J Evol Econ (1992) 2:313–326

Journal of
Evolutionary
Economics

Coordination and organizational learning in the firm

L. Marengo

Science Policy Research Unit, University of Sussex, Mantell Building, Falmer, Brighton BN1 9RF, UK

Abstract. This paper discusses the role of the organizational structure in shaping the organizational learning process. Learning is modelled by means of a computational model in which search takes place in the space of problem representations and cannot be reduced to mere probability updating within a given and constant representation.

When the assumption of a unique and given representation of the problem is dropped, organizational learning emerges from the coordination of individual learning processes. Some simulations analyze the performance, in different environmental conditions, of centralized and decentralized coordination modes.

Key words: Coordination – Theory of the firm – Organizational learning

1. Introduction

Only recently has the economic theory of the firm begun to address the issue of the role of different organizational structures. Traditionally, the neoclassical theory explained all economic phenomena in terms of individual agents (households or firms) and markets. The latter, by means of the price mechanism, convey all the information which is necessary to individual decision makers for the coordination of all their interactions. Alternative modes of coordination were not conceived within such a theory.

The basic assumptions of the standard neoclassical theory are summarized with remarkable clarity by Arrow:

I would like to thank Giovanni Dosi, Keith Pavitt, Nick von Tunzelmann and two anonymous referees for their insightful suggestions. An earlier version of this paper was presented at the Third Meeting of the Joseph A. Schumpeter Society, Airlie House, Virginia, June 1990.

> The neoclassical model is founded on two concepts, which are considerably different in nature. One is the notion of the individual economic agent, whose behavior is governed by a criterion of optimization under constraints which are partly peculiar to the agent, such as production functions, and partly terms of trade with the economic system as a whole. The other is the market; here the aggregate of individual decisions is acknowledged, and the terms of trade adjusted until the decisions of the individuals are mutually consistent in the aggregate, i.e. supply equals demand [Arrow (1974), pp. 1–3].

Therefore the neoclassical theory is based on two fundamental principles: 1) the reduction of individual rationality to optimization, through the axiomatic theory of the maximization of subjective expected utility and 2) the reduction of the modes of coordination to the market, through the General Equilibrium Theory.

These two logical operations allow the neoclassical theory to omit entirely from the analysis the consideration of two otherwise crucial questions. The reduction of rationality to maximization makes it possible for the neoclassical economist to ignore the psychological and cognitive aspects of decision making [Simon (1976)]: the rationality of a decision resides solely in the optimality of the decision itself, regardless the procedure which led to it. In the same way, the reduction of the coordination modes to the market, exemplified by the general equilibrium theory, brings the issue of organization outside the domain of economic theory: organizations such as firms do indeed exist, but their internal structure is immaterial for the allocation of resources at the level of the entire economic system.

The pioneering work of Coase (1937) pointed out the role of firms as a means of coordination of the actions of individual economic agents which does not use – or uses only to a limited extent – prices to direct the actions of individuals. In his view the function of firms parallels and complements that of markets as coordinating institutions, rather than that of consumers as utility maximizing individuals. But an in-depth inquiry into the distinctness of the coordination mechanisms implemented by different organizational set-ups did not begin until the 70's.

It is in this period that the neoclassical theory discovers the "organization", that is a mechanism of coordination based mainly on rules, partly explicit and codified, partly only implicit and tacit. The neoclassical theory extents in this way its scope by adding organizations to markets and individuals and broadens as well the scope for rational choice of individuals under consideration to include the choice of the organizational form which is more efficient at coordinating their interactions.

A relevant part of such recent developments (cf. especially the literature based on the agency theory) has operated a reduction of the essence of the firm to a contractual relationship which links the employer to the employees in the form of hierarchical authority. Individuals who form the organization cannot foresee all contingencies which might arise and affect their operations. It is therefore impossible to design complete contracts which exactly specify ex-ante the detailed performance which will be required in every circumstance [cf. Hart (1988)].

A way out of this impossibility is for the employees to surrender part of their discretion to superordinates, who are given the right to decide, within the contractual limits, the tasks which the worker must execute. The authority relation is therefore assumed to be the result of rational and utility maximizing behaviour on the side of both the employer and the employees.

As Coase has pointed out, the widespread reduction of the nature of the firm to a bundle of contractual relations "has led to or encouraged an undue emphasis on the role of the firm as a purchaser of the factors of production and on the choice of contractual arrangements which it makes with them" [cf. Coase (1988), p. 37]. The firm looses in this way its features of an institution which coordinates decisions and does not have a real organizational structure, that is a set of procedures and mechanisms which are designed to make this coordination possible.

Two streams of research have to some extent challenged the reduction of the essence of the firm to a bundle of bilateral relations: transaction cost economics, originating in the work of Williamson, and a series of studies which could be grouped under the heading of "organizational architecture", especially developed by the recent work of Aoki, Crémer, Sah and Stiglitz.

The former assesses hierarchical organizations in terms of their efficiency at economizing on the costs which might emerge in the processes of allocation of resources; the latter compares the efficiency of different informational structures – basically a vertical and centralized structure versus a horizontal and decentralized one – in terms of their abilities to compensate for mistakes and the occurrence of uncertain events which may affect their normal course of operations.

Aoki, in particular, has provided an extensive comparative analysis on the role of different coordination modes in determining the organizational capability to adapt to various environmental conditions [cf. especially Aoki (1986) and (1988)]. By relaxing the strongest neoclassical assumption on perfect and symmetric information and perfect market surrogation of human errors, he shows that the organizational structure does matter, because the way in which information is distributed within an organization affects the efficiency of its utilization.

But these approaches still focus on the role of firms as allocators of resources and processors of information: learning and adaptation are limited to those activities which can be reduced to Bayesian updating of probabilities, within given and constant information processing capabilities. As a consequence, they cannot account for a more fundamental kind of learning, which involves problem-solving, generation of new skills and routines, building new representations of the environment.

An alternative view [cf., for instance, Nelson and Winter (1982), Dosi, Teece and Winter (1991), Dosi and Egidi (1991)], which is being developed within the paradigm of evolutionary economics and largely draws from behavioural and organizational sciences, stresses instead the role of firms as repositories of knowledge. According to this view, firms are social institutions which develop a body of productive knowledge [cf. Winter (1982)], largely tacit and not codified, with the purpose of doing useful things. Such a body of knowledge is largely specific to the organization and is embodied in the set of routines and skills which characterize it.

Thus, organizational learning – i.e. the accumulation of knowledge – cannot be reduced to Bayesian probability updating, but requires agents to build new representations of the environment and develop new skills and routines which were not known to them before.

Organizational learning is itself highly firm-specific and, although fed by the learning processes of its individual members, cannot be reduced to their sum. In particular, the relations among different parts of the organization, as defined by its structure, play a fundamental role in driving and shaping the collective learning process.

Within this framework, the problem of coordination goes far beyond the distribution of information flows, to include coordination of different and possibly conflicting representations of the environment and different and possibly diverging learning processes.

This paper tackles the analysis of this aspect of coordination, by means of a computational model of learning in which learning is described as a search in the space of representations ("models of the world"). This model will be described in the third section of the paper, and will allow to run a few simulations – presented in the fourth section – of the learning behaviour of different organizational structures in different environmental conditions. In the meantime, the next section defines what is meant in this paper by coordination problem.

2. The coordination problem

Let us begin by outlining a simple organizational coordination problem. Consider a firm which is facing an unknown environment. The latter can be in one out of n possible states:

$$S \equiv \{s_1, s_2, \ldots, s_n\}$$

This set of states of the world can represent, for example, the n possible product types which can be alternatively required by an exogenous demand.

The firm formulates a forecast on the product type which is being demanded and implements a production process. In the most general terms, a production process can be considered as a sequence of operations which embed individual and organizational skills and routines, to be coordinated in such a way as to yield the desired product type. Therefore a production process can be formally represented by a string over an alphabet A which encodes all such basic operations:

$$a_1 a_2 \ldots a_k \quad \text{with} \quad a_i \in A$$

The set A_k^* of all possible strings of length k over this alphabet represents the set of all possible production processes of length k which could be virtually implemented starting from the set A of basic operations.

Production processes map into product types:

$$F: A^* \to P$$

This very general mapping can include also production processes which are technically meaningless (although technically viable) and produce product types which are never demanded. Of course it is also possible that some of the product types included in S cannot be produced with the present set of skills and therefore are not in the range of F. In general we can also have many different production processes which yield the same product type.

The payoff function for the organization is simply:

$$\pi: P X S \to R$$

the organizational payoff when the state of the world s_j occurs and the product type p_i is produced will be indicated by π_{ij}.

Suppose now that the production process is divided into h segments [1], each of them implemented by a different decision unit, or "shop". Suppose also that per-

[1] Obviously learning processes take place also through changes in the division of labour. However this paper does not face this problem, and considers the latter as given and fixed.

fect planning is not feasible, meaning that shops always have some degree of discretion in their decision-making. Discretionary behaviour may be limited to the necessity to interpret and implement the central plan, which, because of the cognitive and computational limits of the central office, cannot make provisions for all the contingencies which can arise at the operational level. But discretionary behaviour may also involve the possibility for the shops to contribute to the definition of the production plan itself, by making use of the knowledge they can autonomously accumulate.

The organizational decision problem has therefore two aspects: first of all it is necessary for the organization to make a forecast of the behaviour of the demand, and then to coordinate the production processes of the various shops in such a way as to obtain the desired product type. The former is a "cognitive" problem, which involves the organizational capabilities to "understand" the characteristics of the environment, the latter is instead a "coordination" problem, which requires the various decision units, in which the organization is divided, to coordinate their actions.

The two aspects interact in a complex way: coordination is generally made easier when the various parts of the organization share a common body of knowledge [cf. also Crémer (1991)], but on the other hand the overall scope of organizational learning is broader when the organization can rely on a diversified knowledge basis.[2] This trade-off will be analyzed in the fourth part of the paper, by means of the model of learning which will be introduced now.

3. A computational model of learning

This section briefly outlines a computational model of learning, in which the model of the world detained by the decision maker (the *representation* level) and the decision rules he follows (the *action* level) co-evolve in a process of adaptation to the environment which is being faced.

The idea that optimum decision rules cannot be defined in absolute terms, but only relatively to the decision maker's cognitive capabilities has received increasing attention in the economic literature [cf., for instance, Heiner (1983)]. This section presents a simulation model which focuses on the links between the decision maker's cognitive capabilities and the actions he can effectively select.

Consider a standard problem of individual decision making and suppose that it is faced repeatedly by the same agent. The decision maker, by using his previous experience, makes a forecast of the state of the world which he thinks will occur next and chooses an action which he considers as appropriate. At the outset the player has no knowledge either of the payoff matrix or of the "laws" which determine the changes in the environment. The decision process consists therefore

[2] The so-called "loss of control" literature analyzes the case when members of the organization do not share the same model of the world, but considers only ex-ante information flows, represented by the top-down transmission of plans. In this case diversity of representations generates a chain of misinterpretation of the plans and determines a loss of efficiency. The fact that diversity of representations could also have the positive effect of improving on the process of plan revision, by making use of the ex-post information coming from subordinates, is not taken into consideration, since it would require a modification of the superordinates' model of the world, which cannot be carried out within a Bayesian framework.

of two elements: the state of knowledge about the environment, represented by the agent's forecasting capabilities, and the rules for choosing an action, given this forecast.

The basic component of our learning system [3] is a condition-action rule, where the execution of a certain action is conditional upon the agent's perception of a certain state of the world. The condition part is a string of symbols which encodes a subset of the states of nature and is activated when the last detected state of the world belongs to this subset.[4] Formally, the condition can be represented as a string of n symbols (as many as the states of the world) over the alphabet $\{0, 1\}$ and it is satisfied when and only when the last observed state of the world corresponds to a position where a "1" appears. Thus, the condition:

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

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

$$c_k = 1$$

All in all, a given set of such conditions defines a set of criteria or categories according to which the environment is classified by the decision maker. It is important to stress that each condition defines one subjective state of the world, as perceived by the agent and defines its relationship with the objective 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:

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

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

$$\Theta_1\text{:}\ \ 110$$
$$\Theta_2\text{:}\ \ 101$$

The agent conceives two "subjective" states of the world, Θ_1 and Θ_2. The agent thinks he is in the former when the real state of the world is either s_1 or s_2, whereas he 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 knows are the two Θ's.

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

$$a_1 a_2 \ldots a_p \quad \text{with} \quad a_i \in \{0, 1\}$$

[3] The learning model employed here is an adaptation of the classifiers system methodology, a highly general learning system which processes a set of condition-action rules in order to achieve high adaptation of its behaviour to complex and largely unknown environmental conditions. A presentation of the original model and its main applications can be found in the works by John Holland (see especially Holland (1975) and (1986)), a discussion of some possible applications to economics can be found in Arthur (1991).

[4] The system could be given a longer "memory" simply by introducing more conditions which depend on previous states of nature, but this would be a redundant complication.

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 therefore be represented by a set of such condition-action rules:

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

where:

$$R_i\colon c_1 c_2 \dots c_n \rightarrow a_1 a_2 \dots a_p \qquad \text{with} \quad c_i, a_i \in \{0, 1\}$$

In addition each rule is assigned a "strength" and a "specificity" (or "generality") measure. The *strength* measures the past usefulness of the rule, that is – approximately [5] – the net payoffs, which have been cumulated every time the rule has been applied. 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 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 absolutely ignorant about the characteristics of the environment he is going to face: all the rules initially available to him have the highest generality, that is 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:

1. *Condition matching:* a message is received from the environment which informs the system about the last state of the world. This message is compared with the condition of all the rules and the rules which are matched, i.e. those which apply to that particular state of the world, enter the following step.
2. *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 this 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):

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

   Where k_1, k_2 and k_3 are constant coefficients.

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

[5] A precise expression for computing the strength will be given later on.

3. *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 (which in turn depends on which state of the world actually occurs). If the j-th rule is the winner of the competition, we have:

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

4. *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 search is always made available.

The search for new rules is driven by the system's past history: genetic operators generate new rules which explore other possibilities in the proximity (in a sense which will be precisely defined) of the presently most successful ones, in order to discover the elements which determine their success and exploit them more thoroughly. These newly generated rules 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:

1. *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.
2. *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 attitudes which tend to drive the learning system towards, respectively, specific rules which classify the environment into finer categories and more robust rules which instead cover a wider set of states of the world. Different degrees of specification and generalization 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 remaining part of this 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 operator's routine is called.

The next section of the paper will apply this computational model of learning and decision making to an example of the organizational decision problem outlined in the previous section.

4. Organizational structure, coordination and learning: some simulation results

In this section we consider a very simple example of the kind of organizational decision problems outlined in the second part of this paper and run a few simulations by means of the framework presented in the previous section.

The outcome for the organization depends on the actions of several agents in a non-additive way. Agents are supposed to share the same collective objective and therefore try in good faith to achieve the organizational goal, but their information processing capabilities are limited and evolve along the lines expounded in the previous part of the paper. In this way we can assume away any incentive consideration and concentrate on the role of shared knowledge and communication among the members of the organization.

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 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 a payoff of -2).

We will consider two different types of organizational structures: one in which knowledge about the environment is centralized and the other in which it is partly decentralized. More precisely, we have the following two models:

Centralized structure is formed by three decisionmaking units: the management and two shops. The former observes and interprets the environmental signal, trying to forecast the demanded product type, and sends a message to the shops. The latter observe and interpret independently the message they receive from the management and implement one of the possible actions (segments of the organizational production process).

All three agents behave according to the computational model of learning outlined in the previous section. In particular, each of them is represented by a set of rules, whose conditions classify environmental messages (in the case of the management) and managerial messages (in the case of the shops) and whose actions are, respectively, messages sent to the shops and segments of the production process.

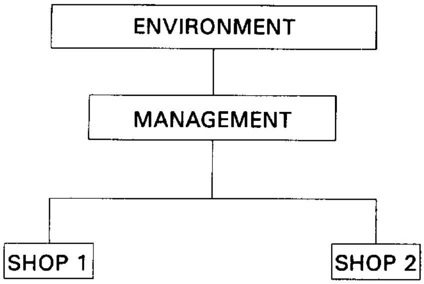

1

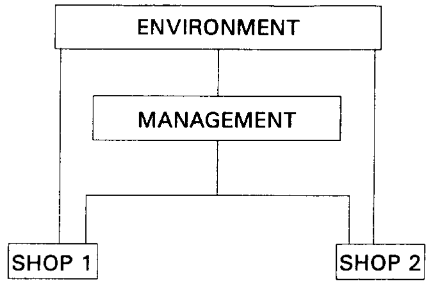

2

Fig. 1. Centralized structure – information flows

Fig. 2. Decentralized structure – information flows

Decentralized structure is also formed by a management and two shops. The former observes and interprets the environmental signal and sends a message to the shops, exactly as in the centralized structure. But the shops, in addition to observing and interpreting the managerial message, base their decision also on the independent observation of the environmental signal.

All three agents are modelled as in the centralized structure, with the addition of a condition which classifies environmental messages for each of the rules representing the two shops.

Thus, in the centralized structure the organizational knowledge of the environment is entirely detained by the management and the two shops do not form any autonomous knowledge of the firm's environment. In the decentralized structure it is instead distributed among management and shops.

The adaptive performance of the two organizational structures has been simulated in different environmental conditions, characterized by varying degrees of uncertainty and non-stationarity.

It is important to stress that all the following simulations assume as initial condition a situation of complete lack of *a priori* knowledge for all the agents. At the outset agents are in fact modelled by a set of rules whose conditions are entirely formed by 1's and whose actions are selected randomly.

Simulation experiment 1 considers the case of a stationary environment, in which the demanded product type is held constant. Coordination on the optimum production process is quickly achieved by both structures. However some differences do indeed emerge.

When the agents are not seeking to make their knowledge more specific (i.e. their specificity coefficient equals 0) the performance of both structure is identical. This is an obvious consequence of the fact that in this case no knowledge of the environment is actually formed, but decisions are selected by random trials until the correct one is chosen. The way in which knowledge is distributed among the parts of the organization is therefore completely irrelevant. It must be also pointed out that, in the case of a stationary environment, random trials, until the best action is chosen, prove to be a very efficient behaviour: learning is in fact a wasteful process in this case in which there is actually nothing to be learned.

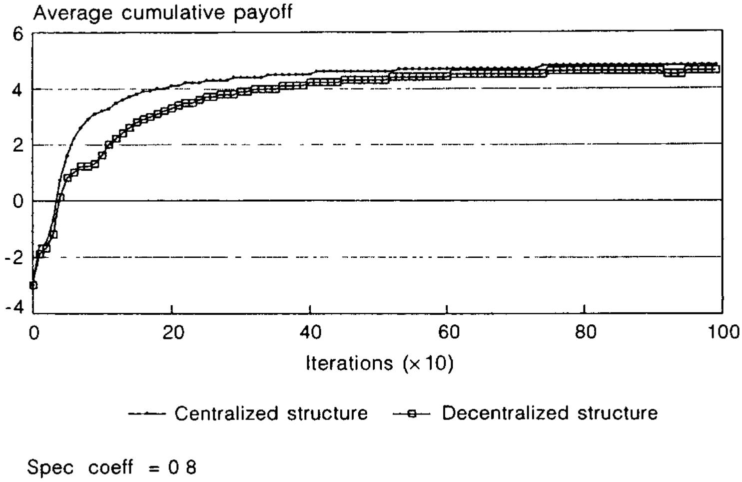

Fig. 3. Stationary environment – constant product type

When instead agents do try to improve their knowledge of the environment, by building decision rules more and more specific (i.e. they are characterized by high specificity coefficients) the distribution of knowledge among the parts of the organization becomes relevant. Figure 3 shows the cumulated payoff (divided by the number of iterations) of the two structures when all agents are characterized by specificity coefficient 0.8. It emerges that the decentralized structure is slower in achieving coordination on the optimum actions. As already pointed out, the stationarity of the environment makes the accumulation of knowledge a useless and wasteful activity, therefore the decentralized structure, which is characterized by a higher degree of overall accumulation of knowledge about the environment, is penalised.

A differentiation between the behaviour of the two types of organization emerges more clearly when we consider changing environments. This is shown by the following two simulations.

Simulation experiment 2 considers an environment which is always changing, but according to a regular and predictable pattern. Suppose that the demanded product type switches from "3" to "4" and vice versa at every iteration.[6] In spite of a high specificity coefficient (0.8), the centralized structure cannot exploit this regularity and settles into constantly producing either types, with an average payoff of 2. On the contrary the decentralized structure is able to discover and exploit the environmental regularity and – after an initial period of learning – attains the highest possible payoff 5.

[6] Since agents are modelled by decision rules which base their forecast only on the last observed state of the world, a pattern of environmental change has to be characterized by a "memory" of only one period in order to be predictable.

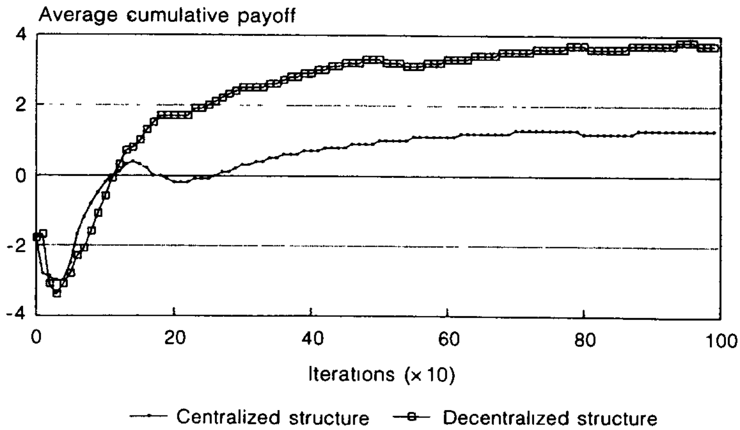

Fig. 4. Regularly changing environment – Every iteration s.o.w. changes 3–4

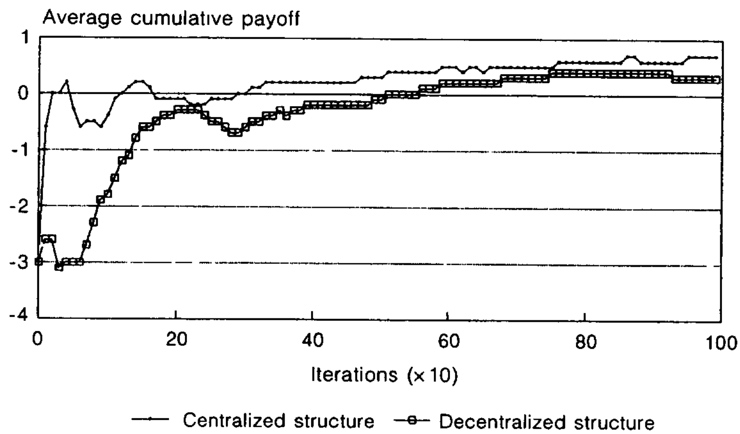

Fig. 5. Randomly changing environment – Every iteration random select. 3–5

Simulation experiment 3 considers instead continuous and 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 possibilities ("3", "4" and "5") at every iteration. Environmental changes are therefore confined to a subset of the possible states of the world, but are unpredictable inside such a subset: the optimum action is therefore producing always type "4".

Simulations (summarized by Fig. 5) show that, with specificity coefficient 0.8, the centralized structure is able to discover the "robust" solution and cling to it, whereas the decentralized one always explores new possibilities and abandons the optimum combination representation-action.

By comparing the results of the previous two simulations some interesting conclusions can be drawn. In order to exploit a regularly changing environment, a high amount of knowledge about the environment itself is required: the organizational model of the world must be able to distinguish between the states of the world and connect them diachronically. It is not surprising therefore that the decentralized structure is more appropriate in such circumstances: it is the structure 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, exemplified by the last experiment, requires strong coordination in order to make the entire organization implement coherently such a set of robust routines. By decentralizing the accumulation of knowledge and allowing therefore autonomous experimentation, this coherence is weakened. Centralization of the accumulation of knowledge, on the contrary, emphasizes the coordination around a unique central body of knowledge.

We could think of a decentralized structure as one which allows shops to explore independently new possibilities and feed them back to the entire organization through the centralized coordination system. On the other hand, a centralized structure emphasizes inter-shop coordinational around a centralized organizational representation of the environment.

5. Conclusions

The approach developed in this paper considers firms as learning organizations. When learning involves problem-solving and representation-building and cannot be reduced to mere probability updating, the role of organizations goes beyond coordination of information flows to encompass coordination of individual learning processes.

Organizational learning becomes an emergent property of the interactions among the members of the organization and, although it is the outcome of the individual learning processes, it cannot be entirely reduced to them. Such an outcome is in fact strictly dependent upon the organizational structure, that is upon the set of rules and relations which determine how knowledge is distributed within the organization.

This paper has suggested a model of learning which involves the construction and continuous revision of representations of the world and applied it to a collective decision problem.

It has emerged from simulations of this model that when agents do not necessarily share the same model of the world, and they are always updating their own model, in order to improve their action upon the environment, there exists a trade-off between the need to have homogeneous representations and the need to promote diversity. Homogeneity of representations favours the coordination of individual actions, diversity increases the scope for organizational learning.

A centralized structure attains higher homogeneity and seems more appropriate both for stationary and for irregularly changing environments, whereas a decentralized one allows higher degrees of diversity and seems more appropriate for regularly changing environments.

References

Aoki M (1986) Horizontal vs. vertical information structure of the firm. American Economic Review 76: 971–983

Aoki M (1988) Information, incentives and bargaining in the Japanese economy. Cambridge, Cambridge University Press

Arrow KJ (1974) The limits of organization. New York, London, W. W. Norton

Arthur WB (1991) Designing economic agents that act like human agents: a behavioral approach to bounded rationality. American Economic Review, Papers and Proceedings 81: 353–364

Calvo G, Wellisz S (1978) Supervision, loss of control, and the optimal size of the firm. Journal of Political Economy 87: 943–952

Coase RM (1937) The nature of the firm. Economica 4: 386–405

Coase RM (1988) Lectures on "The Nature of the Firm: Influence". Journal of Law, Economics and Organization 4: 33–47

Cohen MD (1991) Individual learning and organizational routine: emerging connections. Organization Science 2: 135–139

Crémer J (1980) A partial theory of optimal organization of a bureaucracy. Bell Journal of Economics 11: 683–693

Crémer J (1990) Common knowledge and the co-ordination of economic activities, In: Aoki M, Gustafsson B, Williamson OE (eds.), The firm as a nexus of treaties, London, Sage Publications

Cyert RM, March JG (1963) A behavioral theory of the firm. Englewood Cliff, NJ, Prentice Hall

Dosi G, Egidi M (1991) Substantive and procedural uncertainty. An exploration of economic behaviours in complex and changing environments. Journal of Evolutionary Economics 1: 145–168

Dosi G, Teece DJ, Winter SG (1991) Toward a theory of corporate coherence. In: Dosi G, Gianetti R, Toninelli PA (eds.), Technology and the enterprise in a historical perspective. Oxford, Oxford University Press

Hart O (1988) Incomplete contracts and the theory of the firm. Journal of Law, Economics and Organization 4: 119–139

Heiner RA (1983) The origin of predictable behavior. American Economic Review 73: 560–595

Holland JH (1975) Adaptation in natural and artificial systems. Ann Arbor, University of Michigan Press

Holland JH (1986) Escaping brittleness: the possibilities of general rule-based systems. In: Michalski RS, Carbonell JH, Mitchell TM (eds.) Machine learning II, Los Altos, CA, Morgan Kaufmann, pp. 593–623

Holland JH, Holyoak KJ, Nisbett RE, Thagard PR (1986) Induction: processes of inference, learning and discovery. Cambridge MA, MIT Press

Nelson RR, Winter SG (1982) An evolutionary theory of economic change. Cambridge, MA, Harvard University Press

Sah RK, Stiglitz JE (1986) The architecture of economic systems: hierarchies and polyarchies. American Economic Review 76: 716–727

Simon HA (1976) From substantive to procedural rationality. In: Latsis SJ (ed.), Method and appraisal in economics. Cambridge University Press, pp. 129–148

Simon HA (1979) Rational decision making in business organizations. American Economic Review 69: 493–513

Williamson OE (1975) Markets and hierarchies: analysis and antitrust implications. New York, The Free Press

Williamson OE (1985) The economic institutions of capitalism: firms, markets, relational contracting. London, MacMillan

Winter SG (1982) An Essay on the theory of production. In: Hymans H (ed.), Economics and the world around it. Ann Arbor, University of Michigan Press, pp. 55–93