for decentralized experimentation, which could prove a vital source of organizational learning.

Hence, there exists a tension between centralization and decentralization in the organizational learning process. Firms require both centralization and decentralization to operate successfully in changing environments. Decentralization in the acquisition of knowledge is a source of variety, experimentation and, ultimately, a fundamental source of learning. But, eventually, knowledge has to be made available for exploitation to the entire organization. When agents differ with regard to their representations of the environment and their cognitive capabilities, there must exist an organizational body of knowledge which guarantees the coherence of the various learning processes. In order to cope with changing environments, the process of generation and modification of such a body of knowledge, although fed by the decentralized learning processes, has to undergo some form of centralization. Thus, a tension inevitably arises between the forces which keep the coherence of the organization and the forces which promote decentralized learning.

This paper puts forward a preliminary investigation of these issues by means of a few simulation experiments which make use of a methodology similar to the so-called classifiers systems. The next section outlines a very basic organizational decision-making problem and the simulation methodology which can model substantive learning and the formation of collective

knowledge and languages. The following section will run a few simulations of this model in which the adaptive performance of different organizational designs – characterized by different modes of intra-organizational distribution of knowledge – will be tested against simple environments characterized by varying degree of variability and predictability.

II. A COMPUTATIONAL MODEL OF LEARNING AND DECISION-MAKING

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

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

be the set of n possible states of nature and

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

the set of the k possible actions the decision-maker can undertake.

The payoff to the agent is given by a function:

$$\Pi = AXS \rightarrow R$$

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

The action the agent chooses depends obviously on the level of his or her knowledge about the state of the world. The agent's state of knowledge (or "information processing capabilities") can be represented by a collection of subsets

$$P(s_i) \subseteq S$$

where $P(s_i)$ is the set of states of the world which the agent considers as possible (or cannot tell apart) when the real state is s_i . The set of subsets $P(s_i)$ represents a sort of collection of categories which the agent employs in order to classify the environment and to find in those regularities which can be usefully exploited by his or her actions.

Standard Bayesian decision theory and game theory postulate that the collection of such categories forms a partition of the set of states of the world, ruling out in this way the possibility of substantive ignorance and/or partial knowledge¹ of some parts of the environment, that is ignorance which cannot be reduced by simply acquiring new information, but only by modifying the very categories on which information collection and

interpretation are based (more on this in Marengo, 1992 *b*). But excluding substantive ignorance implies excluding the possibility of substantive learning, that is the construction of new, more useful mental models for the interpretation of reality, which cannot be reduced to simple probability updating, *i. e.* the acquisition of new information within a constant model of the world.

More general representations of the agent's information processing capabilities than the probabilistic one can be found in the theories based on the concept of fuzzy sets (*cf.*, for instance, Shafer, 1976; Dubois and Prade, 1988). In these theories information processing capabilities of decision makers are more generally represented by subsets of the *power set* of S (partitions being a special case). Substantive ignorance, surprise, inconsistencies can be naturally modelled in this framework. Learning, seen as category modification, can be therefore represented as a search in the space of subsets of the power set of S, *i. e.* the space of models of the world.

The computational model which is here described is one the (infinite) possible models which consider learning as a movement in this space. It is based on *classifiers systems* (*cf.* especially Holland, 1975 and 1986 and Holland *et al.*, 1989) but with some substantial differences and simplifications.

The basic component of this learning system is a condition-action rule, where the execution of a certain action is conditional upon the agent's perception that the present state of the world falls in one of the categories he or she has defined in his "mental model". The condition part is a "category", that is a subset of the states of nature and is activated when the last detected state of the world falls in such a subset. Practically, the condition is a string of *n* symbols (as many as the states of the world) over the alphabet { 0,1 } and it is satisfied whenever the last state of the world corresponds to a position where a "1" appears. All in all, the condition

$$c_1 c_2 \cdots c_n \quad \text{with} \quad c_i \in \{0, 1\}$$

is satisfied when, if s_k is the last observed state of the world, we have:

$$c_k = 1$$

Thus, a set of conditions defines a subset of the power set of S. It is important to notice that each condition defines one subjective state (or category) of the world, as perceived by the agent and defines its relationship with the objective ("true") states of the world. This relationship remains anyway unknown to the decision maker, who "knows" only the subjective states.

This important point deserves an example: suppose there exist three “real” states of the world:

$$S = \{ s_1, s_2, s_3 \}$$

and the agent’s state of knowledge is represented by the following two conditions:

$$\vartheta_1 : 110$$

$$\vartheta_2 : 101$$

The agent conceives two “subjective” states of the world, ϑ_1 and ϑ_2 . The agent thinks he or she is in the former when the real state of the world is either s_1 or s_2 , whereas he or she believes to be in the latter when the real state is either s_1 or s_3 . This correspondence between subjective and objective states can only be described by an omniscient external observer and is not actually known by the agent, who ignores even the existence of the elements of the set S . All he or she knows are the two ϑ ’s.

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

$$a_1 a_2 \cdots a_k \quad \text{with} \quad a_h \in \{0, 1\}$$

has one and only one position which equals “1”:

$$a_h = 1$$

and $a_i = 0$ at every other position, meaning that the action “ h ” is chosen.

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

$$R = \{ R_1, R_2, \cdots, R_q \}$$

where:

$$R_j : c_1 c_2 \cdots c_n \Rightarrow a_1 a_2 \cdots a_n \quad \text{with} \quad c_i a_h \in \{0, 1\}$$

In addition, each rule is assigned a “strength” and a “specificity” measure. The strength measures the past usefulness of the rule, that is the payoffs cumulated every time the rule has been applied (minus some other quantities which will be specified later); the specificity measures the strictness of the condition: in our case the highest specificity (or lowest generality) value is given to a rule whose condition has only one symbol “1” and therefore is satisfied when and only when that particular state of the world occurs,

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

At the beginning of each simulation the decision maker is supposed to be completely ignorant about the characteristics of the environment he or she is going to face: all the rules initially generated have the highest generality, meaning that all their conditions are formed entirely by 1's. The action parts are instead randomly generated, to represent the fact that, because of the condition of absolute ignorance, the decision maker does not have any reason to prefer an action to another.

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

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

– *Condition matching*: a message is received from the environment which informs the system about the last state of the world. Such a message is compared with the condition of all the rules and the rules which are matched, *i. e.* those which apply to such a state of the world enter the following step.

– *Competition among matched rules*: all the rules whose condition is satisfied compete in order to designate the one which is allowed to execute its action. To enter his competition each rule makes a bid based on its strength and on its specificity. In other words, the bid of each matched rule is proportional to its past usefulness (strength) and its relevance to the present situation (specificity):

$$\text{Bid}(R_j, t) = k_1 (k_2 + k_3 \text{Specificity}(R_j)) \text{Strength}(R_j, t)$$

Where k_1 , k_2 and k_3 are constant coefficients.

The winning rule is chosen randomly, with probabilities proportional to such bids.

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

$$\text{Strength}(R_j, t+1) = \text{Strength}(R_j, t) + \text{Payoff}(t) - \text{Bid}(R_j, t)$$

– *Generation of new rules*: the system must be able not only to select the most successful rules, but also to discover new ones. This is ensured by applying “genetic operators” which, by recombining and mutating elements of the already existing and most successful rules, introduce new ones which could improve the performance of the system. In this way new rules are constantly injected into the system and scope for new opportunities is always made available.

Genetic operators generate new rule which explore other possibilities in the proximity (in a sense which I am going to define precisely) of the presently most successful ones, in order to discover the elements which determine their success and exploit them more thoroughly: the search is not completely random but influenced by the system’s past history. New rules so generated substitute the weakest ones stored in the system, so that the total number of rules is kept constant.

Two genetic operators have been used for the condition and one for the action part. The latter can be defined “local search” and is simply a mutation in the vicinity: the action included in the newly generated rule is chosen (randomly) in the close proximity of the one included in the parent rule. The interpretation of this operator is straightforward: decision makers tend to explore alternatives in the vicinity of the ones already employed.

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

– *Specification*: a new condition is created which increases the specificity of the parent one: wherever the parent condition presents a “1”, this is mutated into a “0” with a given small probability.

– *Generalization*: the new condition decreases the specificity of the parent one: wherever the latter presents a “0”, this is mutated into a “1” with a given small probability:

Specification and generalization are two possible cognitive strategies which tend to drive the learning system towards, respectively, specific rules which apply to more specific states of the world and more general rules which instead cover a wider set of states of the world. Different degrees of specification and generalizations can be simulated both by means of different combinations of these two genetic operators and by varying the coefficient k_3 with which specificity enters the bid equation: the higher this coefficient, the more highly specific rules will be likely to prevail over general ones. The simulations discussed in the rest of the paper will use a specificity coefficient to summarize the overall inclination of the system toward the search for specific rules, such

coefficient will represent both the value k_3 in the bid equation and the probability of application of the genetic operator "specification" every time the genetic operators routine is called.

III. HOMOGENEITY VS DIVERSITY OF KNOWLEDGE IN ORGANIZATIONAL LEARNING: SOME SIMULATION RESULTS

This section employs the computational model of learning outlined in the previous section in order to analyze, by running a few simulations, how different modes of knowledge distribution within an organization can influence the direction and speed of collective learning.

As already mentioned in the introduction, organizations face a trade-off between the need for coordination, which enables the exploitation of the available knowledge, and the need for expanding and modifying the available knowledge, which is an essential condition for search and exploration of new possibilities.

Coordination requires a collective knowledge basis, consciously shared by the agents involved in a given interaction. In a world – like the one postulated by standard neoclassical economics – where agents share the same model of the world or know each other's model, the only obstacle to effective coordination could derive from some form of lack, bias or strategic use of information. In a world instead where decision makers do not entirely share a given model and do not know *a priori* each other's models, the first issue becomes that of building a collective knowledge basis which enables agents to communicate effectively and eventually achieve coordination. If, for instance, one individual or part of the organization communicates to another that, to the best of their knowledge, the present state of the world is X and such communication is faithful (and known as such to the other), the meaning of such a piece of information can still be misunderstood because the receiver has different information processing capabilities from the sender's. For example the proposition "the state of the world is X" can have for the receiver a different meaning (when the considered subset of the states of the world's power set is not the same for the two agents) or even no meaning at all (when X does not exist in the receiver's information processing capabilities).

But if, on one side, allowing for diversity of knowledge opens new problems for organizational design and is a possible source of inefficiency (*cf.* also the so-called "loss of control" literature *e. g.* Calvo and Wellisz, 1978), on the other side diversity of knowledge is a fundamental source for new ideas,

new bodies of knowledge which can be acquired and exploited by the entire organization.

The model outlined in the previous section can help cast some light on this dilemma.

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

More specifically, suppose that there exist eight possible product types (states of the world), called respectively "1", "2", ... "8". The firm's production possibilities set is represented by sequences of operations which can be of two types (A and B). Such sequences have all the same length and map into a product type, which is conventionally designated by the number of operations of type A which are utilized in its production. For example the product of type "8" is produced by all and only the production processes which contain eight operations of type A. Each production process is divided into two parts (of the same length) which are carried out separately by each of two shops. The problem of the firm is therefore to forecast the product type which will be demanded by the market and to implement the correct production process by coordinating the operations of the two shops. The payoff is the following: if the firm produces the correct product type it receives a payoff of 5 units; if it does not produce the correct output it receives a negative payoff, given by the distance of the actual product type from the required one (for example, if the market demands type "7" but the firm produces type "5", it will receive the payoff 2).

This is a rather naive model, but it already represents quite a complex coordination problem; in game theoretic terms we have two players (the two shops) who have four possible strategies each (*i. e.* implementing a production process with a number of operations of type A which can vary from one to four) and can play one out of eight possible different games, each of them with a different payoff matrix.

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

of knowledge is represented by a subset of the power set of the set of states of world. Moreover they are completely ignorant at the outset (they cannot distinguish among the eight possible states of the world) and refine their knowledge structure according to their experience and their cognitive capabilities.

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




Figure 1. Organizational information flows.

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

The weights with which these three types of messages enter the shop's decision processes define the organizational balance between differentiation

and commonality of knowledge. In our model, such weights are represented by the specificity coefficients which express the agent's search for a precise model which interprets the corresponding type of message.

A high specificity coefficient for the shops' condition parts which classify messages coming from the environment (messages of type 1B in Figure 1) implies that shops are aiming at building a detailed individual model of the world. A low coefficient implies instead that shops do not care to "understand" the environment. When the coefficient is equal to zero we have an organization in which shops do not form any autonomous model of the world but rely entirely on the interpretation of the world given by the management (messages of type 1 and 2).

A high specificity coefficient for the condition part which classifies messages coming from the management (messages of type 2 in Figure 1) implies that shops attribute great importance to the correct interpretation of the coordinating messages which are sent by the management. A low coefficient implies instead that shops are not seeking careful coordination on the organizational collective knowledge. When the coefficient is equal to zero we have an organization without any form of centralized coordination, *i. e.* the management has no role.

Finally, a high specificity coefficient for the condition part which classifies messages coming from the other shop (messages of type 3 in Figure 1) implies that shops are attaching high importance to mutual, decentralized coordination. When the coefficient is equal to zero we have an organization without any form of decentralized coordination, *i. e.* no inter-shop communication.

Thus, by designing experiments with different combinations of the three specificity coefficients, it is possible to test the performance of different organizational balances between centralized and decentralized mechanism of coordination and knowledge distribution:

Simulation 1. Let us first consider a stationary environment with constant state of the world. Simulations show that:

- coordination can be achieved if there is no model-learning at all (*i. e.* specificity coefficients are all equal to zero) due to the action of selection mechanisms. Agents do not "understand" anything of the environment in which they operate, but select randomly actions until they find the good ones and then stick to them. This appears, in our simple example, as the fastest way to achieve coordination in stationary environments.

- coordination can be equally achieved, although after a greater number of iterations, when either the management or the shops are learning about

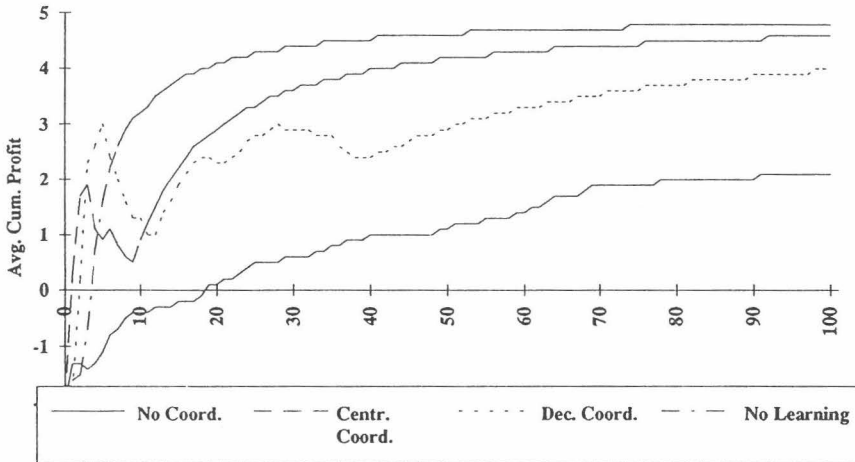


Figure 2. Stationary Environment.

the environment, provided this learning is paralleled by learning about the interpretation of coordinating messages, either centralized or decentralized;

– coordination cannot instead be properly maintained when agents are individually learning about the environment, but not about any of the coordinating devices. In this case selection mechanisms tend to make coordination temporarily emerge, but this action is counteracted by individual search, which constantly breaks such coordination.

To summarize, agents can coordinate – in stationary environments – by randomly selecting among actions and stick to a good one when it has been selected, without building any model of the environment. If instead they try to learn, *i. e.* to build such a model and constantly improving it, they need also to learn a model for the interpretation of coordinating messages: messages 1 and/or 1B are not sufficient, and messages 2 or 3 are also needed. Figure 2 summarizes these results, by plotting the average cumulated payoff of the different organizational set-ups, with different specificity coefficients on different messages.

Simulation 2. A sudden and big environmental shock is introduced in a previously stationary environment. The experiment is designed in the following way: for the first 500 iterations the product type “3” is constantly demanded, at iteration 501 the demand suddenly switches to the type “7” and remains there thereafter. The problem is therefore to reorganize radically the routines which, after 500 iterations with the same product type, are already deeply embedded in the organization.

Simulations show that a search for specific rules is in this case necessary to respond to the environmental change. In particular, a high specificity coefficient for the conditions which classify the environmental message is essential for speeding up the adaptation process.

Figure 3 plots these results.

Simulation 3. The previous experiment considered one environmental shock which require a radical change of organizational routines, now we consider instead an environment which is always changing, but according to a regular pattern. The experiment supposes that the demanded product type switches from “3” to “4” and vice versa at every iteration.

Simulations (cf. Fig. 4) show that only when the specificity coefficients on the shops’ conditions which classify environmental messages are high can the organization exploit the environmental regularity. Otherwise the organization cannot exploit this regularity and settles into constantly producing either types, with an average payoff of 2.

Simulation 4. Let us consider now continuous but unpredictable environmental changes, so that a precise forecast of the demanded product type is impossible. The product type which is being demanded varies randomly among three possible ones (“3”, “4” and “5”) at each iteration. Environmental changes are therefore confined to a subset of the possible states of the world, but are unpredictable inside such a subset: the learning problem for the organization types is therefore to define an “internal state”




Figure 3. Environmental Shock.




Figure 4. Predictably Changing Environment.

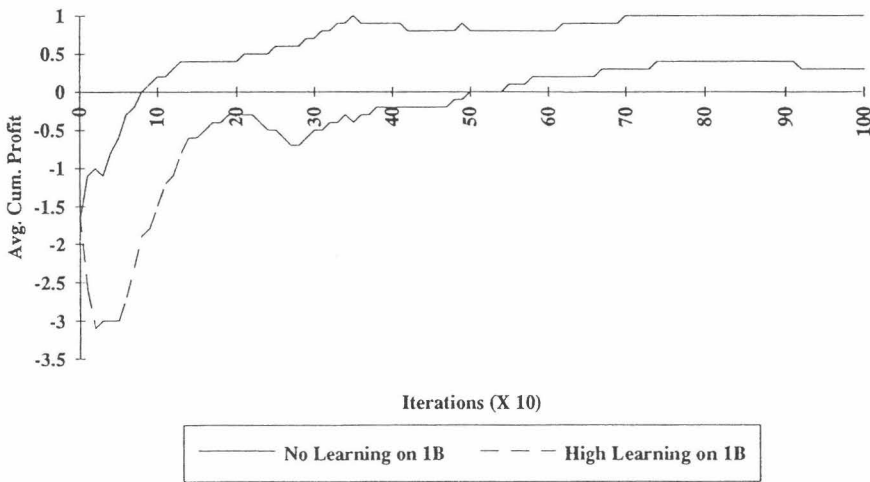


Figure 5. Unpredictably Changing Environment.

which corresponds to the three possible environmental states and link it to the constant action of always producing type “4”.

Contrary to the previous case, high specificity coefficients on the shops’ conditions which classify environmental messages reduce the organizational performance, which is instead increased by high specificity coefficients on the shops’ conditions which classify managerial signals. Results are reported in Figure 5.

By comparing the results of the previous two simulations some interesting conclusions can be drawn. To exploit a regularly changing environment a high amount of knowledge about the environment itself is required: the model must distinguish between the states of the world and connect them diachronically. It is not surprising therefore that the most appropriate organization in such circumstances is the one which, by partly decentralizing the acquisition of knowledge about the environment, can achieve higher levels of sophistication in its model of the world, provided the coordination mechanisms – which are here centralized – are powerful enough to enable the organization to solve conflicts of representations.

On the other hand, this very decentralization of the acquisition of knowledge can be a source of loss when it is more profitable for the organization to cling to a robust and stable set of routines. This situation requires strong coordination in order to make the entire organization implement coherently such a set of robust routines. Autonomous and decentralized experimentation can only disrupt such a coherence.

IV. CONCLUSIONS

Diversity and coherence are general features of every evolving system, and in social organizations they are reflected by the difficult balance between the need for coordination of individual actions and the need for diversity of individual knowledge bases from which collective learning emerges. The simple simulation model outlined in this paper shows that coherence is a non-trivial problem whenever one departs from the assumption that agents share a common representation of the environment in which they operate. On the other hand, distributed knowledge and plurality of representations is an essential mechanisms for promoting collective learning and adaptation.

Simulations show that centralized coordination mechanisms seem especially important both in simple and stationary environments and in turbulent and unpredictable ones: in both cases centralized coordination mechanisms could prove essential to prevent the organization from losing its overall coherence because of diverging individual learning processes. When instead adaptation to regularly changing environments is needed, decentralized learning processes are fundamental, provided that coordination mechanisms make the results of such individual learning processes shared by the entire organization.

Further research in this direction should investigate the properties of specific institutional mechanisms which concretely realize this balance, by

defining the social division of labour, incentive schemes and the directions and characteristics of information flows. Of course each of these issues has been already widely analyzed (cf. for instance Williamson, 1985) but in a usually static context in which the role such mechanisms play in the dynamic processes of social learning is not taken into account. Division of labour, incentives and the distribution of knowledge and information do not only affect the static efficiency of an economic system, but also its capability to generate new knowledge and adapt to changing environmental conditions. In general a trade-off between will likely exist between static efficiency and learning capabilities.

Finally, the distinction between knowledge and information which has been suggested in this paper could prove a very useful dimension for institutional analysis. For instance, a perfect market strongly centralizes informations, by reducing to a single parameter – market price – all the information agents need, but leaves knowledge of tastes and technologies (“representations”) widely distributed (cf. also Hayek, 1937). Organizations such as firms implement some form of intermediate centralization/decentralization of information and knowledge (complete centralization of knowledge being impossible because of bounded rationality, whereas centralization of information is more and more feasible thanks to new communication and data-processing technologies). Comparisons of markets vs. hierarchies from this perspective could provide new insights which cannot emerge from the static view implicit in transaction costs economics (cf. also Winter, 1982).

Notes and references

1. Among the phenomena which standard decision theory rules out with the partition postulate are surprise in the sense of Shackle (*i. e.* a state of world which was not even conceived), systematic mistakes (*i. e.* excluding the state s_i when it occurs, although such a state is considered as possible in other cases), categories which partially overlap and are not necessarily disjoint.

M. AOKI, Horizontal vs. Vertical Information Structure of the Firm, *American Economic Review*, Vol. 76, 1986, p. 971-83.

K. J. ARROW, *The Limits of Organization*, New York and London, W. W. Norton and Co., 1974.

K. J. ARROW, Informational Structure of the Firm, *American Economic Review*, Vol. 75, 1985, p. 303-307.

G. CALVO, S. WELLISZ, Supervision, Loss of Control, and the Optimal Size of the Firm, *Journal of Political Economy*, Vol. 87, 1978, p. 943-952.

- H. H. CLARK, C. R. MARSHALL, Definite reference and mutual knowledge, in A. K. JOSHI, B. L. WEBBER and I. A. SAG eds., *Elements of Discourse Understanding*, Cambridge University Press, 1981, p. 10-63.
- R. M. COASE, The Nature of the Firm, *Economica*, Vol. 4, 1937, p. 386-405.
- M. D. COHEN, Conflict and Complexity: Goal Diversity and Organizational Search Effectiveness, *American Political Science Review*, Vol. 78, 1984, p. 435-454.
- M. D. COHEN, *Adaptation of Organizational Routines*, Ann Arbor, Institute of Public Policy Studies, University of Michigan, mimeo, 1987.
- M. D. COHEN, Individual Learning and Organizational Routine: Emerging Connections, *Organization Science*, Vol. 2, 1991, p. 135-139.
- J. CRÉMER, Common Knowledge and the Co-Ordination of Economic Activities, in M. AOKI, B. GUSTAFSSON and O. E. WILLIAMSON eds., *The Firm as a Nexus of Treaties*, London, Sage Publications, 1990.
- R. M. CYERT, J. G. MARCH, *A Behavioral Theory of the Firm*, Englewood Cliff, NJ, Prentice Hall, 1963.
- G. DOSI, L. MARENGO, Some Elements of an Evolutionary Theory of Organizational Competences, in R. W. ENGLAND ed., *Evolutionary Concepts in Contemporary Economics*, Ann Arbor, University of Michigan Press, 1993.
- G. DOSI, D. J. TEECE, S. G. WINTER, Toward a Theory of Corporate Coherence, in G. DOSI, R. GIANNETTI, P. A. TONINELLI eds., *Technology and the Enterprise in a Historical Perspective*, Oxford, Oxford University Press, 1991.
- D. DUBOIS, H. PRADE, Modelling Uncertainty and Inductive Inference: A Survey of Recent Non-Additive Probability Systems, *Acta Psychologica*, Vol. 68, 1988, p. 53-78.
- J. GEANAKOPOLOS, *Game Theory without Partitions, and Applications to Speculation and Consensus*, Santa Fe Institute, Working Paper No. 90-018, 1990.
- J. GORDON, E. H. SHORTLIFFE, *A Method for Managing Evidential Reasoning in a Hierarchical Hypothesis Space*, Stanford University, Department of Computer Science, Report No. STAN-CS-84-1023, 1984.
- T. H. HAMMOND, *The Agenda of the Firm: Structure and Incentive in Institutional Design*, Conference on "Fundamental Issues in Strategy: a Research Agenda for the 1990s", December 1990, Napa, California, 1990.
- F. A. HAYEK, Economics and Knowledge, *Economica*, Vol. 4, 1937, p. 33-54.
- J. H. HOLLAND, *Adaptation in Natural and Artificial systems*, Ann Arbor, University of Michigan Press, 1975.
- J. H. HOLLAND, Escaping Brittleness: The Possibilities of General Purpose Learning Algorithms Applied to Parallel Rule-Based Systems, in R. S. MICHALSKI, J. G. CARBONELL, T. M. MITCHELL eds., *Machine Learning II*, Los Altos, Morgan Kaufmann, 1986, p. 593-623.
- J. H. HOLLAND, K. J. HOLYOAK, R. E. NISBETT, P. R. THAGARD, *Induction: Processes of Inference, Learning and Discovery*, Cambridge Mass., MIT Press, 1986.
- P. H. LOUNAMAA, J. G. MARCH, Adaptive Coordination of A Learning Team, *Management Science*, Vol. 33, 1987, p. 107-123.
- J. G. MARCH, *Decisions and Organizations*, Oxford, Basil Blackwell, 1988 a.
- J. G. MARCH, Variable Risk Preferences and Adaptive Aspirations, *Journal of Economic Behavior and Organization*, Vol. 9, 1988 b, p. 5-24.

- J. G. MARCH, Exploration and Exploitation in Organizational Learning, *Organization Science*, Vol. 2, 1991, p. 71-87.
- L. MARENGO, Coordination and Organizational Learning in the Firm, *Journal of Evolutionary Economics*, Vol. 2, 1992 a, p. 313-326.
- L. MARENGO, *Structure, Competence and Learning in an Adaptive Model of the Firm*, Papers on Economics and Evolution No. 9203, edited by the European Study Group for Evolutionary Economics, Freiburg, 1992 b.
- R. R. NELSON, S. G. WINTER, *An Evolutionary Theory of Economic Change*, Cambridge, Mass., Harvard University Press, 1982.
- G. SHAFER, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- H. A. SIMON, Rational Decision Making in Business Organizations, *American Economic Review*, Vol. 69, 1979, p. 493-513.
- M. Y. VARDI ed., *Theoretical Aspects of Reasoning about Knowledge*, Los Altos, Morgan Kaufmann, CA., 1988.
- O. E. WILLIAMSON, *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*, London, MacMillan, 1985.
- S. G. WINTER, An Essay on the Theory of Production, in H. HYMANS ed., *Economics and the World around It*, Ann Arbor, University of Michigan Press, 1982, p. 55-93.

