ISIGrowth

INNOVATION-FUELLED, SUSTAINABLE, INCLUSIVE GROWTH

Working Paper

The formalization of organizational capabilities and learning: results and challenges

Giovanni Dosi
Scuola Superiore Sant'Anna

Marco Faillo
Department of Economics and Management, University of Trento

Virginia Cecchini Manara
Scuola Superiore Sant'Anna

Luigi Marengo
Department of Business and Management, LUISS

Daniele Moschella
Scuola Superiore Sant'Anna

3B/2017 February

This project has received funding from the European Union Horizon 2020 Research and Innovation action under grant agreement No 649186

INSTITUTE
OF ECONOMICS

Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

The formalization of organizational capabilities and learning: results and challenges

Giovanni Dosi °
Marco Faillo *
Virginia Cecchini Manara °
Luigi Marengo §
Daniele Moschella °

° Istituto di Economia, Scuola Superiore Sant'Anna, Pisa, Italy
* Department of Economics and Management, University of Trento, Trento, Italy
§ Department of Business and Management, LUISS University, Rome, Italy

2017/08 February 2017
ISSN(ONLINE) 2284-0400

The formalization of organizational capabilities and learning: results and challenges[*]

Giovanni Dosi[1], Marco Faillo [2], Virginia Cecchini Manara[1], Luigi Marengo[3] and Daniele Moschella[1]

[1] Institute of Economics, Sant'Anna School of Advanced Studies, Pisa, Italy

[2] Department of Economics and Management, University of Trento, Trento, Italy

[3] Department of Business and Management, LUISS University, Rome, Italy

Abstract

This work offers an overview of recent formalizations of organizational capabilities and learning. We first present the main characteristics both of NK models and of the approach based on Classifier Systems, focusing on their early applications to organization studies. We then discuss how the use of these models has contributed, in the recent years, to the formal analysis of the development and change of firm's dynamic capabilities by improving our understanding of processes of organizational learning and adaptation, and of the relationship between cognitive and governance issues.

Keywords: Capabilities, Decomposability, Organizational structure, Problem-solving

JEL classification: D21, D23, D83, L23

1. Introduction

One way to study organizations is to conceive them as "behavioral entities", largely characterized by routinized patterns of action and by capabilities stemming from ensembles of them. This is the

[*] This article has been prepared for *The Oxford Handbook of Dynamic Capabilities*, edited by David J. Teece and Sohvi Leih (Oxford University Press). Daniele Moschella received financial support by the Italian Ministry of Education and Research under the SIR Programme (Project code RBSI14JAFW). This paper is produced as part of ISIGrowth project on Innovation-fuelled, Sustainable, Inclusive Growth that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 649186 ISIGrowth. The work draws upon other works of the authors, in particular Dosi et al. 2011, which the reader is referred to for further details.

approach shared by evolutionary, capability, and behavioral theories of the firm, who place the primitives of the essence of organizations in their problem-solving features, which are in turn grounded on imperfect and boundedly rational processes of learning and search, as well as on mechanisms of distribution of cognitive labor (March and Simon, 1958; Cyert and March, 1963; Nelson and Winter, 1982; Dosi *et al.*, 2000, 2008; Dosi and Nelson, 2010).

In this work we review some recent contributions which provide a formalization of those aspects. In the models we present, firms are seen as information-processing and problem-solving entities which have to interpret signals coming from the environment and perform sequences of activities. Such activities stand for physical or cognitive acts eventually leading to the solution of a "problem", being it e.g. the production of a car or the identification of a malaria-curing molecule.

These intuitions can be captured by placing organizations over a *fitness landscape*, as in NK models, on which most of the research we discuss here is based, and which are particularly suitable, as we will see, for the description of complex processes of search for optimal configurations of organizational policies, strategies or decisions.

We also introduce a small class of models based on the "Classifier Systems" approach (Holland, 1975), in which firms are characterized by "if... then…" rules which evolve in response to external feedbacks. We will show how this class of models allows for a richer representation of the link between cognition, learning and environmental feedback.

This strand of research is of great importance for the study of dynamic capabilities. It provides an effective description of the complex processes that lead to the development and change of existing capabilities, and moreover it addresses the relationship between these processes, studying for example the interplay between the division of cognitive labor and the design of the governance structure of the firm, or analyzing the link between diverging interests, managerial discretion, and organizational performance.

The work is organized as follows. In ***Section 2***, we introduce the general characteristics of NK models and Classifier Systems and some of their earlier applications, which explore the conditions under which organizations adapt to changing environments, the possible trade-offs between search and stability associated to different forms of hierarchical governance, the role of organizational architecture in shaping learning processes and determining organizational performance. ***Section 3*** presents the most recent applications of these models. These address how the degree of decentralization and the complexity of the problems at hand impact on the capacity to innovate and learn and on the balance between exploration and exploitation within organizations, under different environmental conditions (Section 3.1); how the cognitive representation of the environment affects organizational performance, the evolution of routines and the possibility to decompose problems (Section 3.2); how the degree of centralization of decision power in presence of conflict between principal's and agents' interests influences organizational learning and performance (Section 3.3). ***Section 4*** concludes.

2. Formal models of organizational capabilities and learning

2.1 The NK Model

The first family of models of organizations that we consider has been inspired by Kauffman's (1993) "NK model", introduced to study the evolution of populations of biological entities made up of many elements (e.g. genes in a genotype)[1]. Its basic features are captured by two variables: *N*, that refers to the number of parts of a system, and *K*, that reflects how complex the interdependencies within the system are. A system is conceived as a string of *N* elements ($i = 1,..., N$); for each element *i*, there are *A* possible states. In the simplest case *A* is made of two states (in many applications $A=\{0,1\}$). The set of all possible configurations of the system's elements is given by A^N. The fitness value of the system depends on the contribution of its elements, which is drawn from a uniform distribution between *0* and *1*. In the simplest case, when $K=0$, the contribution of each element depends only on its state. When *K* is greater than zero, the fitness value of each element depends not only on its state, but also the state of other *K* elements. Thus, the number of fitness values that a single element can assume is equal to A^{K+1}. The system's fitness is usually computed as the mean of the fitness values of its elements.

The distribution of fitness values assigned to all the possible configurations constitutes the *fitness landscape* (Wright, 1932). Evolution is represented as a process of exploration of the fitness landscape in search of the configuration with the highest fitness value, moving from one configuration (a point in the fitness landscape) to another by changing the value of an element; the search ends when a configuration is reached that has no neighbors with higher fitness.

The value of *K* determines the smoothness of the fitness landscape. When $K=0$ the landscape has a single global optimum that can be reached by every configuration by moving to the one-mutant neighbors (the neighbor with the same configuration except the state of one single element). In this case the fitness of all the one-mutant neighbor configurations of a specific configuration is almost the same (the landscape is smooth). When $K=N-1$ (maximum complexity) the landscape is characterized by a large number of local optima, but only a small fraction of the them can be reached starting from

[1] See Frenken (2006) for a review of the applications of NK models.

any specific configuration. The fitness of one-mutant neighbor configuration in this case is entirely uncorrelated with the starting point (the landscape is maximally rugged). In addition, as *K* increases, the size of the basin of attraction of each of the local optima tends to shrink. Thus, it could be that none of the configurations might be located in the basin of attraction of the global optimum.

2.2 Classifier Systems

In NK models some key cognitive and behavioural aspects of learning and adaptation processes are collapsed in the link between elements' status and fitness values. In this section we briefly present a small class of models based on a richer and less stylized representation of the relationships between cognition, action and environmental feedback. This models are inspired to Classifier Systems (Holland, 1975; Holland, 1986; Holland *et al.*, 1986), defined as systems of "if… then…" rules which evolve in response to environmental feedback.

We introduce, in particular, the model of "structural" learning developed by Marengo (1992) (see also Marengo (1996)) in which agents are adaptive learners who adjust their information processing capabilities through trial-and-error and in response to the feedback received both by the environment and by the other members of the organization. In this model, the use of condition-action rules implies that the execution of a certain action depends on the agent's perception how the state of the world he is facing fits one of the categories already defined in his mental model. The aim of the system is to select the most successful rules, and at the same time to discover new rules by recombining and mutating elements of the already existing rules.

With some simplification with respect to Marengo (1992), we say that the environment is described by a set of *n* possible states $S=\{s_1, s_2, ..., s_n\}$ whereas organizational behavior is characterized by a set of *k* possible actions $A=\{a_1, a_2, ..., a_k\}$, with $a_i \in \{0,1\}$. The payoff function for the organization is described as $\pi : SXA \rightarrow [0,1]$: the payoff thus depends both upon organizational acts and environmental states.

The basic component of this learning system is, as mentioned, a condition-action rule, *i.e* an "if… then…" rule that maps detected environmental profiles into actions. Each rule takes the form: $c_1 c_2 ... c_n \rightarrow a_1 a_2 ... a_n$, with $c_i \in \{0,1\}$; a rule with a condition $c_k = 1$ is activated whenever the state of the world is s_k . The condition part can have a *1* in more than one position of the string, meaning that the agent classifies different objective states of the world with the same subjective category; the action part is instead a string with one and only one position equal to *1*, say $a_k = 1$, meaning that the action a_k is chosen.

Notice that two (or more) different rules may apply to the same state of the world: in this case, one of them is selected by means of a competition in which each rule makes a metaphorical "bid" based on its *strength* – i.e. the past effectiveness of the rule represented by its cumulated payoff – and on its *specificity* – i.e. the size of the of condition to which it applies. At the initialization stage, for each rule, all the conditions are formed entirely by *1*'s (all the rules have the highest generality). This captures the idea that at the beginning the decision maker is completely ignorant about the characteristics of the environment. At each round of the simulation, the strength of the rule that is executed is reduced by the amount of the bid and increased by the payoff that the action receives.

The system evolves both by selecting the most effective among the existing rules and by generating new rules. Typically, new rules are generated by mutating and recombining the elements of the most successful ones. Hence, search is not completely random but influenced by the past history of the system.[2]

In Marengo (1992) the genetic operator used for the action part simply consists in a mutation in the "vicinity": the action prescribed by the newly generated rule is (randomly) chosen in the close proximity of the one prescribed by the parent rule. For the condition part two genetic operators are used: the first is a *specification* operator, which increases the specificity of the parent rule (by mutating *1*'s into *0*'s with some probability); the second is a *generalization* operator, which decreases the specificity of the parent rule (by mutating *0*'s into *1*'s with some probability).

2.2 Earlier Applications to Organizational Dynamics

Levinthal (1997) can be considered as one of the earliest attempts to adapt the NK model to the study of organizations. In Levinthal's model, the *N* elements capture some organizational attributes, like a policy or a strategy choice, that can take two alternative states (the policy is carried out or not). At the population level, the evolution of organizations is driven by a selection mechanism based on the replacement of the least fit organizations with new ones, whose configuration is either copied from existing forms or randomly generated, depending on the overall fitness of the population.

In Levinthal's model two alternative patterns of organizational change are analyzed: local search (organizations modify *one* attribute at a time, to achieve a higher level of fitness in their immediate neighborhood) and long jumps (organizations change *all* their attributes, drawing a new organizational form at random).

Levinthal observes that with local adaptation the diversity of organization forms rapidly declines,

[2] In incumbent models, new rules take the place of the currently weakest ones, so that the total number of rules is kept constant.

since bad performers are selected out and replaced by copies of good performers. Good performers reach the closest fitness picks through local mutation, but the final outcome, in terms of overall fitness value, depends on the value of *K*. With *K=0* and local search, all the organizations quickly walk to the only global optimum, and selection wipes out the initial heterogeneity of the population and yield convergence to unique optimal organizational form. With higher values of *K*, the number of local optima on which subsets of organizations converge increases, and the adaptation process assume a path dependent pattern. In this case heterogeneity is reduced, but does not disappear.

Within this model, persistence of heterogeneity in organizational forms can be explained as a consequence of the degree of interrelatedness of organizational attributes, the characteristics of the search process and the nature of the selection mechanism. Levinthal (1997) shows that heterogeneity can even emerge out of homogeneity: in the case of local search, if *K>0*, random mutations will take organizations in the basins of attraction of different local optima and selection and adaptation will only partially reduce diversity. The opposite outcome is obtained when organizations can perform long jumps, by mutating many features at once. Even with large *K*, – assuming that *N* is not so large – heterogeneity disappears since organizations which are trapped in local optima can reach – even if with low probability – the global optimum through a sequence of radical mutations.

The model can also account for the consequences of environmental change. The change in the environment is modeled as a redrawing of the fitness contribution of some of the organizational attributes once the population has reached an equilibrium. The result of such changes depends on the value of *K* and on the number of attributes whose fitness contribution has been changed. When *K=0* and only one attribute is affected by the change, the impact on the population, both in terms of overall fitness and composition, is very weak. With large *K*, even the modification of the fitness contribution of just one attribute can cause a large alteration of the shape of the landscape with a significant change in the distribution of local optima. This induces more radical changes and a high rate of mortality of incumbent organizations. If the change in fitness values affects many attributes, when *K=0*, a new global optimum could emerge that is far away from the previous one, and the new randomly generated and better performing organizations will replace the incumbent ones. With large *K*, the population tends to remain distributed over a large number of local optima but with some probability a subset of the population might well find itself not too far from the high fitness portion of the new landscape. Thus, diversity helps the population to adapt to dramatic environmental changes.

A different representation of the organizational structure using the NK model was introduced by Rivkin and Siggelkow (2002) and then applied in several other works (see, for example, Rivkin and

Siggelkow, 2003; Siggelkow and Rivkin, 2006; and Siggelkow, 2011 for a review) which have studied the effects of interdependencies within an organization on how decisions are made, with a particular emphasis on the trade-off between the broad search for good combinations of decisions and the need to stabilize on some decisions, once a good combination has been reached. These models also allow the study of the relationship between organizational performance and different forms of hierarchical governance, determined by the division of decisions between departments, the amount of information that is submitted to top management, the degree of discretion of the CEO, the managers' incentives to look beyond their field of action, the cognitive skills of managers.

As for the applications based on Classifier Systems, Marengo (1992) develops a model to study the role of organizational structure in shaping the organizational learning process. In particular, the paper investigates how organizational learning emerges from the coordination of individual learning processes, and how the performance of different coordination devices depends on environmental conditions.

Consider a firm that can produce a certain number of product types demanded by an exogenous market; the production process is divided into several parts, carried out by different "shops". The problem is therefore to correctly detect which product type is being demanded (the "state of the world") and to coordinate the actions of the shops so that the correct production process is implemented. Suppose now that the firm is composed by two shops and a manager. All three agents behave according to the classifier system outlined above. In a *centralized structure* 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. In a *decentralized structure,* 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, instead, knowledge of the environment is distributed among management and shops.

In general, the model shows how the relative performance and the learning patterns of centralized *vs.* decentralized structures depend on the characteristics of the environment.

First, in stationary environments (when the demanded product is held constant) agents can achieve coordination without building any model of the environment (*i.e.* with no specificity in their condition-action rule): learning is in fact a wasteful process, since there is actually nothing to be learned. In this case, if agents try to learn by building more and more specific decision rules, the decentralized structure is slower in achieving coordination on the optimal actions.

Second, if the environment undergoes predictable changes (for example of a cyclical type), the centralized structure cannot exploit this regularity, even with very specific condition-action rules. On the contrary, the decentralized structure is able to discover and exploit the environmental regularity and attains the highest possible payoff.

Third, in the case of frequent and unpredictable environmental changes, the organization must develop stable routines that provide a “satisficing” average result in most conditions. Decentralized learning is detrimental, since individual efforts to grasp unpredictable environments limit the stability of these routines. In this case shops are better off by relying on the management’s message.

In a somewhat similar modelling vein, Pentland and Reuter (1994) formalize organizational routines as a set of functionally similar patterns represented via rule-based grammar models. So a routine is a “grammar” which defines all the action patterns which are, so to speak, “legal”, having different action patterns as possible instantiations triggered by different environmental or intra-organizational signals (the “if” part).

3. Recent applications

3.1 Organizational learning: exploration, exploitation and imitation

The use of formal models of organizational search like the ones we have reviewed so far have proven particularly fruitful in the study of the trade-off between exploration and exploitation (March 1991) as well as of the role of imitation in organizational change. The use of simulations can provide an effective operative definition of ambidexterity, broadly conceived as the organization’s ability to both exploit and explore (Tushman *et al.*, 1996; O’Reilly and Tushman, 2004; O’Reilly and Tushman, in press) that according to some is a key component of dynamic capabilities (O’Reilly and Tushman, 2008).

Here we review some very recent contributions aimed at providing a formal analysis of organizational adaptation to changing and complex environments. Some of these works focus on how some specific features of the organizational structure influence the effectiveness of exploration activities in complex and instable environment. Others pay more attention to the role of imitation in organizational learning and its long-run effects as the degree of complexity of the target practices increases. The great majority of these applications is based on the use of some versions of the basic NK model.

Fang *et al.* (2010) investigate the relationship between the degree of decentralization and the ability of the firm to identify innovative solutions to organizational problems. The key research question is whether partial decentralization in the form of subgroup structural semi-isolation benefits the organization's long-term performance, and if its benefits depend on factors like the complexity of the problems, the degree of environmental change or the rate of personnel turnover. The study is inspired by March's (2005) idea that having an organizational structure based on units or groups working in isolation increases the degree of diversification of solutions. The limited connections between groups would allow the selection and diffusion, at the organization level, of superior solutions that increase the organization's innovation capacity. As noted by the authors, similar dynamics have been described in the organizational literature on ambidexterity (Benner and Tushman, 2003; Bower and Christensen, 1995; O'Reilly and Tushman, 2004), and in studies based on NK models that show how some degree of decentralization could prevent the organization to be trapped in local optima (Ethiraj and Levinthal, 2004; Siggelkow and Levinthal, 2003, 2005; Siggelkow and Rivkin, 2005).

The model used for the simulation is a variation of the connected caveman model (Watts, 1999)[3] combined with a learning mechanism inspired by March (1991) and with the addition of parameters that identify the number of decisions and the interdependence between decisions in their contribution to the final payoff that are equivalent to the N and K parameters of NK models. By varying the degree of group isolation, the problem complexity, the environmental dynamics and the personnel turnover, the authors find that under most conditions the best performance is reached with a semi-isolated subgroup structure with moderate interaction between groups.

More recently, Bocanet and Postiglione (2012), provided an NK version of March's (1991) model of organizational learning where N represents individuals' beliefs. They replicate March's results and extend his model by allowing direct interpersonal learning, together with the possibility of learning from the organizational code, as in the original March's model, and the possibility for the code to learn from individuals. The dependent variables are the probabilities of learning, at the individual, organizational and code level, and the value of K. The main results are that the increase in complexity reduces organizational performance, independently of the degree and type of learning, and, most importantly, mutual learning has a positive effect on organizational performance.

Posen and Levinthal (2012) bring into question the common assumption of a causal link going from environmental change to the need for an increase in exploration and adaptation effort by the

3 In network analysis this correspond to an interaction structure represented by a graph with fully connected clusters of agents in which one agent for each cluster is connected to another agent in the adjacent cluster.

organization. They identify a set of conditions under which the appropriate response to unstable and rapidly changing environments would lead to focus on exploitation of existing knowledge instead of on exploration. In fact, if on the one hand environmental change may induce a reduction of the knowledge stock of the organization as the existing strategic practices are no longer effective (as stressed by the majority of scholars in strategic management), on the other hand recurring and unpredictable changes may also erode the benefits of additional knowledge acquired through exploration.

They use a model whose main feature is that the returns to an alternative can be evaluated only by choosing that alternative. At each time t the organization must choose between N alternatives. The choice of an alternative i results in an outcome $\sigma=1$ (positive) with probability p_i and $\sigma=0$ (negative) with probability $(1-p_i)$. The actual probabilities are not known to the organization and are replaced by beliefs q_i. Accuracy of the organizational beliefs system is measured by the sum of squared errors (differences between p_i and q_i). Turbulences in the payoff vector are introduced by allowing a random reset of the payoffs. The organization updates its beliefs on the basis of the outcome of each choice in every period. In addition, it is characterized by a search strategy that maps its beliefs to a specific choice from the set of the alternatives. Since expected payoffs are not known, the organization must use the information collected by choosing alternatives and updating its beliefs to choose the next alternative. Posen and Levinthal assume that the organization uses a softmax strategy (Luce, 1959) corresponding to a random choice from a Gibbs (Boltzmann) distribution. One key parameter of the latter is τ, measuring the degree of exploration. When τ is close to zero, then an alternative is selected with probability close to one even if the belief associated with it is only marginally bigger than the ones associated to the other alternative. When τ is close to one, differences in beliefs are ignored and the choice is random. With intermediate values of τ a good alternative has a higher probability of being sampled.

Simulations show that in a static environment March's (1991) conclusion about the need to find a balance between exploration and exploitation is confirmed. When turbulence is introduced the authors find support to March's conclusion about the positive correlation between organizational performance and exploration, but only with very low levels of turbulence. The result is the erosion of existing knowledge due to environmental change. But as turbulence increases, the environmental change erodes the returns to exploratory efforts at accumulating new knowledge.

Acquisition of knowledge through imitation and the effects of complexity in organizational decision making design on innovation and imitation strategy are studied by Ethiraj, Levinthal, and

Roy (2008). Their simulation is based on a classical NK model in which they compare organizational decision making structures characterized by the same number of interdependencies among decisions, but that differ in the distribution of interdependencies. In particular, they consider a non-modular structure, in which interdependencies between the *n* decisions are randomly distributed; a perfectly modular structure, in which interdependencies are limited to decisions belonging to same module; and a nearly-modular structure, in which the majority of links are within modules with a limited number of interdependencies between decisions belonging to different modules. The impact of the complexity of the decision structure on innovation capacity is assessed by allowing an intra-module incremental design change in which a new decision configuration is adopted if it improves the performance of the module.

The aim of the authors is to go beyond the short-term impact of complexity on firm's performance, by looking at the long-run effect of imitation on the firm's organizational design. This is done by allowing low performing firms to replace their modules and/or inter-module linkages with those of the most successful firm. As expected, the modular structure is superior in terms of innovation performance to the non-modular one. The performance of the near-modular structure is not different from that of the perfect modular structure. But once the innovation process ends, and firms have the possibility to observe and copy the decisional structure of the best performer (the leader), then the cost of the adoption of a perfect modular structure may become very high. In particular, when low performing firms can perfectly copy modules (decision configuration and intra-module linkages), but not inter-module linkages, imitators can easily catch up with the performance of the leading firms when the structure is perfectly modular, and can do even better with a nearly modular structure. In the case of non-modular structure the distance between imitators and leaders first increases and then remains stable until the end of the simulation. This implies that modular structures, and in particular perfectly modular ones, are less resistant to imitation than the non-modular ones, and their initial advantage in terms of innovation performance vanishes in the long-run.

When firms can imitate not only the modules but also inter-module interdependencies, then the distance between imitators and leaders, even if it is greater than the one observed with modular structures, decreases over time also in the case of non-modular structures. Interestingly, by considering a case very close to reality, like the one in which firms can perfectly copy modules but can produce only imperfect copies of the linkages between modules, the deterrence capacity of nearly-modular structures is significantly higher than that of the perfect modular ones. The authors conclude that nearly modular structure performs better than non-modular ones in terms of innovation capacity, and at the same time they have a better imitation deterrence as compared with perfectly modular structures.

In a related study, Csaszar and Siggelkow (2010) use a simulation based on an NK model to explore the relation between firm's performance and breadth of imitation (number of practices copied, $\beta \leq N$) as dependent on three factors: *i)* the degree of interaction between the practices of a firm (K); *ii)* firms similarity, measured as the number of practices shared by the imitating and the imitated firm (number of elements of the practice vectors with the same values, $H \leq N$); *iii)* and context similarity, measured in terms of correlation between the performances of firms that share the same vector of practices (number of shared contribution functions, $S \leq N$).

The authors distinguish between short run and long run effects and analyze both the case in which the imitating firm is aware of the number of practices that it shares with the imitated firm (discriminant imitation), and the case in which the imitating firm does not know what is shared (non-discriminant imitation). The main result of the simulation is that when context similarity is high, increasing the breadth of imitation is generally beneficial, while it is harmful when the firms live on very different landscapes. What is of particular interest is that imitation in this model has the double function of quickly mimicking high performing competitors and dislodging the imitator from local peeks. The authors show that when the level of complexity (K) is low, firms should copy (mimic) high performing firms, reducing the activity of independent search. Instead, when the level of complexity is higher, the low performing firm should copy small chunks of practices and use the new configuration to start a search aimed at escaping from local peaks. In the case of mimicking, even if by importing large chunks of practices from other firms the imitator escapes the current local peak, it does it by adopting a strategy that differs from exploration, and consists of a kind of exploitation of the practices of high performing firms.

3.2. Cognition and problem solving

3.2.1. The role of cognition and representations in organizational problem solving

Classifier systems give a prominent role to actors' cognitive map of the environment through condition-action rules. Gavetti and Levinthal (2000) add a similar perspective to the NK framework by introducing the distinction between forward-looking choices based on off-line evaluation of alternatives and backward-looking choices based on on-line local search.

In their NK model, organizations have a simplified and incomplete "cognitive model" of the environment, which is captured by a representation of the fitness landscape of lower dimensionality than the actual landscape ($N1<N$). For each point of the represented landscape there are 2^{N-N1} points in

the actual fitness landscape that are consistent with this point. The fitness value assigned to each point of the cognitive representation corresponds to the average fitness values of these 2^{N-N1} points. Thus, organizations use "cognitive templates" to explore not single points of the landscape but *regions*: they identify a pick in their perceived *N1*-dimensional landscape (by cognitive or off-line search) and then explore the remaining *N-N1* alternatives through a local (or on-line) search based on one bit-mutations.

Simulation results show that organizations which adopt a joint cognitive and experiential search dominate the population; this is especially evident under rugged landscapes, in which organizations which use purely experiential search are trapped into local optima. The cognitive search is important in identifying the superior, on average, basins of attractions, from which the local search can start, whereas the role of experiential search becomes more and more important as the accuracy of the cognitive representation (as measured by the dimensionality *N1*) decreases.

The paper also considers the effects of adaptation through changes in the cognitive representation. When organizations only use cognitive search, the effects of changes in the representation depend on the complexity of the landscape (the value of *K*). If *K* is high these changes may produce good performances, as they can compensate for a poor representation of the landscape. However, if one considers organizations which use joint off-line and on-line search, the shift to a new representation may also destroy the accumulated (on-line) experience.

In case the landscape itself changes, the loss of experiential wisdom can be more than compensated by the positive effects of changes in cognitive representation: indeed, the new representation may be able to identify also new and possibly superior basins of attraction. A typical example is when a firm enters an industry which is novel either for the managers or for the firm itself. Gavetti *et al.* (2005) model this situation by opposing the local, on-line search to the *analogical reasoning*, which is a form of cognitive search in which agents choose their initial configuration landscape based on some representation of the similarities between a novel industry (the "target" landscape) and a set of familiar industries. Simulation results show that analogical reasoning can enhance firm performance especially when the underlying decision problem is not easily decomposed.

Organizational structure also matters in determining the effects of a change in representations. Gavetti (2005) investigates in which way different hierarchical arrangements, characterized by different allocations of "cognitive rights", influence organizational performance. In his NK model, managers have to decide which representation is more adequate when the organization starts a new line of business. They do so by comparing the cognitive models they already have with the payoff generated by the local search on the new landscape. The organizational hierarchy exactly determines at which level (firm or divisional) and in which way this choice is made. Simulation results show that the way in which the initial "matching" between old representations and new landscape is done determines to a large extent organizational performance. Managers who hold a higher position in the organizational hierarchy find it more difficult to match the outcome of the local search with an appropriate cognitive representation. Moreover, representations that fit the old business tend to be preferred to representations that capture the new domain.[4]

More recent literature has shown that the cognitive dimension interacts also with other characteristics of the decision process, and that this interaction is important in determining organizational performance. Knudsen and Levinthal (2007) look at the capacity of evaluating alternatives. Using an NK model, they show that an imperfect evaluation of alternatives can improve organizational performance as it avoids a rapid identification of the local peak within the initial basin of attraction. On the other hand, Gary and Wood (2011) empirically show that, when the search for alternative policy configurations is driven by the (simplified) representations managers have of the landscape, organization performance crucially depends on the accuracy of the representations themselves.

Martignoni *et al.* (2016) argue that the relationship between the accuracy of representations and performance is mediated by three moderators: type of interdependence, type of misspecification, and degree of complexity. In their model, which is similar in spirit to the standard NK structure, the manager has both an *interdependence representation*, which captures her belief about the interdependences among the elements in the mental model, and a *performance representation,* that captures the mapping from actions to outcomes. They show that managers who overspecify their mental models ("complexifiers"), i.e. with a mental model that is more complex than the true performance landscape, tend to perform differently from managers who underspecify their models ("simplifiers"). In particular, when interdependencies are external, i.e. when the value of an organizational choice is affected by a variable which is not under the control of the manager, complexifiers outperform simplifiers. The intuition is that overspecification has a dislodging effect to the extent that it forces the manager into situations for which he or she has not formed an opinion. This effect vanishes when the manager has all activities under her control (internal interdependencies). In this case, underspecification is less detrimental.

[4] Gavetti (2005) reports the case of Polaroid's transition from instant to digital imaging. In particular, he observes how, during the early 90's, Polaroid's senior managers opposed a digital camera with no printing device as it clashed with the old razor/blade representation. See also, more in detail on the Polaroid case, Tripsas and Gavetti (2000).

Csaszar and Levinthal (2015), building on Gavetti and Levinthal (2000), present a NK model in which managers are allowed to search both over representations and over policy choices. The model delivers two important implications. *First,* the balance between the two types of search depends on the time horizon and on the complexity of the landscape. In particular, an increase in search over representations is optimal if the time horizon is sufficiently long and if the landscape is complex (high *K*). In this case, search over representations can help in dislodging from local peaks. *Second,* and in contrast to the result of Gary and Wood (2011), they show that in some conditions increased accuracy of mental representations can actually worsen firm performance. This happens, for example, when there are many dimensions which are not very relevant to firm performance. In these cases, having a more accurate representation (=exploring more dimensions) can leave the search stranded at one of many local peaks because of the increased ruggedness.

Representations play a central role also in the evolutions of routines. Marengo (2014) uses a simple case of problem solving, the Tower of Hanoi, to show that the way in which we partition the set of states (our representations) modifies the landscape on which the search takes place and, therefore, the possible outcomes of this search (the routines). Routines are formalized in terms of if-then or condition-action rules (see above the section on classifier systems). This work also links the notion of representation to that of decomposition by introducing hierarchical representations, i.e. representations which decompose a problem into a hierarchy of sub-problems. Hierarchical representations are shown to perform better than simple representations in terms of flexibility, evolvability, and scalability.

Finally, Dosi, Marengo, Paraskevopoulou and Valente (2017) represent the *memory* of an organization, both its collective “cognitive” memory and its “operational” one, in terms of structured ensembles of “if…then…” rules (see the classic Walsh and Ungson, 1991). They investigate the performance of both types of memory in environments characterized by different degrees of complexity and non-stationarity. They show that in simple and stable environments memory does not matter, provided it satisfies some minimal requirements; in more complex and gradually changing environments, having more memory is better. However, there is some critical level of environmental instability above which forgetfulness is evolutionary superior from the point of view of long term performance.

3.2.2 Decompositions and problem solving

In explaining the inner features and boundaries of the economic organization, traditional organizational economics has focused upon the governance of transaction and contractual relations between given "technologically separable" units. No explanation is provided about the origin of these technologically separable units, and the theory is silent with regard to key facts: first, that most processes of division of labor take place *within* organizations and, second, that most of the times technologies are born in a highly-integrated fashion, and possibly undergo subsequent vertical disintegration both within and among firms. In other words, contrary to the conclusion reached in transaction cost economics, one can say that "in the origin there were organizations" and then markets develop along the lines defined by the processes of division of labor.

Following Simon (1981), Marengo and Dosi (2005)[5] tackle this issue by comparing different organizational structures with varying degrees of vertical integration in terms of patterns of division of labor and problem decomposition. Solving a problem, in this context, requires the coordination of *N* atomic "elements" or "actions" or "pieces of knowledge", which we can generically call components, each of which can assume a number of alternative states. The process can be formalized by using an NK model in which the one-bit mutation represents the case in which the problem is fully decomposed and the search process is fully decentralized: each sub-problem consists of a single component (bit). As showed by Kaufmann (1993), this is the quickest, but at the same time it is the one that leads to the local optimum whose basin of attraction contain the initial configuration. On the opposite extreme, there is the zero-decomposition strategy in which all the components (bits) are simultaneously mutated. In this case the global optimum can be reached through a slower exploration of all the possible configurations. In between there are all the other possible strategies.

The performance of different decomposition strategies depends on the existence of interdependences among the components of the problem (the value of K). Thus, separating interdependent components and then solving each sub-problem independently will prevent the very possibility of overall optimization. Another crucial issue stressed by Simon is the difference between the real structure of the interrelations between the components of the problem and the structure perceived by boundendly rational agents. The knowledge of the former is the necessary condition to design an optimal decomposition strategy corresponding to a division of labor that separates into sub-problems only the components that are independent from each other. Agents normally are bound to aim at *near-decompositions*, that is decompositions that try to put together within the same sub-problem only those components whose interdependences are perceived as "more important" for the performance of the system.

Note that different decompositions schemes entail different degrees of decentralization of the search process. The finer the decompositions, the smaller the portion of the search space which is

5 See also Marengo *et al*. 2000.

being explored and tested by market selection. Thus, there is a trade-off between finer decompositions coupled with decentralization, which make search and adaptation faster, and the size of the portions of the search space that are explored, that becoming smaller induce a decrease of the likelihood of reaching and testing optimal solutions. One of the main implications for organizational design is that the faster adaptation associated with decentralization usually implies a cost in terms of likelihood to reach the global optimum.

3.3. Capabilities and modes of governance with conflicting interests and asymmetric power distributions

Traditionally, one of the most well-known limits of the capabilities approach is the failure to take into account the existence of conflicting interests within the organization, but also the governance mechanisms adopted to handle with such conflicts and more generally the relationship between heterogeneity in the distribution of knowledge, differences in the perception of the organizational context and heterogeneity of preferences.

The problem is well known, and often the only relief comes from the acknowledgment that the alternative view, which characterizes both the neoclassic and the new-institutional theories of organization, has the opposite problem, as it gives almost exclusive emphasis on the alignment of incentives, forgetting everything that has to do with organizational problem solving. A possible merger between the two views seems unlikely: they are grounded on different foundations, they use different models and they finally end up explaining, at least in part, the same phenomena starting from completely different assumptions.

In the literature we find attempts to explicitly bring together the two approaches (see Coriat and Dosi, 1998, Dosi el at 2003, Foss and Foss, 2000) and some important advances have been made in recent works (Marengo and Pasquali, 2012; Dosi and Marengo, 2015), trying to tackle the existence of conflicts within the capability-based view of the firm.

Drawing on the ideas and the model originally developed in Dosi *et al.* (2003), Marengo, and Pasquali (2012) introduce political conflict within a typical evolutionary framework, building a model of an organization where a principal has the possibility to constrain agents' decisions and learning. The most interesting aspect of this work has to do, however, with the introduction of a further dimension of conflict, a cognitive one: conflict does not only arise from divergent interests but also from alternative representations of the world. In this context, the principal/manager can exert two forms of power, acting on the organizational structure, i.e. on the allocation of decisions, or through incentives and fiat interventions. Organizational behavior is modeled through an NK model and it is quite complex, with decisions that generate both positive and negative strong externalities on the outcomes of others' decisions. In the model, the set of policies on which a decision has to be taken is defined by a string of length *n*, $P=\{p_1, p_2, \ldots, p_n\}$. Each element of the string (a policy) can take only two values: $p_i \in \{0,1\}$. The set of possible policies is thus given by 2^n vectors. Different configurations of policies generate different performances that capture the degree of adaptation to the environment, and can be ordered according to a preference relation.

The organization is composed of a principal/manager and a number of agents that can range from 1 up to *n*. The manager has no direct control over any policy, but he delegates decisions to agents, each of which has decision rights on a subset of policies. The right of decision over a policy is assigned to one and only one agent. The structure of the organization is thus defined by the distribution of decision rights over policies and by an agenda, that is, the order in which decisions are taken.

Managers and agents have idiosyncratic orderings of policies that may not coincide with the "true" ordering that reflects the organization's performance in the environment; as a consequence, there may be a conflict generated by the existence of different preferences for the various configurations of policies. Everyone, when placed in front of two alternative configurations, will choose the one he prefers and agents will act accordingly for the policies under their control, unless the manager intervenes to change the decision that was taken through monetary incentives or overruling it by fiat. This kind of interventions, however, are expensive for the principal, and the cost is proportional to the distance between the rankings of the alternative preferred by the agent and that induced or imposed by the manager.

An alternative consists in the intervention on the organizational structure, changing the distribution of decision rights or the agenda. In the simplest case, the manager knows her preferred configuration or she knows the true ordering of alternatives.

The preferences of the agents and the principal, the organizational structure and the interventions of the manager define the organizational landscape, i.e. a description of configurations that can be reached starting from any policy configuration.

The decision takes place in a sequential manner starting from a "status quo", where a policy vector is (randomly) predetermined; the first agent then modifies the policies under her control by generating all the possible configurations and choosing the policies that generate the configuration he prefers; then the second agent does the same, and so on, following the agenda. The procedure can be terminated when all the agents have decided, or it can start again and be repeated until an equilibrium or a cycle is reached. In equilibrium, no agent finds it convenient to modify the policies under her control; while a cycle is a subset of configurations that agents keep repeating in the same order.

Different organizational structures can lead to cycles and to multiple equilibria, precisely because of externalities. In their simulations, Marengo and Pasquali (2012) test alternative organizational structures, starting from all the possible status quo. The structures that they test differ for the degree of decentralization of decisions, ranging from the two extreme cases in which a single agent decides on all the policies and the one where each agent holds control over a single policy. The main result that they get is that, in the most decentralized structures a greater number of equilibria are generated, and therefore the likelihood that any of these is the preferred by the manager is higher. At least in principle, he will be able to choose configurations that are in the basin of attraction of the equilibria that he prefers. Authority can be used to prevent cycles, that are quite common in decentralized structures.

A completely analogous result is observed in the case where the agenda is not repeated. Thus decentralization might allow the manager to manipulate organizational decisions, obtaining control without using authority, which is costly.

Obviously, since there is no correlation between the preferences of the manager and the real ordering, there is no relationship between control and organizational performance. The most interesting result derives from the analysis of the learning processes that are typical of cases in which the manager is aware of not knowing the real ordering and must learn it through a trial-and-error mechanism: given two configurations, he compares their performance in the environment and he keeps or not her preference for one of the two, depending on the feedback he received. This possibility generates a trade-off between control and learning: if the manager, through intervention, focuses on the alignment with her favorite policies, he might reduce the chances of exploration by agents, which could reduce the likelihood of learning. Finally, a further aspect that is addressed concerns the removal of externalities among agents' decisions, which can be obtained through an organizational structure that reflects the real interdependences between policies, and therefore assigns all the interdependent decisions to the same agent. In this way the costs for control would be minimized because the number of interventions required to induce a specific configuration is reduced. One of the results from the literature on near-decomposability is that by doing so a unique equilibrium can be reached and thus learning would become impossible. Consequently, if the internalization of externalities is positive in terms of costs for control, it is not for learning.

Dosi and Marengo (2015) extend the model, in particular by broadening the analysis of learning processes. As in the Marengo and Pasquali (2012) the principal can overrule agent's decision either by *veto*, inducing a return to status quo, or by *fiat*, replacing the agents' decisions with her own preferred configuration. The authors study three different learning processes based on simple trial-and-error adaptation rules. In the first two the principal and the agent, respectively, adapt their rankings to the real one, while in the third the agents adapts their ranking to that of the manager. During the processes of learning of the real ranking, once an equilibrium or an outcome (when the agenda is not repeated) is reached, the agents and the principle receive a feedback from the environment and act consequently. In particular, if the new configuration is better than the previous one, in terms of the real ranking, with a certain probability those who preferred the old one change their ordering. In the third learning process, instead, when the principle intervenes, by *fiat* or *veto*, the agent observes the imposed configuration and, deducing that the manager prefers that configuration to her configuration, with a certain probability she shows docility by changing her ordering and adapting it to that of the manager.

The simulations are run on the same set of organizational structures studied by Marengo and Pasquali (2012). A first, interesting, result has to do with the existence of a non-monotonic relation between strength of the interventions and organizational performance. When the control is zero, there is no coordination and the performance is low. With more incisive interventions, both the coordination and the performance increase, but when the exercise of authority becomes too deep, and full control is reached, the organizational landscape becomes similar to that of the principle, which is single-peaked by definition, and the performance decreases.

In this article one finds interesting results also about the relation between the level of competence – defined as the correlation between the agent/principal's ranking and the real one – and the allocation of decision power. When the principal is fully competent, then she must exercise the maximum control, overruling all the decisions taken by the agents that do not reflect her preferences. Instead, when the principal is less than fully competent, but she is still more competent than the agents, the best result is reached when decision rights are decentralized and an intermediate level of authority is used. With regard to agents' competence, in general, delegating the decisions to the more competent agents it is not a good idea, unless there is a single fully competent agent. As noted by the authors, this result, that at a first glance might appear counterintuitive, depends on the high complexity of the environment that characterizes these models, in which very different performances might be associated to very similar configurations of policies. In such contexts, the paths that lead to the optimal configuration might be very tortuous.,

As already shown by Marengo and Pasquali (2012), the capacity of adaptation to the environment increases with the number of equilibria, and consequently with the degree of decentralization of the organizational structure, with a central role being played by managerial interventions in preventing cycles. In the case in which only the principal can learn and with highly decentralized decisions, the exercise of soft authority induces a lower degree of adaptation, because cycles become more frequent. But strong interventions reduce the exploration possibilities of the agents, reducing the overall level of adaption. With a coarser partition, as in the case in which the decision rights are equally split between two agents, this trade-off disappears and learning increases steadily with the use of *veto* or *fiat.*

When also agents have the possibility to learn from the environment, learning by the principal improves only in less partitioned organizational structures. This is due to the negative effect on coordination induced by multiple adaption in a highly-decentralized structure. Finally, with regard to docility, the adaptation of agents to the rankings of the principal, reduces the possibility of exploration, and has a negative effect on learning by the principal, and consequently on the performance of the organization.

4. Conclusions

We presented different formal instruments and models which try to account for organizational capabilities, with a specific attention to problem-solving activity, in terms of sequences of procedures nested into specific organizational architectures characterized by division of cognitive labor and decision rights. According to these models, there exist strong interdependencies among the various activities carried out by an organization, and the boundedly rational agents involved in the search processes are only partially able to understand them.

A crucial implication of the partial “opaqueness” of the mappings between actions and outcome is that firms might not be able to reach the global optimum. In fact, in most models we reviewed, firms have to be satisfied with local optima, relying on partial cognitive representations of the environment and imperfect decompositions of the problem-solving space.

In this perspective, heterogeneity in firm performance originates from the different ways that organizations adopt to search over the landscape. The results showed that problem-solving efficiency and learning depend on: how hierarchy is distributed by the organizational architecture; the cognitive maps and representations that are available to agents and organizations; how different degrees of centralization are associated to the possible trade-offs between exploration and exploitation. In a nutshell, the models allow for differential heterogeneity in firm performance be driven by different degrees of organizational capabilities, which stem from the ensemble of routinized patterns of actions.

One can also conceive *changes* in organizational capabilities, partly as a result of deliberate search and learning, which in turn depend upon the coordination among many interdependent subunits and their cognitive representations of the world. This is in line with the ongoing stream of research on dynamic capabilities (Helfat *et al.*, 2007; Teece *et al.*, 1997; Winter, 2003) that precisely addresses the criteria and processes by which capabilities evolve at least partly steered by the effort of strategic management.

References

Benner, M. J., and Tushman, M. L. (2003). Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review*, *28*(2), 238–256.

Bocanet, A., & Ponsiglione, C. (2012). Balancing exploration and exploitation in complex environments. *VINE*, *42*(1), 15–35

Bower, J. L., and Christensen, C. M. (1995). Disruptive Technologies: Catching the Wave. *Harvard Business Review*, *73*(1), 43–53.

Coriat, B., G. Dosi. 1998. “Learning how to govern and learning how to solve problems: On the co-evolution of competences, conflicts and organizational routines”. A. D. Chandler Jr., P. Hagström,O. Sölvell, (eds.) *The Dynamic Firm: The Role of Technology, Strategy, Organization, and Regions*. Oxford University Press, Oxford, UK, 103–133.

Csaszar, F. A., and Levinthal, D. A. (2015). Mental representation and the discovery of new strategies. Strategic Management Journal.

Csaszar, F. A., & Siggelkow, N. (2010). How Much to Copy? Determinants of Effective Imitation Breadth. *Organization Science*, *21*(3), 661–676.

Cyert, R. M. and J. G. March (1963), *A Behavioral Theory of the Firm*. Cambridge: Blackwell.

Dosi, G., D. A. Levinthal, L. Marengo (2003). “Bridging contested terrain: Linking incentive-based and learning perspectives on organizational evolution”, *Industrial and Corporate Change* 12(2) 413–436.

Dosi, G., M. Faillo and L. Marengo (2008), “Problem-Solving and Governance in the Capability-based view of the Firm: The Roles and Theoretical Representations of Organizational Routines”, in M. Becker (ed.), Handbook of Organizational Routines. Cheltenham, UK and Northampton, MA: Edward Elgar.

Dosi, G., R. Nelson and S. G. Winter (eds.) (2000), The Nature and Dynamic of Organizational Capabilities. Oxford and New York: Oxford University Press.

Dosi, G., M. Faillo, L. Marengo, and D. Moschella (2011). “Toward Formal Representations of Search Processes and Routines in Organizational Problem Solving: An Assessment of the State-of-the-Art”, *Seoul Journal of Economics*, 24(3), 247-286.

Dosi G., and Marengo, L. (2015). The dynamics of organizational structures and performances under diverging distributions of knowledge and different power structures. *Journal of Institutional Economics*, *11*(3), 535–559.

Dosi, G., L. Marengo, E. Paraskevopoulou and M. Valente (2017), *A model of cognitive and operational memory of organizations in changing worlds,* Cambridge Journal of Economics. Forthcoming.

Dosi, G. and R. Nelson (2010), “Technical Change and Industrial Dynamics as Evolutionary Processes”, in B. Hall and N. Rosenberg (eds.), Handbook of Innovation.

Ethiraj, S. K., and Levinthal, D. (2004). Bounded Rationality and the Search for Organizational Architecture: An Evolutionary Perspective on the Design of Organizations and Their Evolvability. *Administrative Science Quarterly, 49*(3), 404 – 437.

Ethiraj, S. K., Levinthal, D., & Roy, R. R. (2008). The Dual Role of Modularity: Innovation and Imitation. *Management Science*, *54*(5), 939–955.

Fang, C., Lee, J., and Schilling, M. a. (2010). Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning. *Organization Science*, *21*(3), 625–642

Foss, K., and N. Foss. (2000). “Competence and governance perspectives: How much do they differ? And how does it matter?” in N. J. Foss, V. Mahnke, eds. Competence, Governance, and Entrepreneurship (eds.), *Advances in Economic Strategy Research*. Oxford University Press, Oxford, UK, 55–80.

Frenken, K., 2006. *Innovation, Evolution and Complexity Theory*. Edward Elgar, Northampton.

Gary, M. S. and Wood, R. E. (2011). Mental models, decision rules, and performance heterogeneity. *Strategic Management Journal*, 32(6), 569–594.

Gavetti, G. (2005), “Cognition and Hierarchy: Rethinking the Microfoundations of Capabilities’ Development”, *Organization Science*, 16, 599-617.

Gavetti, G. and D. Levinthal (2000), “Looking Forward and Looking Backward: Cognitive and Experimental Search”, *Administrative Science Quarterly*, 45, 113–137.

Gavetti, G., Levinthal, D. A., and Rivkin, J. W. (2005). Strategy making in novel and complex worlds: the power of analogy. *Strategic Management Journal*, *26*(8), 691-712.

Helfat, C., Finkelstein, S., Mitchell, W., Peteraf, M., Sing h, H., Teece, D. and Winter, S., (2007),

Dynamic Capabilities: Understanding Strategic Change in Organisations, Blackwell Publishing, Malden.

Holland, J.H. (1975). *Adaptation in natural and artificial systems*. Ann Arbor: University of Michigan Press.

Holland, J.H. (1986). “Escaping brittleness: the possibilities of general rule-based systems” in R.S. Michalski, J.H. Carbonell, T.M. Mitchell (eds.), *Machine learning II*. Los Altos, CA, Morgan Kaufmann, pp. 593-623.

Holland, J.H., H.J. Holyoak, R.E. Nisbett, P.R. Thagard (1986). *Induction: Processes of Inference, Learning and Discovery.* Cambridge MA: MIT Press.

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

Knudsen, T. and D. A. Levinthal (2007), “Two Faces of Search: Alternative Generation and Alternative Evaluation”, *Organization Science*, 18, 39-45.

Levinthal, D. (1997), “Adaptation on Rugged Landscapes”, *Management Science*, 43 (7): 934-950.

Luce, R. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York.

March, J. G. (1991). Exploration and Exploitation in Organizational Learning. *Organization Science* 2(1): 71-87.

March, J. G. and H. A. Simon (1958), *Organizations*, New York: Wiley.

March, J. G. (2005). Parochialism in the Evolution of a Research Community: The Case of Organization Studies. *Management and Organization Review*, *1*(1), 5–22.

Marengo, L. (1992), “Coordination and Organizational Learning in the Firm”, *Journal of Evolutionary Economics*, 2, 313-326.

Marengo, L. (1996), “Structure, Competence and Learning in an Adaptive Model of the Firm”, in G. Dosi and F. Malerba (eds.), *Organization and Strategy in the Evolution of the Enterprise*. London: MacMillan.

Marengo, L. (2014). Representation, search, and the evolution of routines in problem solving. *Industrial and Corporate Change*, *24(5), 951-980.*

Marengo, L. and G. Dosi (2005), “Division of Labor, Organizational Coordination and Market Mechanisms in Collective Problem-solving", *Journal of Economic Behavior and Organization,* 58, 303-326.

Marengo, L., G. Dosi, P. Legrenzi and C. Pasquali (2000), “The Structure of Problem-Solving

Knowledge and the Structure of Organizations”, *Industrial and Corporate Chang*e, 9, 757-788.

Marengo, L., and Pasquali, C. (2012). How to Get What You Want When You Do Not Know What You Want: A Model of Incentives, Organizational Structure, and Learning. *Organization Science*, *23*(5), 1298–1310.

Martignoni, D., Menon, A., and Siggelkow, N. (2016). Consequences of misspecified mental models: Contrasting effects and the role of cognitive fit. *Strategic Management Journal*.

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

O’Reilly, C. A., and Tushman, M. L. (2004). The Ambidextrous Organization. *Harvard Business Review*. 82(4) 74–81.

O’Reilly, C. A., & Tushman, M. L. (2008). Ambidexterity as a dynamic capability. *Research in Organizational Behavior*, *28*, 185–206.

O’Reilly, C. A., & Tushman, M. L. (in press). Ambidexterity: Organizational Ambidexterity: Past, Present and Future. *Academy of Management Perspectives.*

Pentland, B. T. and H. H. Reuter (1994), “Organizational Routines as Grammars of Action”, *Administrative Science Quarterly*, 39, 484-510.

Posen, H. E., & Levinthal, D. A. (2012). Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments. *Management Science*, *58*(3), 587–601.

Rivkin, J. W. and N. Siggelkow (2002), “Organizational Sticking Points on NK Landscapes”, *Complexity*, 7, 31-43.

Rivkin, J. W. and N. Siggelkow (2003), “Balancing Search and Stability: Interdependencies Among Elements of Organizational Design”, *Management Science*, 49, 290-311.

Siggelkow, N., and Levinthal, D. a. (2003). Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation. *Organization Science*, *14*(6), 650–669

Siggelkow, N., and Levinthal, D. A. (2005). Escaping real (non-benign) competency traps linking the dynamics of organizational structure to the dynamics of search. *Strategic Organization*, *3*(1), 85–115

Siggelkow, N., and Rivkin, J. W. (2005). Speed and Search: Designing Organizations for Turbulence and Complexity. *Organization Science*, *16*(2), 101–122.

Siggelkow, N. and J. W. Rivkin (2006), “When Exploration Backfires: Unintended Consequences of Multilevel Organizational Search”, *Academy of Management Journal*, 49, 779-795.

Siggelkow, N. (2011), " Firms as Systems of Interdependent Choices", Journal of Management

Studies, 48 (5): 1467-6486.

Simon, H. A. (1981), *The Science of Artificial*, 2nd Ed.. Cambridge, MA: MIT Press.

Teece, D. J., Pisano, G.P. and Shuen, A. (1997) "Dynamic Capabilities and Strategic Management." *Strategic Management Journal*, 18(7): 509-33.

Tushman, Michael L. and O'Reilly, Charles A. (1996). "The ambidextrous organization: managing evolutionary and revolutionary change". *California Management Review*, 38: 1-23

Tripsas, M. & Gavetti, G. (2000). Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal*, 21(10/11), 1147-1161.

Walsh, J. P. and G. R. Ungson (1991), "Organizational Memory", *The Academy of Management Review*, 16, 57-91.

Watts, D. J. (1999). *Small Worlds: The Dynamics of Networks Between Order and Randomness*.

Winter, S. (2003) "Understanding Dynamic Capabilities." *Strategic Management Journal* 24(10): 991-995.

Wright, S. (1932). The roles of mutations, inbreeding, crossbreeding and selection in evolution. *Proceedings of the 11th International Congress of Genetics*.